Forward Return Expectations*

Mihir Gandhi, Niels Joachim Gormsen, and Eben Lazarus

April 2026

Abstract

We measure investors' short- and long-term stock-return expectations using both options and survey data. These expectations at different horizons reveal what investors think their own short-term expectations will be in the future, or forward return expectations. Across data sources, investors' forward return expectations are countercyclical. This countercyclical variation is, however, excessive, as the changes in forward return expectations are too large relative to subsequent realized short-term expectations. The excess variation in forward return expectations can account for a substantial share of the stock-price declines during the global financial crisis and the Covid-19 crisis.

Keywords: asset pricing, stock returns, subjective return expectations

JEL classifications: G10, G12, G40

*A previous version of this paper was circulated under the title "Excess Persistence in Return Expectations." We thank Anna Pavlova and three anonymous referees for very helpful feedback, as well as Nick Barberis, Francesca Bastianello, Pedro Bordalo, Nina Boyarchenko, John Campbell, Fousseni Chabi-Yo, Mikhail Chernov, Nicola Gennaioli, Stefano Giglio, Ben Hébert, Ralph Koijen, Martin Lettau, Ian Martin, Peter Maxted, Tyler Muir, Stefan Nagel, Emi Nakamura, Stijn Van Nieuwerburgh, Dino Palazzo, Frans de Roon, Andrei Shleifer, David Sraer, Jón Steinsson, David Thesmar, Chen Wang, and seminar and conference participants at UC Berkeley, Oxford, Warwick Business School, Chicago Booth, MIT Sloan, BI Oslo, Copenhagen Business School, UW Foster, Yale SOM, LSE, Brown, the AFA Annual Meeting, ASU Sonoran Winter Finance Conference, ITAM Finance Conference, Junior Finance Conference on Valuations at USC, Tilburg Finance Summit, and Virtual Derivatives Workshop for helpful comments and discussions. We are grateful to the Fama-Miller Center, Fama Faculty Fellowship, Asness Junior Faculty Fellowship, Fischer Black Doctoral Fellowship, the Danish Finance Institute (DFI), and the Danish National Research Foundation (DNRF) through the Center for Big Data in Finance (grant no. DNRF167) and the DNRF Chair program (grant no. DNRF199) for financial support. Gandhi is at Copenhagen Business School and the Danish Finance Institute (mihir.a.gandhi@gmail.com); Gormsen is at Copenhagen Business School, the University of Chicago Booth School of Business, NBER, and the Danish Finance Institute (niels.gormsen@chicagobooth.edu); and Lazarus is at the Haas School of Business, University of California, Berkeley (lazarus@berkeley.edu).

1 Introduction

One of the organizing facts in asset pricing is that expected stock returns vary over time. A key question is how investors perceive this variation ex ante. One extensively studied aspect of this question is how investors perceive the contemporaneous level of the equity premium. Another aspect — and the focus of this paper — is what investors expect the equity premium to be in the future, or their *forward return expectations*. These forward return expectations are important for stock price determination, as prices depend on all future return expectations and growth expectations. Our analysis reveals predictable forecast errors in forward return expectations, which can account for a large share of stock-price movements during crashes.

We study forward return expectations using both option prices and survey data. First, using a global sample of equity index options, we extract proxies for the equity premium at multiple horizons as perceived by sophisticated investors. With both long- and short-term expectations in hand, we construct forward expectations of returns based on the difference between the two. New theoretical results show that we can identify time-series changes in these expectations, which are our main focus, under relatively weak conditions. We complement the options data with return expectations from surveys of professional forecasters, CFOs, and retail investors. These surveys provide return expectations across different time horizons, enabling us to extract forward expectations in a more model-free setting.

The data reveal that forward return expectations are countercyclical among all investor types: sophisticated, professional, and retail investors. This finding aligns with estimates of the objective variation in expected returns, which are also known to be countercyclical (Campbell and Shiller 1988). The robust countercyclical variation in forward return expectations stands in contrast to the evidence on spot return expectations, for which the evidence on cyclicality in past literature is mixed.

The countercyclical variation is, however, excessively strong. In bad times, investors believe that expected returns will be substantially elevated in the future, but these elevated expectations are too high relative to investors' own subsequent beliefs about short-term returns. (The reverse applies in good times.) During the 2008 financial crisis and the 2020 Covid crisis, for instance, investors believed that equity premia would be substantially elevated in the future; ex post, however, their realized short-term expectations were only mildly elevated, and they reverted to their usual levels relatively quickly. This pattern applies consistently over time, and across all the settings we study, in both options and survey data.

The excess variation in forward expectations represents a new source of excess volatility in stock prices. If investors did not overestimate forward return expectations during bad times, prices would fall less in crises and would generally exhibit more modest cyclical swings. We show in particular that predictable mistakes in forward expectations can account for a significant share of the decrease in prices observed in major crisis episodes. Our lower-bound estimates explain about 20–25% of the index-level price declines in the global financial crisis and the early-2020 Covid crash, while our largest estimates imply that forward-expectation errors can account for the majority of the price drops in these periods.

Our findings on forward expectations can also help explain or reconcile a number of facts documented in recent work, including disagreements in previous research on survey expectations of returns, stylized facts about the equity term structure, and more speculatively, investors' inelastic demand for equities in response to a given change in prices. We explore these implications further below and in Section 5.

Framework and Stylized Facts. To better understand our analysis, we begin by defining the following notation for the expected log return on the market over n periods:

$$\mu_t^{(n)} = \mathbb{E}_t[r_{t,t+n}].$$

We refer to these expected returns as *spot rates*. Later in the paper, we net out the risk-free rate and consider risk premia in addition to expected returns; for now, we consider expected returns for notational simplicity.[1]

[1] To abstract from the trend and cyclical behavior of the risk-free rate, we consider forward risk premia to be our main measure of forward return expectations in the analysis below.

If at time t we observe the n-period and $(n+1)$-period spot rates, we can back out the expected returns between these two periods, which we call the *forward return expectation*, or *forward rate* for short:

$$f_t^{(n)} = \mu_t^{(n+1)} - \mu_t^{(n)} = \mathbb{E}_t\left[\mu_{t+n}^{(1)}\right].$$

After n periods, we can compare this forward rate to the realized one-period spot rate, $\mu_{t+n}^{(1)}$. Under the law of iterated expectations, the forward rate should be an unbiased predictor of the realized spot rate, which means that *forecast errors*,

$$\varepsilon_{t+n} = \mu_{t+n}^{(1)} - f_t^{(n)},$$

should have an expected value of zero and be unpredictable by any time-t information.

The forward rates and forecast errors can be estimated straightforwardly with survey data, as long as the given survey provides return expectations at multiple horizons. This is the case for three well-known surveys in the finance literature: the Livingston survey of professional forecasters, the Duke–Fed CFO survey of financial executives, and the Vanguard survey of retail investors. In the Livingston survey, for instance, we observe expected returns at the 6-month and 12-month horizon. These allow us to calculate a 6-month, 6-month forward rate, namely today's expectation of the 6-month spot rate, 6 months from now. By comparing this forward rate to the realized 6-month spot rate, as observed 6 months later, we can calculate forecast errors and evaluate how well forward rates predict future spot rates.

These forecast errors for return expectations can also be estimated from option prices under relatively mild conditions. We start by considering a simple special case, in which an investor has log utility over the market return. For this investor, we can directly estimate expected returns at different horizons from option prices, allowing us to estimate forward rates and forecast errors over time. But an important contribution of our paper is to show that we can extract the forecast errors under much more general conditions than log utility: for any investor, we can identify expected forecast errors up to a small risk premium term that can be managed empirically or through theory. Our forecast-error estimates are thus useful for a wide range of specifications of preferences and the data generating process. The advantage of the option-based data is that we observe forward rates and forecast errors at many different horizons, in a relatively long sample that spans many countries.

We provide a consistent set of stylized facts about forward rates and forecast errors across investors, which hold both for forward risk premia and forward expected returns. Forward rates are decent predictors of future realized spot rates. Across data sources, a 1-percentage-point increase in ex ante forward rates predicts around a 0.7-point increase in future realized spot rates, with R^2 values around 0.5. But forward rates consistently overshoot: when forward rates are high, subsequent spot rates are lower than expected by investors ex ante, generating predictable forecast errors. In predictive regressions, these forecast errors are statistically significant and economically meaningful. There is accordingly excess cyclicality in return expectations: in bad times, investors' beliefs about future discount rates are elevated by significantly more than justified by their subsequent beliefs.[2]

To illustrate this excess cyclicality, Figure 1 shows forward expected returns and realized spot rates from the option-based measure during three crises: the 1998 Russian debt crisis, the 2008 global financial crisis, and the 2020 Covid-19 crisis. The blue circles show one-month forward rates maturing $n = 0$, 1, 2, 3, 4, and 5 months from the first plotted date t (which is soon after each crisis onset). The red triangles show ex post realized one-month spot rates as of months t, $t + 1$, . . ., $t + 5$. In all cases, the spot expected return (the first point) increases substantially as of the crisis onset. Forward rates also increase, but the forward curve is strongly downward sloping, suggesting that investors expect significant mean reversion of spot rates in subsequent months; the slope of the term structure of return expectations is procyclical. Going forward, the spot rate indeed decreases substantially as the crisis recedes.

In all cases, however, forward rates are too high at the peak of the crisis relative to future realized spot rates. While investors understood that expected returns would decrease in the

[2] This cyclicality is excessive relative to their *own* subsequent beliefs: we are testing the internal consistency of expectations over time, and not just relative to an estimate of cyclicality in objective expected returns.

Figure 1
U.S. Forward Expected Returns and Realized Spot Rates in Crises

This figure plots forward expected returns $f_t^{(0)}$, $f_t^{(1)}$, $\cdots$, $f_t^{(5)}$ (blue) and the corresponding ex post spot expected returns $\mu_t^{(1)}$, $\mu_{t+1}^{(1)}$, $\cdots$, $\mu_{t+5}^{(1)}$ (red) in the U.S. sample. For the Russian debt crisis, forward rates are as of 08/1998 and realized spot rates are from 08/1998 to 01/1999. For the financial crisis, forward rates are as of 11/2008 and realized spot rates are from 11/2008 to 04/2009. For the Covid-19 recession, forward rates are as of 03/2020 and realized spot rates are from 03/2020 to 08/2020.

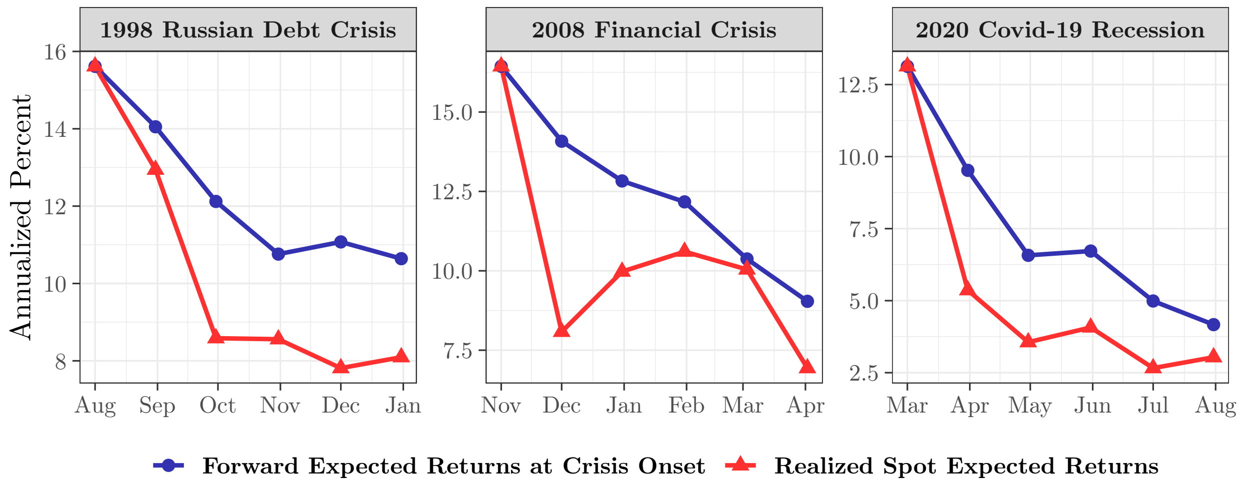

future, they appear to have underestimated the speed of mean reversion. One interpretation is that the rebound following each crisis was driven by unexpected news (e.g., a series of policy shocks). Another interpretation is that investors overreact to crises: they think that spot rates will be elevated by more and for longer than what one should rationally expect.

We find consistent evidence of this excess cyclicality across all data sources and investor groups. The forward rates also share a common countercyclical component and comove with standard indicators such as the earnings yield, implied volatility, and a recession indicator, as well as many other indicators. Forecast errors also share cyclical variation: in bad times, all investor groups consistently overestimate how high expected returns and risk premia will be in the future. And option-based forward rates strongly predict forecast errors in the surveys. Patterns in forward expectations are thus shared across a broad range of investors.

Implications. Our results have implications for multiple literatures in asset pricing.

The debate on the cyclicality of subjective expectations. A recent literature studies the cyclical dynamics of subjective expected returns and risk premia. Greenwood and Shleifer (2014) document procyclical variation; Nagel and Xu (2023) document generally acyclical variation; and Martin (2017) and Dahlquist and Ibert (2024, 2026) document countercyclical variation, all studying different measures and investor types. We find that the degree of cyclicality depends on the horizon: short-run expectations differ across investor groups, but our focus on forward rates eliminates such disagreement. In all cases, longer-run forward rates are robustly and excessively countercyclical. We therefore provide a unified body of evidence on the dynamics of long-run return expectations, which are key for determining prices.

Table 1 summarizes how our contribution fits in with and helps unify previous findings across different horizons and investor groups. The second ("Forward") column illustrates that across all investor types, we find that forward expectations consistently move excessively relative to the predictable component of future expectations. This contrasts with the inconsistent evidence on the behavior of short-run expectations, with different cyclical properties across investor groups.[3] Put together, we find strong evidence that long-run expectations vary countercyclically and co-move positively with objective risk premia.

A new source of excess volatility in prices. Our results suggest that when crises hit, prices drop by an excessive amount because investors believe future risk premia will be excessively elevated. Without such expectation errors, price fluctuations would be more modest: in our baseline option-based estimation, the expectation errors can account for most of the stock price declines during crises, and lower-bound estimates based on the survey data also account for a significant share of these crashes. Our results thus provide a new interpretation of stock-price variation, and crashes in particular, in recent decades.[4]

Stylized facts about the equity term structure. The literature on the equity term structure studies prices and returns for dividend claims with different maturities, finding that the

[3]For the past literature, the results for sophisticated and professional forecasters are based on equity premium expectations, while those for CFOs and retail investors are based on equity return expectations. This is one proposed reconciliation for the different cyclicality results, although Martin (2017) and Dahlquist and Ibert (2024) document that their results hold for expected returns as well.

[4]We note that our results *account* for a share of excess volatility without necessarily *explaining* its fundamental cause. Forecast errors also account for somewhat less of the *unconditional* variation in valuations.

equity term premium — the relative return for long- versus short-maturity claims — is negative on average (Binsbergen and Koijen 2017, Binsbergen, Brandt, and Koijen 2012) and countercyclical (Gormsen 2021). We show that our forecast-error results can help rationalize such findings: if the ex ante equity term structure is flat, the forecast errors we document result in a realized-return term structure that is downward-sloping on average and countercyclical.

Our results also build on and relate to a range of other work, as reviewed in Appendix B.

2 Methodology

We begin with a theoretical analysis of the term structure of expected log equity returns and risk premia. We first set up notation for spot rates, forward rates, and forecast errors. We then present our main theoretical results on forecast-error identification from option prices. We end with a discussion of identification of forecast errors from surveys.

2.1 Notation

We start by generalizing our notation slightly relative to the introduction. Continue to define the spot rate as the time-t subjective log return expectation between periods t and $t+n$:

$$\mu_t^{(n)} = \mathbb{E}_t[r_{t,t+n}], \tag{1}$$

where $r_{t,t+n} = \ln(R_{t,t+n})$ is the log return on the market. For any horizon n, the expected return in (1) can be written as the one-period spot rate plus a series of one-period forward rates:

$$\mu_t^{(n)} = \mu_t^{(1)} + \sum_{i=1}^{n-1} f_t^{(i,1)},$$

where forward rates are now defined as

$$f_t^{(n,m)} = \mu_t^{(n+m)} - \mu_t^{(n)} = \mathbb{E}_t\left[\mu_{t+n}^{(m)}\right]. \tag{2}$$

We refer to $f_t^{(n,m)}$ as the $n \times m$ forward rate.

We define forecast errors as the difference between realized spot rates and ex ante forward rates,

$$\varepsilon_{t+n}^{(m)} = \mu_{t+n}^{(m)} - f_t^{(n,m)}. \tag{3}$$

Under the law of iterated expectations, the time-t conditional expectation of forecast errors is zero:

$$\mathbb{E}_t\left[\varepsilon_{t+n}^{(m)}\right] = 0.$$

In much of the analysis, we study expectations about future risk premia (net of the n-period log risk-free rate $r_{t,t+n}^f$), as opposed to expected returns (gross of the risk-free rate, as above). When considering risk premia, we use the notation $\tilde{\mu}$, $\tilde{f}$, and $\tilde{\varepsilon}$ to refer to spot rates, forward rates, and forecast errors, respectively. More specifically, in place of (1),

$$\tilde{\mu}_t^{(n)} = \mu_t^{(n)} - r_{t,t+n}^f = \mathbb{E}_t\left[r_{t,t+n} - r_{t,t+n}^f\right].$$

Then $\tilde{f}$ and $\tilde{\varepsilon}$ are defined as in (2)–(3), with $\tilde{\mu}$ in place of μ. In line with much of the past literature, we consider forward risk premia to be our main measures of forward return expectations, and we present these first in our tables and show them in the remaining figures.

For forward risk premia, the law of iterated expectations implies $\mathbb{E}_t\left[\tilde{\varepsilon}_{t+n}^{(m)}\right] = \theta_t^{(n,m)}$, where $\theta_t^{(n,m)}$ is the term premium on the $n+m$–period bond in excess of the n-period bond. The expected value of the forecast errors for risk premia is therefore zero only when the pure expectations hypothesis for interest rates holds. Tests for forecast-error predictability, meanwhile, hold under the slightly weaker expectations hypothesis in which $\theta_t^{(n,m)}$ is constant.[5]

Finally, we write $M_{t,t+n} = M_{t,t+1}M_{t+1,t+2}\cdots M_{t+n-1,t+n}$ for the n-period stochastic discount factor (SDF).

[5]This assumption is typical in the literature on derivatives term structures; see Giglio and Kelly (2018).

2.2 Identification from Option Prices

This section describes how to estimate spot rates, forward rates, and forecast errors (all of which depend on physical expectations) from option prices. As a starting point, using that $\mathbb{E}_t[M_{t,t+n}R_{t,t+n}] = 1$ under the law of one price, we can write expected log returns as

$$\begin{aligned} \mathbb{E}_t[r_{t,t+n}] &= \mathbb{E}_t[r_{t,t+n}]\,\mathbb{E}_t[M_{t,t+n}R_{t,t+n}] \\ &= \underbrace{\mathbb{E}_t[M_{t,t+n}R_{t,t+n}r_{t,t+n}]}_{\mathscr{L}_t^{(n)}} - \underbrace{\operatorname{cov}_t(M_{t,t+n}R_{t,t+n}, r_{t,t+n})}_{\mathcal{C}_t^{(n)}}. \end{aligned} \tag{4}$$

This follows Gao and Martin (2021), and we build on their analysis. They show that the first term on the right side of (4) is observable from option prices, and they label this term the LVIX.[6] We denote this term by $\mathscr{L}_t^{(n)}$, as in (4). The covariance term, $\mathcal{C}_t^{(n)}$, is an unobserved risk adjustment. The size and sign of $\mathcal{C}_t^{(n)}$ can be controlled by theory, as discussed below.

To build intuition, we first make the simplifying assumption that $M_{t,t+n} = 1/R_{t,t+n}$. This stochastic discount factor would, for instance, arise if the representative investor is fully invested in the stock market and has log utility over terminal wealth. In this case, the covariance term in (4) is equal to zero, and the LVIX directly identifies expected excess returns at different horizons. Given the identity in (4) and the definitions in Section 2.1, we can thus calculate spot rates, forward rates, and forecast errors from the data as follows.

Proposition 1 (Log Utility Identification). *Assuming that $M_{t,t+n} = 1/R_{t,t+n}$, spots, forwards, and forecast errors are given, respectively, by:*

$$\begin{aligned} \mu_t^{(n)} &= \mathscr{L}_t^{(n)} \\ f_t^{(n,m)} &= \mathscr{L}_t^{(n+m)} - \mathscr{L}_t^{(n)} \\ \varepsilon_{t+n}^{(m)} &= \mathscr{L}_{t+n}^{(m)} - \mathscr{L}_t^{(n+m)} + \mathscr{L}_t^{(n)}. \end{aligned}$$

[6] We use a slightly different definition of LVIX, as theirs is in terms of excess returns and we start with total returns. Note that the relevant risk-free rate is the rate from t to $t+n$.

Proofs for all theoretical results can be found in Appendix A. Proposition 1 follows straightforwardly from the fact that $M_{t,t+n}R_{t,t+n} = 1$ under log utility. Thus, with this assumption, we can directly identify forecast errors from options and test whether expectations are intertemporally consistent (i.e., whether $\mathbb{E}[\varepsilon_{t+n}^{(m)}] = 0$ and $\mathbb{E}[Z_t\, \varepsilon_{t+n}^{(m)}] = 0$ for Z_t observable as of t). This result does not require the existence of a representative agent: the LVIX-based estimates reflect the expectations of any unconstrained log investor who is content to hold the market portfolio. These expectations can, for instance, be thought of as those of Mr. Market in the heterogeneous-agent log utility model of Martin and Papadimitriou (2022).

Going beyond log utility, estimated spot and forward rates are contaminated by the covariance terms in (4). When considering spot rates alone, it may be reasonable to assume that $\mathcal{C}_t^{(n)} = \text{cov}_t(M_{t,t+n}R_{t,t+n}, r_{t,t+n}) \leq 0$, so that the LVIX provides a lower bound for $\mu_t^{(n)}$; this is the tack taken by Gao and Martin (2021). But when considering forward rates, $f_t^{(n,m)} = \mathscr{L}_t^{(n+m)} - \mathscr{L}_t^{(n)} - (\mathcal{C}_t^{(n+m)} - \mathcal{C}_t^{(n)})$, it is unclear whether $\mathcal{C}_t^{(n+m)} \leq \mathcal{C}_t^{(n)}$ or vice versa.

Our main innovation, however, is to consider *forecast errors*, for which we show that the unobserved covariance terms largely cancel in expectation. We present two versions of this result. First, to continue building intuition, Proposition 2 considers identification in a log-normal world given general $M_{t,t+n}$. Proposition 3 then provides a fully general identification result, which shows that the main insights from Proposition 2 carry through.

Before presenting these results, we define our empirical proxy for forecast errors:

$$\widehat{\varepsilon}_{t+n}^{(m)} = \mathscr{L}_{t+n}^{(m)} - \mathscr{L}_t^{(n+m)} + \mathscr{L}_t^{(n)},$$

which are the forecast errors obtained under log utility. We also define $MR_{t,t+n} = M_{t,t+n}R_{t,t+n}$. We now show that expected forecast errors can be measured up to a single covariance term.

Proposition 2 (Log-Normal Identification). *Assume a general SDF $M_{t,t+n}$, and assume that $M_{t,t+n}$ and $R_{t,t+n}$ are jointly log-normal. Given a true-forecast-error process $\varepsilon^{(m)}_{t+n}$, the expected value of the forecast-error proxy satisfies*

$$\mathbb{E}_t\left[\widehat{\varepsilon}^{(m)}_{t+n}\right] = \mathbb{E}_t\left[\varepsilon^{(m)}_{t+n}\right] - \operatorname{cov}_t(MR_{t,t+n}, \mathbb{E}_{t+n}[r_{t+n,t+n+m}])\,. \tag{5}$$

Proposition 2 shows that in a log-normal world, our forecast-error estimate is equal in expectation to the true forecast error minus an unobserved covariance term related to the pricing of shocks to expected returns (or discount-rate risk). Note that when considering spot rates, the risk adjustment term $\mathcal{C}^{(n)}_t = \operatorname{cov}_t(M_{t,t+n}R_{t,t+n}, r_{t,t+n})$ depends on a covariance with *realized* returns $r_{t,t+n}$. By contrast, when considering forecast errors, the relevant risk adjustment depends on a covariance with *expected* returns $\mathbb{E}_{t+n}[r_{t+n,t+n+m}]$. By replacing (highly volatile) realized returns with (much less volatile) expected returns, the covariance in (5) is likely to be substantially smaller than the covariance in (4). We discuss how one might quantify this covariance in greater detail in Appendix E.

The following proposition shows that the above intuition carries through to more general distributions, with $\mathbb{E}_{t+n}[r_{t+n,t+n+m}]$ replaced by a closely related LVIX-based proxy.

Proposition 3 (Generalized Identification). *For any SDF $M_{t,t+n}$ and any data-generating process for which the relevant expectations exist, the expected value of the forecast-error proxy satisfies*

$$\mathbb{E}_t\left[\widehat{\varepsilon}^{(m)}_{t+n}\right] = \mathbb{E}_t\left[\varepsilon^{(m)}_{t+n}\right] - \operatorname{cov}_t\left(MR_{t,t+n}, \mathscr{L}^{(m)}_{t+n}\right)\,. \tag{6}$$

The appearance of the LVIX $\mathscr{L}^{(m)}_{t+n}$ in the risk adjustment term in (6) has one benefit relative to (5): the LVIX is empirically observable. We can thus quantify the degree to which $\mathscr{L}^{(m)}_{t+n}$ is less volatile than the realized return $r_{t,t+n}$ in (4). Under our main specification (discussed further in Section 3.1), we find that the unconditional volatility of $\mathscr{L}^{(m)}_{t+n}$ is one-tenth the volatility of the realized market return in our sample. As such, our empirical estimate of forecast errors is likely to be useful even for SDFs different from $M_{t,t+n} = 1/R_{t,t+n}$.

The above analysis pertains to expected returns. We also consider expected excess returns (risk premia) net of the n-period risk-free rate; for this case, as discussed in Section 2.1, we assume throughout that the expectations hypothesis holds, so that forward rates for risk premia correspond to expected future risk premia.

We also note in passing that expected log returns $\mathbb{E}_t[r_{t,t+n}]$ are conceptually distinct from log expected returns $\ln \mathbb{E}_t[R_{t,t+n}]$ (with the difference depending on higher moments of the return distribution). Expectations of the former are more relevant for prices, as prices depend on geometric-average expected returns (as can be seen in a Campbell-Shiller decomposition). This motivates our focus on expected log returns, but this distinction should be kept in mind in interpreting our results. We return to this issue in the subsection just below.

2.3 Identification from Surveys

We can measure spot rates, forward rates, and forecast errors using surveys that elicit return expectations at multiple horizons. While this survey-based measurement is more straightforward than with options data, it is still not fully model-free. First, interpreting forecast errors as reflecting intertemporally inconsistent expectations requires that the survey sample is representative over time (or that the set of respondents stays fixed).

Second, and more substantively, the elicited survey expectations allow for measurement of expected returns $\mathbb{E}_t[R_{t,t+n}]$ in non-log terms. As above, our interest is largely in expected log returns. But as we do not observe these expected log returns directly in the survey data, we measure spot rates here as $\ln \mathbb{E}_t[R_{t,t+n}]$. This differs from the true value $\mathbb{E}_t[r_{t,t+n}]$ by

$$\ln \mathbb{E}_t[R_{t,t+n}] - \mathbb{E}_t[r_{t,t+n}] = \sum_{n=2}^{\infty} \frac{\kappa_n}{n!},$$

where κ_n is the n^{th} cumulant of the distribution for the log return $r_{t,t+n}$ (so κ_2 is the variance, κ_3 is skewness, and so on). By working with $\ln \mathbb{E}_t[R_{t,t+n}]$ rather than $\mathbb{E}_t[r_{t,t+n}]$, we are thus assuming that the higher-order cumulants are fixed over time when estimating forecast errors. If this is not the case, forecast errors can be thought of as reflecting simple return expectations. These are of interest as well, but not exactly equivalent to the option-based measures.

3 Data and Implementation

This section describes the option and survey data and discusses implementation of the option-based and survey-based measures of expectations.

3.1 Option-Based Measures

Options Data. We obtain data on option prices largely from OptionMetrics. The data consists of European-exercise put and call options on major stock market indexes around the world. We have options on a total of 20 indexes from at least 15 countries, of which we select the 10 largest and most liquid indexes for the main panel analysis, as shown in Table A1. We also separately focus on the S&P 500, which is the longest available sample. For the S&P 500, the sample starts in 1990. For the global sample, the sample starts between 2002 and 2006. In both cases, the sample runs through 2021. We use end-of-month prices and apply standard filters. Details on the data, sample selection, and filters are relegated to Appendix C.

For the risk-premium analysis, we measure expected returns relative to currency-matched and maturity-matched LIBOR rates, as provided by OptionMetrics.

Measuring the LVIX. We measure $\mathscr{L}_t^{(n)}$ using standard results from Breeden and Litzenberger (1978) and Carr and Madan (1998), following Gao and Martin (2021). We make the simplifying assumption that the ex dividend payment at time $t+n$ is known ex ante at time t.[7] We denote the relevant index price at time t as P_t and the corresponding time-t

[7] More generally, our option analysis can be thought of as focusing on expected capital gains, rather than expected total returns. This follows much of the literature (e.g., Adam, Marcet, and Beutel 2017, Martin 2017), and we think dividends are unlikely to meaningfully affect the analysis: while expected dividend growth is time-varying, the contribution of dividends to short-horizon expected returns depends largely on the dividend yield. The CFO and Vanguard surveys analyzed below do consider total returns.

forward price at time $t+n$ as $F_t^{(n)}$. We denote the time-t prices of put and call options on P_{t+n} with strike K as $\text{put}_t^{(n)}(K)$ and $\text{call}_t^{(n)}(K)$, respectively. With these definitions,

$$\mathscr{L}_t^{(n)} = \frac{1}{P_t}\left[\int_0^{F_t^{(n)}} \frac{\text{put}_t^{(n)}(K)}{K}dK + \int_{F_t^{(n)}}^{\infty} \frac{\text{call}_t^{(n)}(K)}{K}dK\right] + r_{t,t+n}^f, \tag{7}$$

which means that the LVIX is a weighted sum of option prices on the relevant index (as well as a function of the current and forward price of the index).

Implementation for Option-Based Measures. We run our baseline analysis at the 6-month horizon ($n = 6$); for forward rates, our baseline also uses forward maturities of $m = 6$ months. This choice reflects a tradeoff: longer horizons are more economically meaningful, but long-maturity options are generally less liquid. Longer-dated options with multi-year horizons do, however, trade on a relatively liquid basis for the Euro Stoxx 50; we make use of these options when estimating the term-structure dynamics of forecast errors. We also provide a robustness analysis at alternative horizons. The 6-month horizon also allows for useful comparison with the survey-based measures, as one of the surveys (the Livingston survey) allows us to construct spot and forward rates only for $n = m = 6$. We annualize all returns in the empirical analysis.

We do not observe options for all positive strike values, as is needed in (7). We instead truncate the integral after extrapolating well past the range of observable contracts, as is standard practice in calculating option-implied moments. Details are in Appendix C. We examine possible approximation errors in Appendix D.

3.2 Survey-Based Measures

Our analysis uses data from three surveys: the Livingston survey of professional forecasters, the Duke–Fed CFO survey, and the Vanguard retail survey.

Livingston Survey: Data and Implementation. We obtain data on the Livingston survey of professional forecasters from the Philadelphia Fed. Forecasts for the S&P 500 are available twice annually (in June and December), and the sample runs from June 1992 to December 2021. The survey elicits the expected level of the S&P 500 in 6 months and 12 months, as well as the value at the end of the zero month.[8] We use median responses. We calculate the 6-month spot rate as the annualized log expected price change from month 0 to month 6, and we similarly calculate the 6-month, 6-month forward rate as the annualized log expected price change from month 6 to month 12.[9] We compare this forward rate to the next survey's realized spot rate to calculate forecast errors.

We also re-conduct the analysis using forecasts of equity risk premia rather than expected returns. To do so, we use the median Livingston survey forecasts of the 3-month T-bill rate (6-month rate forecasts are not elicited). For our spot rate, we subtract the annualized log T-bill rate as of month 0. For our forward rate, we subtract the annualized forecasted log T-bill rate in 6 months. We then calculate forecast errors as above. Because we observe interest-rate expectations directly, bond term premia do not affect these forward risk premia.

CFO Survey: Data and Implementation. We obtain data on the Duke–Fed CFO survey from the Federal Reserve Bank of Richmond. The sample is quarterly and runs from December 2001 to July 2025.[10] We use mean responses as median responses are not available until 2020. The survey asks respondents about expected total annualized returns on the S&P 500 at the 1-year and 10-year horizons. With these forecasts, we calculate the 1-year, 9-year forward rate, the current expectation of the 9-year expected return in one year. We compare this forward rate to the realized 10-year spot rate in one year, both in annualized terms, to calculate forecast errors. We do not observe the realized 9-year spot rate, but it is likely close to the 10-year spot rate on an annualized basis.

[8]For example, for the December 2018 survey, respondents are asked to forecast the S&P 500 as of the end of December 2018, June 2019, and December 2019.

[9]We ignore dividends in order to align this analysis with the option-based estimates. This is again unlikely to affect our conclusions given the relatively short horizon. Note that we also assume here that using the ratio of forecasted prices is valid to calculate expectations of the 6-to-12-month return.

[10]There are five quarters with missing forecast errors; the survey omitted returns once in 2019 and twice in 2020. We use the full available CFO sample to estimate the post-2020 effect (discussed just below) precisely. Similarly, for the Vanguard data (see Section 3.2), we use the full sample given the short data span.

For the risk-premium analysis, we measure spot risk premia relative to maturity-matched Treasury yields from Liu and Wu (2021) via Cynthia Wu's website.

In using the CFO data, we also account explicitly for a change in the survey administration that took place in 2020. Prior to 2020, the survey was administered by Duke University; in 2020, it transitioned to being jointly administered by Duke and the Federal Reserve Banks of Richmond and Atlanta (Graham et al. 2020). The sample composition, survey design, and elicitation of return expectations changed at that date, which we discuss further in Appendix C.3. These changes generated a discrete upward level shift in the CFO forward rates, as can be seen in Figure A9. To avoid contamination from this break-induced variation in CFO forward rates and remove the level shift, all CFO survey regressions include an indicator variable equal to one for observations from 2020 onward. Figures with CFO forward rates also plot the data net of the post-2020 indicator effect.

Vanguard Survey: Data and Implementation. We also obtain data on the Vanguard survey of retail investors, introduced by Giglio et al. (2020) and extended in Giglio et al. (2021). The survey is conducted every two months, with the sample running from February 2017 to December 2024.[11] The survey asks roughly 2,000 retail investors about their expected total annualized returns on the U.S. stock market at the 1-year and 10-year horizons, and we observe mean forecasts. With these forecasts, we then calculate the 1-year, 9-year forward rate — the current expectation of the 9-year expected return in one year — in exactly the same way as for the CFO survey.

Forecast errors and risk-premium measures also follow an identical construction as for the CFO data. For forecast errors, we again compare the ex ante forward rate to the realized 10-year spot rate in one year, both in annualized terms. For risk premia, we subtract Treasury yields from Liu and Wu (2021) and then calculate forward rates and forecast errors as before.

[11]We obtain the post-2020 data from Vanguard (2025), whose pre-2020 data matches that of Giglio et al. (2021). The data contain one unscheduled survey in March 2020 in addition to the usual bimonthly surveys. We drop this observation from our regression analyses, as its irregular spacing makes forecast-error calculations difficult. Given the large increase in forward rates in March 2020, this makes our results more conservative.

4 Empirical Results

This section studies spot rates, forward rates, and forecast errors for option-based and survey-based measures of expectations. Across the data sources, we find that forward rates are excessively countercyclical relative to realized spot rates. This applies consistently for expected returns and even more strongly for risk premia.

4.1 Empirical Results from Option-Based Measures

After measuring the LVIX, we construct empirical proxies for spot rates, forward rates, and forecast errors using the expressions provided in Proposition 1. As in that proposition, these proxies are exact under log utility ($M_{t,t+n} = 1/R_{t,t+n}$), which provides a useful benchmark for interpreting the following results. We begin by considering the relationship between spot and forward rates. We then turn to forecast errors; given the results in Propositions 2–3, these forecast-error estimates are less heavily reliant on the assumption of log utility. We end with a discussion of the potential role of the risk adjustment term for our results.

Throughout the majority of this subsection, we study the behavior of both forward risk premia and forward expected returns. As discussed in Section 2, for forward risk premia to reflect expected spot premia, the expectations hypothesis must hold for the bond term structure, an assumption we do not need to make for forward expected returns. However, forward risk premia have the advantage that they are stationary in our sample, whereas the decline in risk-free rates induces a trend in forward expected returns. Many of the analyses with risk premia are accordingly somewhat better behaved. For time-series plots of spot and forward rates, we consider risk premia in order to focus on the relevant cyclical variation rather than the trend in interest rates.

Summary of Spot and Forward Rates. We begin with an overview of our estimated spot and forward rates. Table A2 in the Appendix provides summary statistics by exchange; we focus here on descriptive figures for the U.S. sample. The top panel of Figure 2 shows contemporaneous time-series estimates for 6-month spot rates and 6×6-month forward rates for risk premia. Consistent with Figure 1 in the introduction, spot and forward rates increase significantly in crises. The slope of the term structure of expected risk premia is procyclical, as forward rates increase less than one-for-one with contemporaneous spot rates in bad times.

The bottom panel of Figure 2 instead compares forward rates with ex post realized spot rates. The gap between the realized spot rate and the forward rate is the risk-premium forecast error. As in Figure 1, forecast errors tend to be negative after crises, when spot rates decline from their crisis peaks. By contrast, forecast errors are frequently (though not always) positive outside of crises. As a result, forecast errors appear to be procyclical.

To formalize this visual analysis, we move now to a set of regression analyses.

Forward Rates as Predictors of Future Spot Rates. We first consider Mincer-Zarnowitz (1969) regressions of realized spot rates on forward rates:

$$\tilde{\mu}_{i,t+6}^{(6)} = \beta_0 + \beta_1 \tilde{f}_{i,t}^{(6,6)} + e_{i,t+6}, \tag{8}$$

where i now indexes the exchange. We consider both panel regressions and U.S.-only time-series regressions, and we conduct overlapping monthly regressions for $n = m = 6$ months, following Section 3.1. Standard errors are clustered by exchange and date in the panel case; in the U.S.-only case, we use heteroskedasticity- and autocorrelation-robust (HAR) Newey-West standard errors, with lags selected following Lazarus et al. (2018), and fixed-b critical values.

Table 2 presents the results. The first three columns consider spot and forward rates for risk premia (i.e., $\tilde{\mu}$ and $\tilde{f}$), while the last three columns show results for expected returns (i.e., μ and f). The relevant null for forecast predictability features $\beta_0 = 0$, $\beta_1 = 1$; we consider these two restrictions separately as well as jointly. For risk premia, the null holds under rational expectations if $M_{t,t+n} = 1/R_{t,t+n}$ and the expectations hypothesis holds for bonds. For expected returns, it holds under rational expectations if $M_{t,t+n} = 1/R_{t,t+n}$.

Column (1) shows results for risk premia from the main panel, as defined in Section 3.1, without any fixed effects. The slope coefficient is around 0.64, and the intercept is statistically significant at 1.04 annualized percentage points. We can therefore reject the null of $\beta_0 = 0$, $\beta_1 = 1$ with substantial confidence for this specification. The fact that $\widehat{\beta}_1 < 1$ suggests that forward rates tend to overshoot future spot rates in the data. But in spite of this overshooting of forward rates, the R^2 for this regression is around 0.2, so forward rates do predict a substantial portion (one-fifth) of the variation in the equity premium ex ante.

Columns (2) and (3) show that similar results obtain with exchange fixed effects (our preferred panel specification, as it identifies within-exchange predictability) and in the U.S. sample alone, with slope coefficients slightly further below 1 and within-exchange R^2 slightly below 0.2. Columns (4)–(6) consider expected returns rather than risk premia. These regressions suggest a higher degree of predictability, but this is driven by persistent variation in the risk-free rate rather than higher-frequency variation in forward rates.[12]

One might be concerned with attenuation bias given possible measurement error in forward rates. We address this by instrumenting the 6×6 forward rate $f_{i,t}^{(6,6)}$ with the 2×1-month forward rate $f_{i,t}^{(2,1)}$, given that options are more liquid at shorter maturities. Table A3 presents the resultant two-stage least squares estimates.[13] For the risk-premium regressions, the estimated slope coefficients increase slightly and are now quite consistent across specifications at around 0.7–0.8. We continue to reject the null of $\beta_0 = 0$, $\beta_1 = 1$. For the expected-return regressions, the results hardly change.

Overall, we find that while forward rates have substantial predictive power over future spot rates, they tend to overshoot those spot rates. This is particularly true for forward risk premia: when forward rates for risk premia increase by 1% versus their within-exchange mean, this corresponds on average to an increase in future spot rates of only about 0.7%.

[12] Table A7 examines split-sample regressions in the U.S. sample. Both the regression slope and R^2 are substantially lower in the later subsample, suggesting that the reported results are driven by declines in the risk-free rate. While such secular variation in interest rates likely affected equity valuations (as examined by Gormsen and Lazarus 2026), our focus is instead on higher-frequency variation in forward expected returns; we accordingly prefer specifications netting out the risk-free rate.

[13] The first stage in these regressions (not shown) has an R^2 close to 70% when excluding all fixed effects, suggesting the 2×1 forward rate is a strong instrument for the 6×6 rate.

Forecast Errors: Insignificant on Average but Predictable. We next turn to forecast errors, which can be estimated under weaker assumptions than spot and forward rates themselves (see Propositions 2 and 3). We first consider unconditional averages $\bar{\varepsilon}_{i,t+6}^{(6)}$ across different samples in Table 3. For risk premia (columns (1) and (2)), average forecast errors are statistically insignificant and effectively zero: $\bar{\varepsilon}_{i,t+6}^{(6)}$ is below 20 basis points annualized. Realized spot rates for risk premia have thus been very slightly (and insignificantly) higher than forward rates on average. This rough magnitude applies across exchanges, though the U.S. average in column (2) is even lower, at roughly 2 basis points.

For expected returns (columns (3) and (4)), average forecast errors are slightly negative and statistically significant. Realized spot rates for expected returns have thus been slightly lower than forward rates on average. This finding is likely driven by an unexpectedly large decline in interest rates over the sample. Despite the statistical significance of these estimates, we note that the averages are still close to zero.

While forecast errors are close to zero unconditionally, this average masks substantial predictability over time. Figure 3 plots realized risk-premium forecast errors along with forward rates and indications of prominent crises in the U.S. sample. The figure further illustrates the takeaway from Figure 1 in the introduction: following crises, forecast errors tend to be significantly negative, as future realized spot rates are lower than expected ex ante. To the extent that these post-crisis forecast errors are systematically predictable, the errors represent what we refer to as excess cyclicality in forward rates: forward rates are excessively elevated during crises, relative to subsequent realized spot rates. In the remainder of this section, we study whether the forecast errors are indeed systematic through a set of predictive regressions.

For a formal consideration of forecast-error predictability, we begin by testing whether forward rates predict forecast errors in regressions of the following form:

$$\tilde{\varepsilon}_{i,t+6}^{(6)} = \beta_0 + \beta_1 \, \tilde{f}_{i,t}^{(2,1)} + e_{i,t+6}. \tag{9}$$

We use the 2×1 forward rate $\tilde{f}_{i,t}^{(2,1)}$ instead of the 6×6 rate, as we would otherwise have the same forward rate on both sides of the regression, leading to upward bias in the slope in the presence of measurement error.[14] All other aspects of the regressions are identical to Table 2.

The first three columns of Table 4 report the results of regression (9) for risk premia. In the main panel without any fixed effects (column (1)), the slope coefficient on the forward rate is -0.13 and is significant at the 5% level.[15] In our preferred panel specification including exchange fixed effects (column (2)), the slope coefficient is now significant at the 1% level. Results are similar in the U.S.-only sample, although power is naturally slightly lower.

We next consider the forecast errors for expected returns (ε) on the left-hand side. When doing so, we continue to use the forward rate for risk premia as the predictor variable; forward rates for expected returns themselves have a strong downward trend due to the risk-free rate component, so using these forward rates as predictor variables would amount to asking if forecast errors are trending over the sample (e.g., positive in the early sample and negative in the later sample). By instead using the more transitory forward rate for risk premia ($\tilde{f}$) as the predictor variable, the predictor is more directly related to the state of the economy and better suited for capturing the type of crisis-dependent errors plotted in Figures 1 and 3. These regressions (columns (4) to (6) of Table 4) thus have transitory components on both the left- and right-hand side, and they have the benefit that the theoretical predictions for the forecast errors on the left-hand side do not require assumptions about the term structure of interest rates (they require only the assumptions laid out in Propositions 2 and 3).

As can be seen in Table 4, the forecast errors for expected returns are highly predictable. In the main panel, the slope coefficients in columns (4)–(5) roughly double relative to the regressions in columns (1)–(2). The R^2 increases to around 10% with exchange fixed effects. For the U.S.-only regressions, the difference between risk premia and expected returns appears more modest. Overall, the results suggest that the patterns observed in columns (1)

[14]Moreover, with the same conditioning variable, the difference between the slope in Mincer-Zarnowitz and error-predictability regressions would be one by construction.

[15]These magnitudes are not directly comparable to those in Tables 2 given the use of $\tilde{f}_{i,t}^{(2,1)}$ as the predictor. Instead, regression (9) is analogous to the reduced form of the IV specification for (8) in Table A3.

through (3) are not driven by the risk-free rate component (i.e., a failure of the expectations hypothesis for interest rates); if anything, considering the total expected return strengthens the results in this case. And in all cases, forecast errors are robustly predictable from the level of forward rates. This test for forecast-error predictability is evidently more powerful against a rational null than the Mincer-Zarnowitz regressions presented above.

One might be concerned that forecast-error predictability is the artifact of measurement error. For this measurement issue to matter, we would have to *overestimate* forward rates in bad times, generating predictably negative forecast errors. However, in Appendix D, we find that if anything, we *underestimate* forward rates, especially in bad times, as longer-maturity options are generally less liquid.

We conclude that the pattern of excess forward-rate cyclicality observed in Figures 1 and 3 is systematic and statistically significant. Investors have believed that spot rates would be more elevated following crises than would ex post have been rational. This conclusion does require that the risk-premium adjustment in Propositions 2 and 3 is small, so that the options provide useful estimates of subjective forecast errors. We consider the behavior of the risk-premium term in detail in Appendix E, but our main source of confirmation for our interpretation of small risk premium effects is in our survey-data analysis, considered below.

4.2 Empirical Results from Survey-Based Measures

We now turn to spot rates, forward rates, and forecast errors estimated using survey expectations. As a preview, Figure 4 plots forward risk premia for the three surveys, and Figure A1 shows their common variation with the option-based forward risk premia. As with the option-based measures, we find that (i) forward rates are countercyclical, as they increase in bad times, and (ii) forecast errors are predictably negative following crises, as will be shown below. These results are again consistent with excess cyclicality in forward expectations: during bad times, investors think future spot rates will be more elevated than their own beliefs justify ex post. We detail this analysis below for each of the surveys in turn.

Livingston Survey. As explained in Section 3.2, we use the Livingston survey of professional forecasters to estimate the 6-month spot rate and the 6-month, 6-month forward rate. As with the option-based measures, we consider spot and forward rates for both risk premia and expected returns. Since the Livingston survey offers expectations about interest rates, we can measure forward rates for risk premia without assumptions about the expectations hypothesis for interest rates (see Section 2.1). The sample runs from 1992 to 2021.

Columns (1) and (2) of Table 5, Panel A, show the results of Mincer-Zarnowitz regressions of future spot rates on ex ante forward rates. We find slope coefficients of 0.81 for risk premia and 0.68 for expected returns. These estimates echo those from the option-based measure, suggesting that a one-percentage-point increase in forward rates is associated with spot rates that are about 0.75 percentage points higher in the future. The R^2 values are somewhat higher than those from the option-based regressions for risk premia, somewhat lower than those for expected returns; on average, around half of the variation in spot rates can be predicted by the forward rate at the 6-month horizon. Columns (3) and (4) further show that the forecast errors are close to zero on average and statistically insignificant.

Columns (5) and (6) of the panel show results from error-predictability regressions of future realized forecast errors on ex ante forward rates. As for the option-based measures, we find significantly negative slope coefficients around -0.2, suggesting that high forward rates are associated with lower-than-expected future spot rates. The R^2 values are around 0.05, which are again higher than those for option-based risk premia, lower than those for expected returns. Given the countercyclical nature of the forward rates, as shown in Figure A1, the pattern is similar to the one obtained from the options: in bad times, investors think future expected returns will be more elevated than investors' own subsequent beliefs justify.

CFO Survey. As explained in Section 3.2, we use the Duke–Fed CFO survey to estimate the 1-year, 9-year forward rate, which is the 9-year return expectation one year from now. We again consider forecast errors for both risk premia and expected returns; since we measure risk premia relative to the yield curve, we require that the expectations hypothesis holds, as with the option-based measure. The key advantage of this survey is the horizon: longer-horizon expected returns are more important for prices. The sample runs from 2001 to 2025.

Columns (1) and (2) of Table 5, Panel B, report the results of Mincer-Zarnowitz regressions of future spot rates on forward rates. Forward rates predict future spot rates with an average slope coefficient around 0.5. The R^2 values are comparable to that in the Livingston survey, but are somewhat different than that for the option-based measure. Columns (3) and (4) report small average forecast errors. Columns (5) and (6) show the results of error-predictability regressions. For risk-premium forecast errors, we observe strong predictability, with an R^2 of 0.48. For expected returns, this effect is statistically insignificant. The former is much stronger than that from the option-based measure or the Livingston survey, and the latter slightly weaker. The difference might arise from the longer horizon.

We note that the CFO survey exhibits a clear distinction between the cyclical behavior of spot and forward rates. As shown in the top panel of Figure 5, spot rates appear at least weakly *procyclical*: short-term subjective risk premia generally decrease slightly in bad times (as highlighted by Greenwood and Shleifer 2014). As seen in Figure 4, however, forward rates are countercyclical, as are long-term spot rates (unplotted), so that long-term subjective risk premia increase in bad times. One interpretation is that the CFO respondents understand present value logic — future long-run returns are high when prices are low — but during crises, they believe prices will continue to decrease for a while before increasing.

Vanguard Retail Survey. We now consider the sample of Vanguard retail investor expectations, from which we also calculate 1-year, 9-year forward rates and associated forecast errors as in Section 3.2. The 2017–2024 sample is somewhat shorter than the other data sources, but it does include the early-2020 downturn and thus contains substantial cyclical variation. We can thus conduct the same analyses for this survey as for the remaining data sources. This gives us further evidence for long-horizon errors, alongside the CFO analysis.

Panel C of Table 5 presents the results for the Vanguard regressions. Columns (1) and (2) show that forward rates predict future spot rates with an average slope coefficient very similar to that for the CFO survey (around 0.5), albeit with slightly lower R^2 values. In both cases, we reject the joint null of forecast efficiency: forward rates move too much relative to realized spot rates, as with the other data sources. Columns (3) and (4) show that average forecast errors are again economically and statistically insignificant. Columns (5) and (6) show that forecast errors are strongly predictable for risk premia, with a significantly negative slope coefficient of -0.57 and an R^2 of 0.21. This predictability is weaker for expected returns. Overall, the results are quite consistent with those for the other data sources.

Finally, the bottom panel of Figure 5 shows that spot risk premia for these retail investors appear weakly procyclical over this shorter sample, similar to the CFOs. This again contrasts with the forward risk premia in the bottom of Figure 4, which increase meaningfully in the Covid crisis. So there is consistently strong common countercyclical variation in forward expectations. While short-term beliefs differ across investor groups, there is less disagreement about forward rates: across all option-based and survey-based measures, longer-run forward rates are countercyclical, and excessively so. The apparent disagreement in previous literature thus appears related to the cyclical behavior of short-term spot rates. While short-term spot rates are important for understanding the expectations-formation process, long-run risk premia and forward rates are often the key object of interest in determining prices.

4.3 Common Cyclical Variation in Options and Surveys

We conclude this section by showing how the results from the different data sources align with each other and with business- and financial-cycle state variables.

Common Variation in Forward Rates. The analysis so far has considered each data source individually. We now explore their common variation by relating the survey-based forward rates to the forward rates obtained from option prices. Figure A1 documents a relatively strong comovement between option-based forward risk premia and Livingston, CFO, and Vanguard sources, particularly during crisis episodes but also in normal times.

The previous subsections establish that high forward rates predict negative forecast errors within each investor group, as subsequent spot rates are systematically lower than anticipated. If this pattern is driven by a fundamental economic mechanism, we might expect it to work *across* investor groups, such that a high forward rate in a given data source predicts negative forecast errors in another data source. To test this, we run cross-predictability tests in which we predict survey-based forecast errors with option-based forward rates. Table A4 shows that option-implied forward rates predict realized forecast errors for both Livingston and CFO surveys.[16] This is true for forward rates measured both as risk premia and expected returns, and results are significant in all cases except for the CFO-based risk premia ($p = 12.7\%$). These strong cross-predictability results, along with Figure A1, indicate that forward rates across investor groups share a common driver and all comove with the cyclical component of discount rates. We turn to this cyclical variation next.

Cyclical Variation in Forward Rates and Forecast Errors. Our evidence on the countercyclicality of forward rates so far has largely come from visual evidence on their behavior in crisis episodes. We now conduct a more systematic analysis of their comovement with standard cyclical indicators. We then consider the cyclical behavior of forecast errors.

To estimate cyclical behavior, we need proxies for the appropriate state variables. We first consider three benchmark measures. The first two are measures of objective risk premia, following Dahlquist and Ibert (2026): the excess CAPE yield (Shiller, Black, and Jivraj 2020) and the SVIX index (Martin 2017). The third is a measure of the state of the economy: an NBER recession indicator. Figure 6 presents results of regressions of forward risk premia, as measured from the four separate data sources, on each of these three countercyclical state variables separately. All the forward risk premia load positively on each measure, implying countercyclical variation in forward risk premia (in line with Dahlquist and Ibert's results for spot risk premia). In addition to all 12 point estimates being positive, 9 of them are significant at the 5% level. As with the regression tables, significance is calculated based on two-sided tests with HAR standard errors and fixed-b critical values. This approach follows Lazarus

[16]Given the shorter Vanguard sample, we omit it from this exercise to make the table easier to read.

et al. (2018) and Lazarus, Lewis, and Stock (2021) and is more conservative than textbook Newey-West.[17] So while small-sample inference is difficult (and size distortions remain a possibility), the best available evidence suggests clear forward-rate countercyclicality.

In Figure A2, we expand the main measures of the state of the economy to include the five variables that Nagel and Xu (2023) identify as business-cycle indicators: industrial production, the term spread, the default spread, the real factor from Ludvigson and Ng (2009), and the squared VIX. Once including these variables, 31 of the 32 estimates are countercyclical, and the majority are statistically significant. Panel B of Figure A2 considers forward expected returns, and finds that 25 of the 32 estimates are countercyclical.

In addition to these main variables drawn from past literature, Figure A2 considers an additional set of eight variables — including, for instance, asset managers' expected returns, from Dahlquist and Ibert (2024) — reaching a total of 16 state variables. We find that 62 out of 64 estimates are countercyclical for risk premia and 55 are countercyclical for expected returns, although many are insignificant. None of the estimates are significantly procyclical.

We next turn to the cyclicality of forecast errors. Since forward rates are countercyclical and predict forecast errors negatively, the forecast errors vary procyclically. To see this, we regress future realized forecast errors on the main state variables from Dahlquist and Ibert (2026) and Nagel and Xu (2023). Figure A3 shows the results. For risk premia, 21 of the 28 baseline state variables imply procyclical forecast errors; for expected returns, 24 of the 28 imply procyclical forecast errors.[18] Significance is naturally lower in the predictive regressions than in the contemporaneous regressions, but it is nonetheless meaningful: for example, the 24 procyclical estimates for expected-return forecast errors have an average t-statistic of 2.50, while the 4 countercyclical estimates have an average t-statistic of 0.77.

Overall, the results in this section confirm that forward rates are countercyclical, and

[17]In particular, the long-run variance estimator uses substantially more lags than the classic Andrews (1991) rule, and fixed-b critical values are uniformly larger than standard normal quantiles. Even with this conservatism, results for the eight-year Vanguard sample are significant due to the very large increase in forward risk premia in 2020, as seen in Figure 4.

[18]NBER recessions are no longer considered in the baseline set of state variables, as these are backdated and ill-suited for predictions. We keep them in Figure A3 purely for consistency with the preceding analysis.

excessively so: during bad times, forward rates increase more than one-for-one relative to future spot rates, yielding predictably negative forecast errors after crisis spikes for both risk premia and expected returns. This cyclical behavior is strongly related across investor groups, and their forward rates and forecast errors all share comovement with economic state variables. In the following section, we examine the asset-pricing implications of these findings.

5 Implications for Asset Pricing

This section studies additional implications of excess cyclicality in forward rates.

5.1 A New Source of Excess Volatility

If negative asset-pricing shocks lead to predictably larger increases in forward rates than justified by subsequent realizations, then such shocks may lead to excess volatility in prices. To better understand the relevance of this source of excess volatility, we conduct two illustrative quantifications here. We first quantify the impact of option-based forecast errors on prices; we then conduct a similar exercise using survey-based forecast errors. We find that both explain a sizable share of the observed volatility during crises. Each requires different assumptions but builds off the same overarching theory, which we turn to first.

Defining Excess Volatility from Forecast Errors. We focus on the component of excess volatility coming from expectations about risk premia, as opposed to expected returns; this leads to generally smaller estimates. Starting from a Campbell-Shiller (1988) decomposition for the log price-dividend ratio $p_t - d_t$ and using the definition of forward rates:

$$\begin{aligned} p_t - d_t &= k + \sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t[\Delta d_{t+j+1}] \; - \sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t[r_{t+j,t+j+1}] \\ &= k + \underbrace{\sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t[\Delta d_{t+j+1}]}_{CF_t} \; - \underbrace{\sum_{j=0}^{\infty} \rho^j \, \tilde{f}_t^{(j,1)}}_{\mathcal{F}_t} - \underbrace{\sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t\Big[r^f_{t+j,t+j+1}\Big]}_{RF_t}, \end{aligned} \tag{10}$$

where k is a constant and $0 < \rho < 1$.

We further split the forward-rate term $\mathcal{F}_t$ into

$$\underbrace{\sum_{j=0}^{\infty} \rho^j \tilde{f}_t^{(j,1)}}_{\mathcal{F}_t} = \underbrace{\sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t \left[\tilde{\mu}_{t+j}^{(1)} \right]}_{\mathcal{M}_t} - \underbrace{\sum_{j=1}^{\infty} \rho^j \, \mathbb{E}_t \left[\tilde{\varepsilon}_{t+j}^{(1)} \right]}_{\mathcal{E}_t}. \tag{11}$$

Our interest is in the term $\mathcal{E}_t$, which quantifies the component of the log price-dividend ratio attributable to all future predicted forecast errors, or the price impact of forecast errors. Note that the predictable forecast error for $j = 0$ is zero by construction given equation (3), so we have rewritten the sum for $\mathcal{E}_t$ as being from $j = 1$ to ∞. In what follows, we often refer to this $\mathcal{E}_t$ series as *discounted forecast errors*, for short.

The true $\mathcal{E}_t$ is unobservable. The main empirical challenge is that we cannot easily estimate the entire term structure of predicted forecast errors — that is, the entirety of the infinite sum — without additional assumptions, as we typically only observe relatively short-horizon forecast errors. Instead, we consider two proxies $\widehat{\mathcal{E}}_t$ for the true price impact.

Measuring Discounted Forecast Errors. We measure $\widehat{\mathcal{E}}_t$ using forecast errors from both options and surveys. For options, our approach requires more assumptions but allows us to estimate $\mathcal{E}_t$ in full; that is, our estimate captures the entire infinite sum. For surveys, our approach requires fewer assumptions but provides only a lower bound for $\mathcal{E}_t$; that is, our estimate truncates the infinite sum.

We begin with option-based measures. In the baseline U.S. sample, we estimate forecast errors for horizons only up to one year. To obtain estimates for longer horizons, we estimate the term structure of predictable forecast errors using spot and forward risk premia out to an 8-year horizon for the Euro Stoxx 50, which is the only exchange in our sample with index options available at such maturities.[19] For each horizon, we predict forecast errors as in (9)

[19]This is similar to the case for the dividend futures market (Binsbergen and Koijen 2017). We note that we use Euro Stoxx 50 data in this exercise only to estimate the term structure behavior of forecast errors; we do not assume that longer-horizon errors are equal across exchanges on a period-by-period basis. Our assumption is therefore that the term structure dynamics are similar across exchanges on average (as

by regressing realized forecast errors on shorter-horizon forward rates. We then estimate a decay function $\phi^{(n,m)}$ for forecast errors, such that $\mathbb{E}_t[\tilde{\varepsilon}^{(m)}_{t+n}] = \phi^{(n,m)}\,\mathbb{E}_t[\tilde{\varepsilon}^{(1)}_{t+1}]$. Details on the long-horizon options data and on our estimation approach are in Appendix G.2. As reported in Table A12, the decay parameters are close to or greater than 1 at all horizons. We therefore assume that predictable forecast errors have a flat term structure, $\mathbb{E}_t[\tilde{\varepsilon}^{(1)}_{t+j+1}] = \mathbb{E}_t[\tilde{\varepsilon}^{(1)}_{t+j}]$ for all j in (11), to obtain a back-of-the-envelope quantification of $\mathcal{E}_t$ for the U.S. sample:

$$\widehat{\mathcal{E}}_t^{\text{option}} = \frac{\rho}{1-\rho}\,\mathbb{E}_t\left[\tilde{\varepsilon}^{(1)}_{t+1}\right]. \tag{12}$$

The length of one period is arbitrary, so we set it to 6 months to correspond to our baseline option horizon. We measure $\mathbb{E}_t[\tilde{\varepsilon}^{(1)}_{t+1}]$ in (12) using the estimates in column (3) of Table 4.

The forecast-error decay estimates in Table A12 are in many cases well above 1, suggesting that the assumption of a flat term structure of predictable forecast errors may be conservative. On the other hand, since we cannot observe very long-horizon forecast errors, assuming this flat term structure over an infinite horizon may be somewhat restrictive. It also makes $\widehat{\mathcal{E}}_t^{\text{option}}$ dependent on the scaling factor $\frac{\rho}{1-\rho}$ in (12). We estimate $\rho \approx 0.967$,[20] and our scaling factor is thus $\frac{\rho}{1-\rho} \approx 29$. This scaling reflects the same logic as in standard present-value decompositions — persistent changes in discount rates are amplified in prices — and when implementing variance decompositions via a vector autoregression, one multiplies expected returns by a similar constant.

To obtain a set of direct estimates that do not require any scaling, we turn to the survey data. We focus on the surveys with long-horizon forward rates, the CFO and Vanguard surveys, and we quantify the price impact of predictable forecast errors during the 2008 financial crisis and the 2020 Covid crisis. The CFO survey is only usable for this exercise in the 2008 crisis,[21] and the Vanguard survey is only available during the 2020 crisis.

---

documented in the yield-curve context by Diebold, Li, and Yue 2008, among others).

[20]Both forecast errors and ρ are in annualized terms. We estimate $\rho = (1+\exp(\overline{d-p}))^{-1}$, where $\overline{d-p}$ is the mean repurchase-adjusted log dividend yield from Nagel and Xu (2022), to obtain $\rho \approx 0.967$.

[21]As discussed in Section 3.2, the CFO survey underwent a redesign and break between 2019 and 2020. We therefore drop the 2020 forecast errors to avoid splicing pre-break forecasts with post-break spot rates.

For these surveys, we quantify $\mathcal{E}_t$ with a slightly different approach. Rather than estimating the decay in predictable forward-rate errors, we exploit that the surveys directly elicit 10-year expectations and use these to compute the contribution of the first 9 years of forecast errors. As long as predictable errors $\mathbb{E}_t[\tilde{\varepsilon}^{(1)}_{t+j}]$ at horizons $j \geq 10$ have the same sign as those for $j \leq 9$, the contribution from the first 9 years then provides a lower bound on $\mathcal{E}_t$:

$$\widehat{\mathcal{E}}_t^{\text{survey}} = \sum_{j=1}^{9} \rho^j \, \mathbb{E}_t\left[\tilde{\varepsilon}^{(1)}_{t+j}\right] \leq \mathcal{E}_t. \tag{13}$$

For the surveys, we estimate the 1-year, 9-year predicted risk-premium error $\mathbb{E}_t[\tilde{\varepsilon}^{(9)}_{t+1}]$ using column (5) of Table 5, Panels B–C. For the one-period forward rates in (13), we apply iterated expectations and set $\mathbb{E}_t[\tilde{\varepsilon}^{(9)}_{t+1}] = \mathbb{E}_t[\tilde{\varepsilon}^{(1)}_{t+j}]$, as both are in annualized terms.

Price Impact of Forecast Errors. We then compare variation in our estimates of $\mathcal{E}_t$ to variation in $p_t - d_t$. To account for share repurchases and ensure that $p_t - d_t$ is stationary, we use the repurchase-adjusted price-dividend series provided by Nagel and Xu (2022).

Figure A4 visualizes the results for the option-based estimates, with both discounted forecast errors $\widehat{\mathcal{E}}_t^{\text{option}}$ (red) and the log price-dividend ratio $p_t - d_t$ (blue) shown as log differences from their full-sample means. Unconditionally, the forecast errors explain a modest amount of the overall variation in prices. A regression of forecast errors on the log dividend yield gives a slope coefficient of around 0.1, suggesting that 10% of the unconditional variation in prices can be tied back to predictable forecast errors. Moreover, the two series do not always comove positively.[22]

While the unconditional effect of forecast errors may be modest, the conditional effect in bad times appears sizable: during large stock-market declines (low $p_t - d_t$), predictable forecast errors play a significant role in the decomposition in (10)–(11). This holds particularly during the financial crisis and the Covid crisis. Figure 7 zooms in on these episodes. From June 2007 to the February 2009 crisis trough, $\widehat{\mathcal{E}}_t^{\text{option}}$ drops by 36 log points; during the same

[22]During the Russian debt crisis in the late 1990s, for example, forward rates were quite elevated (leading to negative estimated $\mathcal{E}_t$), but the dot-com boom largely continued apace (leading to high $p_t - d_t$).

period, the valuation ratio drops by 42 log points, suggesting that predictable forecast errors can justify a large part of the observed crash in the stock market. For the Covid crisis, $\widehat{\mathcal{E}}_t^{\text{option}}$ drops by around 22 points, which is the exact amount by which the valuation ratio drops, suggesting that the predictable forecast errors can explain effectively all of the crash under the scaling assumption used in (12).

Figure 7 also plots $\widehat{\mathcal{E}}_t^{\text{survey}}$ over the same two crises. The survey-based forecasts comove strongly with their option-based counterparts during these crises. During the global financial crisis, for which we rely on the CFO survey for forecast errors, $\widehat{\mathcal{E}}_t^{\text{survey}}$ drops by around 10 points over the crisis. The predictable forecast errors in the first 9 years of CFO forward rates can thus explain around one fourth of the drop in the stock market. This is again a lower bound for the full effect $\mathcal{E}_t$, under much weaker assumptions than used for the option-based measure.[23] During the Covid crisis, for which we rely on the Vanguard survey, $\widehat{\mathcal{E}}_t^{\text{survey}}$ again drops by around 10 points. The predictable Vanguard errors thus give a lower-bound price impact of nearly one half of the drop in the stock market during this period.

In Table 6, we provide a unified summary of the different numerical estimates of $\mathcal{E}_t$ during crises, along with $p_t - d_t$. The table shows that the option-based measure can explain more than half of the stock-price drop during the global financial crisis, and the full drop during the Covid crisis. The survey-based measures give rise to a more assumption-free lower bound, which confirms a large impact of predictable forecast errors in forward rates on valuations during crashes. In bad times, investors' overestimation of the length and magnitude of risk-premium increases can account for a significant share of observed price drops, and this holds across investor groups and measurement methods.

These price effects all use estimates of the predictable component of forecast errors. One can also more directly ask whether directly observed forward rates can explain movements in prices. The answer here appears to be yes. For the Vanguard survey, for instance, the

[23]The subjective short-horizon spot risk premium also varies, and this variation could exert offsetting effects on prices relative to $\mathcal{E}_t$. This is not our focus: we estimate only the predictable errors in forward rates. The spot rates in Figure 5, however, suggest that the offset is likely to be small in the two crises.

1×9-year forward risk premium increased by 2.5 percentage points in early 2020 (Figure 4). This translates to a decrease in prices of around $2.5 \times 9 \approx 22.5$. As a result, effectively the entirety of the observed Covid crash can be explained by changes in long-term discount rates.

5.2 The Term Structure of Equity

We have argued above that a large part of the variation in stock prices during crises may come from excess cyclicality in forward rates. According to this argument, fluctuations in prices are driven by revisions in expectations that are unexpected by investors (but predictable by traditional state variables). If this mechanism helps drive realized returns on equities, we should expect it to manifest itself in the behavior of the equity term structure, i.e., the prices and expected returns of dividend claims with different maturities.

In crises, when forecast errors are predicted to be negative (investors have overly high expectations about future spot rates), investors will in the future revise their expectations about spot rates downwards, leading to positive realized returns. This effect will be stronger for longer-maturity claims, because these have longer duration, so the forecast errors will push up both the level and slope of the realized equity term structure. Outside crises, positive average forecast errors instead push the level and slope of the term structure downwards.

We provide an estimated quantification of these effects in Figure A5. Specifically, we examine the impact of risk-premium forecast errors on realized returns to dividend strips with different maturities unconditionally as well as in good and bad times, following Gormsen (2021).[24] The figure shows that the estimated forecast errors can produce a term structure that is mildly downward sloping on average (Binsbergen, Brandt, and Koijen 2012) and countercyclical (Golez and Jackwerth 2024, Gormsen 2021). As such, our forecast errors help account for the dynamics of the equity term structure.

[24] While the results are numerically close to the average returns in Gormsen (2021), we mainly interpret the figure qualitatively rather than quantitatively: our estimates rely fairly heavily on modeling assumptions (and represent log returns as opposed to arithmetic returns).

5.3 Impact of Forecast Errors on Demand Elasticities

Our results may also help speak to the recent finding that investors' demand for equities is relatively inelastic with respect to price changes (Gabaix and Koijen 2022). If a drop in prices corresponds to a large, short-term increase in expected excess returns, then this inelasticity is puzzling. But our results suggest that investors often mistakenly attribute a significant share of the decrease in prices to increases in risk premia at relatively long horizons, and perceive short-term spot risk premia to increase only modestly. A modest increase in short-term risk premia will not lead to a large increase in desired portfolio weights in equities, providing a candidate explanation — albeit a speculative one — for the observed inelasticity. We leave potential quantification of this effect to future work.

5.4 A Model of Expectation Errors

Finally, to provide a positive potential explanation for our results, Appendix F presents a simple model of expectation errors in the term structure of spot rates, building on Hirshleifer, Li, and Yu (2015). The model matches both the excess cyclicality in forward rates across all investors and the cyclicality of short-term spot rates across different investors.

6 Conclusion

This paper provides novel analysis of the term structure of return expectations for equities. This term structure is a powerful laboratory for understanding how investors form expectations, and the behavior of this term structure is key for understanding variation in stock prices.

Our results suggest that the term structure of investor expectations is well-behaved in multiple ways: forward rates are reasonably good predictors of future realized spot rates, they embed mean-reversion, and they exhibit countercyclical variation. However, we find strong evidence that forward rates are *excessively* countercyclical: in bad times, investors' expectations of future expected returns and risk premia are more elevated than their own subsequent beliefs justify.

This excess cyclicality has broad implications for asset pricing. It can lead to substantial excess volatility in stock prices, and it is particularly relevant for explaining crashes: we find that the excess cyclicality can account for a large share of the drop in prices during the global financial crisis and the Covid crisis. These crises appear to have been much more short-lived than investors expected ex ante, with meaningful price impacts. The excess cyclicality in forward rates can also help explain the dynamics of the equity term structure and the apparently low price elasticity of demand for equities.

A distinctive feature of these empirical results is their robustness across option-based and survey-based measures. Whereas the dynamics of short-run expected returns vary across data sources, longer-run forward rates are robustly and excessively countercyclical across option-based and survey-based measures, and for both forward risk premia and expected returns, although somewhat more strongly for risk premia. We therefore provide a unified body of evidence on the dynamics of forward return expectations.

Tables and Figures

Table 1
Subjective Equity Premia: A Summary Across Horizons and Investors

This table summarizes findings on subjective equity return and risk premium expectations from past literature and the current paper. Plus signs (**+**) indicate cases in which subjective expectations tend to co-move positively with objective measures, while minus signs (**–**) indicate cases in which subjective expectations move the wrong direction with respect to objective measures. "Excess" refers to excessive positive co-movement with objective measures (or excess cyclicality). In all cases, the "Forward" column refers to findings documented later in the paper, for forward risk premia (our main measure) and forward expected returns. All other table entries are based on past literature, for which the first column in the table provides one or more leading references containing this finding. For the past literature, expectations are for excess returns, aside from the references for the CFOs and retail investors. For the long-run entry for CFOs, De la O and Myers (2021) find that this value is close to acyclical for expected returns, while we find that it is countercyclical (**+**) when focusing on equity premia.

	Horizon of expectations				
	Short-run	+	Forward *(this paper)*	=	Long-run
Sophisticated investors					
"Mr. Market" [Martin]	**+**		**+** (excess)		**+**
Asset managers/institutions [Dahlquist–Ibert]					**+**
Other professionals					
CFOs [Greenwood–Shleifer, De La O–Myers]	**–**		**+** (excess)		**+**
Prof. forecasters [Nagel–Xu]	**+**		**+** (excess)		**+**
Retail investors					
Vanguard investors [Giglio et al.]	**–**		**+** (excess)		**+**
Gallup survey [Greenwood–Shleifer]	**–**				

Table 2
Option-Based Expectations: Mincer-Zarnowitz Regressions

This table reports Mincer-Zarnowitz regressions of future realized spot rates on current forward rates for option-based risk premia (left panel) and expected returns (right panel). For risk premia, the realized spot rate

$$\tilde{\mu}_{t+6}^{(6)} = \mathbb{E}_{t+6}\left[r_{t+6}^{(6)}\right]$$

is the future expectation of the 6-month equity premium, and the forward rate

$$\tilde{f}_t^{(6,6)} = \mathbb{E}_t\left[r_{t+6}^{(6)}\right]$$

is the current expectation of the same risk premium. These expectations are analogously defined for expected returns. The units are annualized percentage points. Panel regressions, in the main sample, report standard errors clustered by exchange and date. This sample is the longest available for each exchange. Time-series regressions, in the U.S. sample, report Newey-West standard errors with $L = \lceil 1.3 \times T^{1/2} \rceil$ lags and fixed-b p-values, following Lazarus et al. (2018). This sample is from 01/1990 to 06/2021 for forward rates and from 07/1990 to 12/2021 for realized spot rates.

	Risk Premia $\tilde{\mu}_{t+6}^{(6)}$			Expected Returns $\mu_{t+6}^{(6)}$		
	(1) Main	(2) Main	(3) U.S. Only	(4) Main	(5) Main	(6) U.S. Only
$\tilde{f}_t^{(6,6)}$	0.64	0.56	0.67			
	(0.049)	(0.055)	(0.096)			
$f_t^{(6,6)}$				0.92	0.91	0.88
				(0.053)	(0.061)	(0.073)
Intercept	1.04		0.74	0.095		0.27
	(0.17)		(0.28)	(0.22)		(0.34)
p-val: $\beta_1 = 1$	<0.001	<0.001	0.003	0.157	0.181	0.114
p-val: $\beta_0 = 0$	<0.001	-	0.019	0.675	-	0.446
p-val: $\beta_1 = 1$, $\beta_0 = 0$	<0.001	-	0.012	0.067	-	0.110
Observations	2227	2227	378	2227	2227	378
Fixed Effects	None	Ex	None	None	Ex	None
Standard Errors	Cluster	Cluster	Newey-West	Cluster	Cluster	Newey-West
Clusters/Lags	Ex/Date	Ex/Date	26	Ex/Date	Ex/Date	26
Adjusted R^2	0.20	0.22	0.22	0.63	0.63	0.71
Within R^2	-	0.15	-	-	0.59	-

Table 3
Option-Based Expectations: Average Forecast Errors

This table reports average forecast errors for option-based risk premia (left panel) and expected returns (right panel). For risk premia, the forecast error is the difference between the future realized spot rate and the current forward rate. The realized spot rate is the future expectation of the 6-month equity premium, and the forward rate is the current expectation of the same risk premium. These expectations are analogously defined for expected returns. The units are annualized percentage points. Panel regressions, in the main sample, report standard errors clustered by exchange and date. This sample is the longest available for each exchange. Time-series regressions, in the U.S. sample, report Newey-West standard errors with $L = \lceil 1.3 \times T^{1/2} \rceil$ lags and fixed-b p-values, following Lazarus et al. (2018). This sample is from 01/1990 to 06/2021 for forward rates and from 07/1990 to 12/2021 for realized spot rates.

	Risk Premia $\tilde{\varepsilon}^{(6)}_{t+6}$		Expected Returns $\varepsilon^{(6)}_{t+6}$	
	(1) Main	(2) U.S. Only	(3) Main	(4) U.S. Only
Average Error	0.17	0.021	-0.28	-0.40
	(0.11)	(0.15)	(0.10)	(0.17)
p-val: $\bar{\varepsilon}_t = 0$	0.136	0.891	0.026	0.035
Observations	2227	378	2227	378
Standard Errors	Cluster	Newey-West	Cluster	Newey-West
Clusters/Lags	Ex/Date	26	Ex/Date	26

Table 4
Option-Based Expectations: Predictability of Forecast Errors

This table reports predictability regressions of future realized forecast errors on current forward rates for option-based risk premia (left panel) and expected returns (right panel). For risk premia, the realized spot rate is the future expectation of the 6-month equity premium, the forward rate is the current expectation of the same risk premium, and the forecast error is the realized spot rate minus the forward rate. These expectations are analogously defined for expected returns. For both risk premia and expected returns, the predictor is the 2×1-month forward risk premium, the current expectation of the 1-month equity premium in 2 months. The units are annualized percentage points. Panel regressions, in the main sample, report standard errors clustered by exchange and date. This sample is the longest available for each exchange. Time-series regressions, in the U.S. sample, report Newey-West standard errors with $L = \lceil 1.3 \times T^{1/2} \rceil$ lags and fixed-b p-values, following Lazarus et al. (2018). This sample is from 01/1990 to 06/2021 for forward rates and from 07/1990 to 12/2021 for realized spot rates.

	Risk Premia $\tilde{\varepsilon}_{t+6}^{(6)}$			Expected Returns $\varepsilon_{t+6}^{(6)}$		
	(1) Main	(2) Main	(3) U.S. Only	(4) Main	(5) Main	(6) U.S. Only
$\tilde{f}_t^{(2,1)}$	-0.13	-0.16	-0.17	-0.25	-0.29	-0.20
	(0.047)	(0.047)	(0.067)	(0.050)	(0.047)	(0.091)
Intercept	0.51		0.39	0.39		0.039
	(0.15)		(0.23)	(0.17)		(0.29)
p-val: $\beta_1 = 0$	0.025	0.008	0.025	0.001	$<$0.001	0.050
Observations	2227	2227	378	2227	2227	378
Fixed Effects	None	Ex	None	None	Ex	None
Standard Errors	Cluster	Cluster	Newey-West	Cluster	Cluster	Newey-West
Clusters/Lags	Ex/Date	Ex/Date	26	Ex/Date	Ex/Date	26
Adjusted R^2	0.02	0.03	0.03	0.08	0.11	0.04
Within R^2	-	0.03	-	-	0.10	-

Table 5
Survey-Based Regressions

This table reports regressions for Livingston (top panel), CFO (middle panel), and Vanguard (bottom panel) survey expectations. Columns (1)–(2) are Mincer-Zarnowitz regressions of future realized spot rates on current forward rates and test $H_0\colon \beta_1 = 1$, $H_0\colon \beta_0 = 0$, and $H_0\colon \beta_1 = 1,\ \beta_0 = 0$, as in Table 2. Columns (3)–(4) are average forecast errors and test $H_0\colon \bar{\varepsilon}_t = 0$, as in Table 3. Columns (5)–(6) are error-predictability regressions of future realized forecast errors on current forward rates and test $H_0\colon \beta_1 = 0$, as in Table 4. The Livingston horizon is the 6×6-month forecast error, the CFO horizon is the 1×9-year forecast error, and the Vanguard horizon is the 1×9-year forecast error. The units are annualized percentage points. Each regression reports Newey-West standard errors with $L = \lceil 1.3 \times T^{1/2} \rceil$ lags and fixed-b p-values, following Lazarus et al. (2018). The Livingston sample is half-yearly from 06/1992 to 06/2021 for forward rates and from 12/1992 to 12/2021 for realized spot rates. The CFO sample is quarterly from 12/2001 to 06/2024 (excluding 09/2019 and 12/2019) for forward rates and from 12/2002 to 06/2025 (excluding 09/2020 and 12/2020) for realized spot rates. Each regression includes an indicator variable that equals 1 on and after 09/2020 for forward rates and likewise 09/2021 for realized spot rates. The Vanguard sample is bimonthly from 02/2017 to 12/2023 for forward rates and from 02/2018 to 12/2024 for realized spot rates.

	Panel A. Livingston Survey					
	Mincer-Zarnowitz		Average Error		Error Predictability	
	(1) RP $\tilde{\mu}^{(6)}_{t+6}$	(2) ER $\mu^{(6)}_{t+6}$	(3) RP $\tilde{\varepsilon}^{(6)}_{t+6}$	(4) ER $\varepsilon^{(6)}_{t+6}$	(5) RP $\tilde{\varepsilon}^{(6)}_{t+6}$	(6) ER $\varepsilon^{(6)}_{t+6}$
$\tilde{f}^{(6,6)}_t$	0.81				-0.19	-0.19
	(0.077)				(0.077)	(0.073)
$f^{(6,6)}_t$		0.68				
		(0.069)				
Intercept	0.97	1.95	0.37	0.11	0.97	0.72
	(0.64)	(0.70)	(0.34)	(0.30)	(0.64)	(0.59)
p-val: Slope	0.052	0.002	-	-	0.052	0.040
p-val: Intercept	0.205	0.032	0.347	0.733	-	-
p-val: Joint	0.125	0.005	-	-	-	-
Observations	59	59	59	59	59	59
Newey-West Lags	10	10	10	10	10	10
Adjusted R^2	0.56	0.38	-	-	0.05	0.06

	Panel B. CFO Survey					
	Mincer-Zarnowitz		Average Error		Error Predictability	
	(1) RP $\tilde{\mu}^{(9y)}_{t+1y}$	(2) ER $\mu^{(9y)}_{t+1y}$	(3) RP $\tilde{\varepsilon}^{(9y)}_{t+1y}$	(4) ER $\varepsilon^{(9y)}_{t+1y}$	(5) RP $\tilde{\varepsilon}^{(9y)}_{t+1y}$	(6) ER $\varepsilon^{(9y)}_{t+1y}$
$\tilde{f}^{(1y,9y)}_t$	0.35				-0.65	-0.18
	(0.14)				(0.14)	(0.083)
$f^{(1y,9y)}_t$		0.64				
		(0.076)				
Intercept	2.27	2.22	0.096	-0.24	2.27	0.37
	(0.50)	(0.49)	(0.072)	(0.095)	(0.50)	(0.27)
p-val: Slope	0.001	0.001	-	-	0.001	0.069
p-val: Intercept	0.001	0.001	0.245	0.036	-	-
p-val: Joint	0.005	0.004	-	-	-	-
Observations	84	84	84	84	84	84
Newey-West Lags	12	12	12	12	12	12
Adjusted R^2	0.59	0.83	-	-	0.48	0.04

(Continued on the next page)

Table 5
Survey-Based Regressions (Continued)

	Panel C. Vanguard Survey					
	Mincer-Zarnowitz		Average Error		Error Predictability	
	(1) RP $\tilde{\mu}^{(9y)}_{t+1y}$	(2) ER $\mu^{(9y)}_{t+1y}$	(3) RP $\tilde{\varepsilon}^{(9y)}_{t+1y}$	(4) ER $\varepsilon^{(9y)}_{t+1y}$	(5) RP $\tilde{\varepsilon}^{(9y)}_{t+1y}$	(6) ER $\varepsilon^{(9y)}_{t+1y}$
$\tilde{f}^{(1y,9y)}_t$	0.43				-0.57	-0.089
	(0.16)				(0.16)	(0.069)
$f^{(1y,9y)}_t$		0.54				
		(0.15)				
Intercept	2.20	3.03	-0.32	-0.10	2.20	0.29
	(0.82)	(1.08)	(0.40)	(0.070)	(0.82)	(0.34)
p-val: Slope	0.016	0.033	-	-	0.016	0.289
p-val: Intercept	0.048	0.040	0.497	0.237	-	-
p-val: Joint	0.061	0.079	-	-	-	-
Observations	42	42	42	42	42	42
Newey-West Lags	9	9	9	9	9	9
Adjusted R^2	0.13	0.26	-	-	0.21	0.03

(Continued from the previous page)

Table 6
Price Impact of Forecast Errors in Crises

This table reports the price impact of forecast errors in the financial crisis (top panel) and the Covid-19 recession (bottom panel) for option-based and survey-based expectations. For option-based expectations, the price impact is the discounted sum of predicted forecast errors from $j = 1$ to $j = \infty$ under the assumption of a flat term structure, as in (12). For survey-based expectations, the price impact is the discounted sum of predicted forecast errors from $j = 1$ to $j = 9$, as in (13). The predicted forecast error is from a time-series regression of future realized forecast errors on current forward rates. Forecast errors and forward rates are for risk premia. The option-based sample is monthly from 01/1990 to 06/2021 for forward rates and from 07/1990 to 12/2021 for realized spot rates. The CFO sample is quarterly from 12/2001 to 12/2018 for forward rates and from 12/2002 to 12/2019 for realized spot rates. The Vanguard sample is bimonthly from 02/2017 to 12/2023 for forward rates and from 02/2018 to 12/2024 for realized spot rates. $p_t - d_t$ is the log repurchase-adjusted price-dividend ratio in the U.S. sample and is obtained from Nagel and Xu (2022) via Zhengyang Xu's website.

	Change in Discounted Risk Premia Forecast Errors			
		Survey-Based Expectations		
	Option-Based (Flat Term Structure)	CFO Survey (Lower Bound)	Vanguard Survey (Lower Bound)	Change in $p_t - d_t$
2008 Financial Crisis				
Peak to Trough: 06/2007 to 03/2009	-0.3576	-0.1127	-	-0.4273
2020 Covid-19 Recession				
Peak to Trough: 12/2019 to 03/2020	-0.2183	-	-0.0940	-0.2060

Figure 2
Option-Based Spot and Forward Rates

The top panel plots contemporaneous 6-month spot rates $\tilde{\mu}_t^{(6)}$ (light blue) and 6×6-month forward rates $\tilde{f}_t^{(6,6)}$ (dark blue). The bottom panel plots instrumented 6×6-month forward rates $\tilde{f}_t^{(6,6)}$ (dark blue) and the corresponding realized 6-month spot rates $\tilde{\mu}_{t+6}^{(6)}$ (red). The instrument is the 2×1-month forward rate, as in Table A3. Spot and forward rates are for risk premia. Gray bands are NBER recessions. The sample is from 01/1990 to 06/2021 in the U.S.

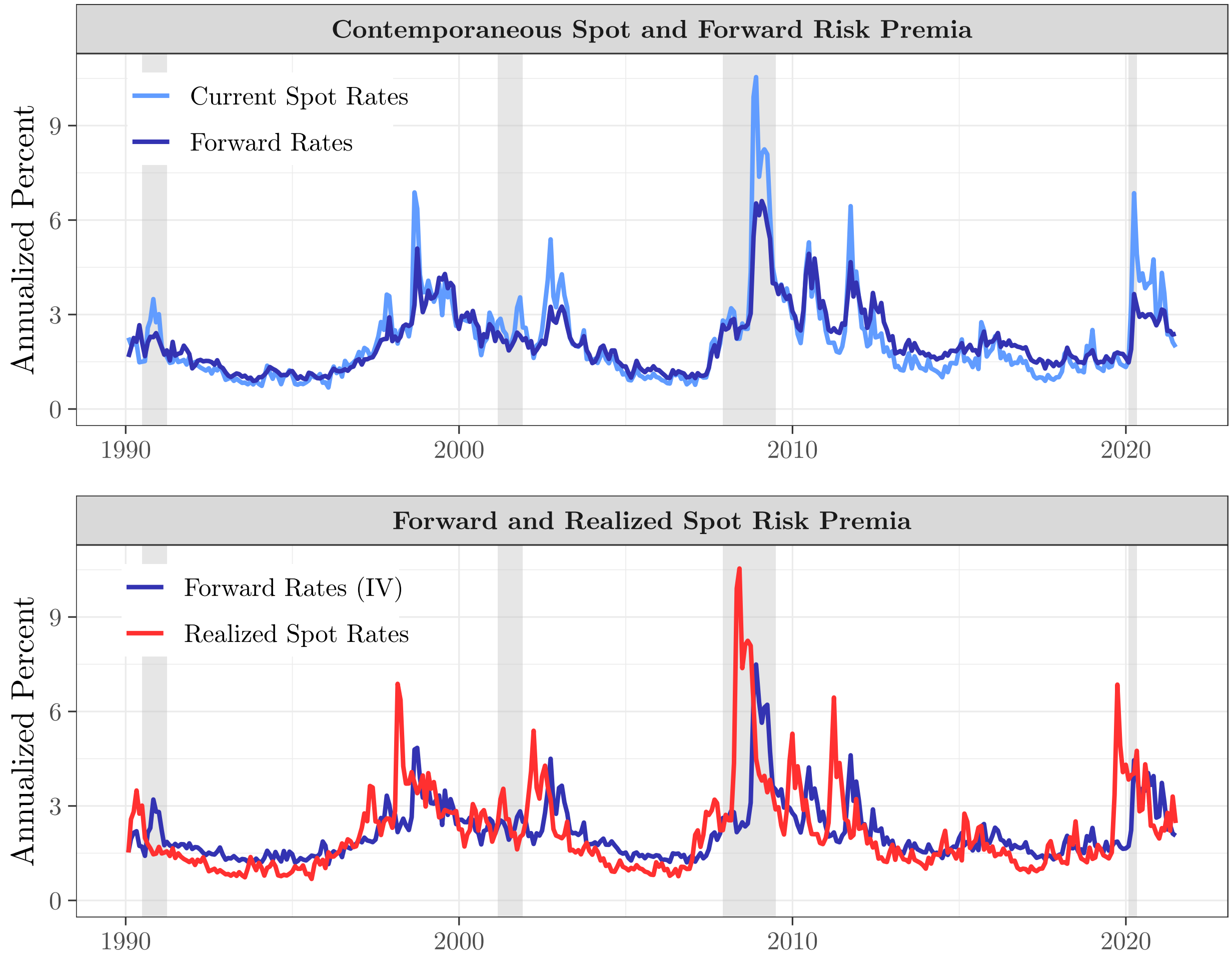

Figure 3
Option-Based Forward Rates and Realized Forecast Errors

This figure plots 6×6-month forward rates $\tilde{f}_t^{(6,6)}$ (dark blue line) and the corresponding realized forecast errors $\tilde{\varepsilon}_{t+6}^{(6)}$ (orange bars). Forward rates and forecast errors are for risk premia. The sample is from 01/1990 to 06/2021 in the U.S.

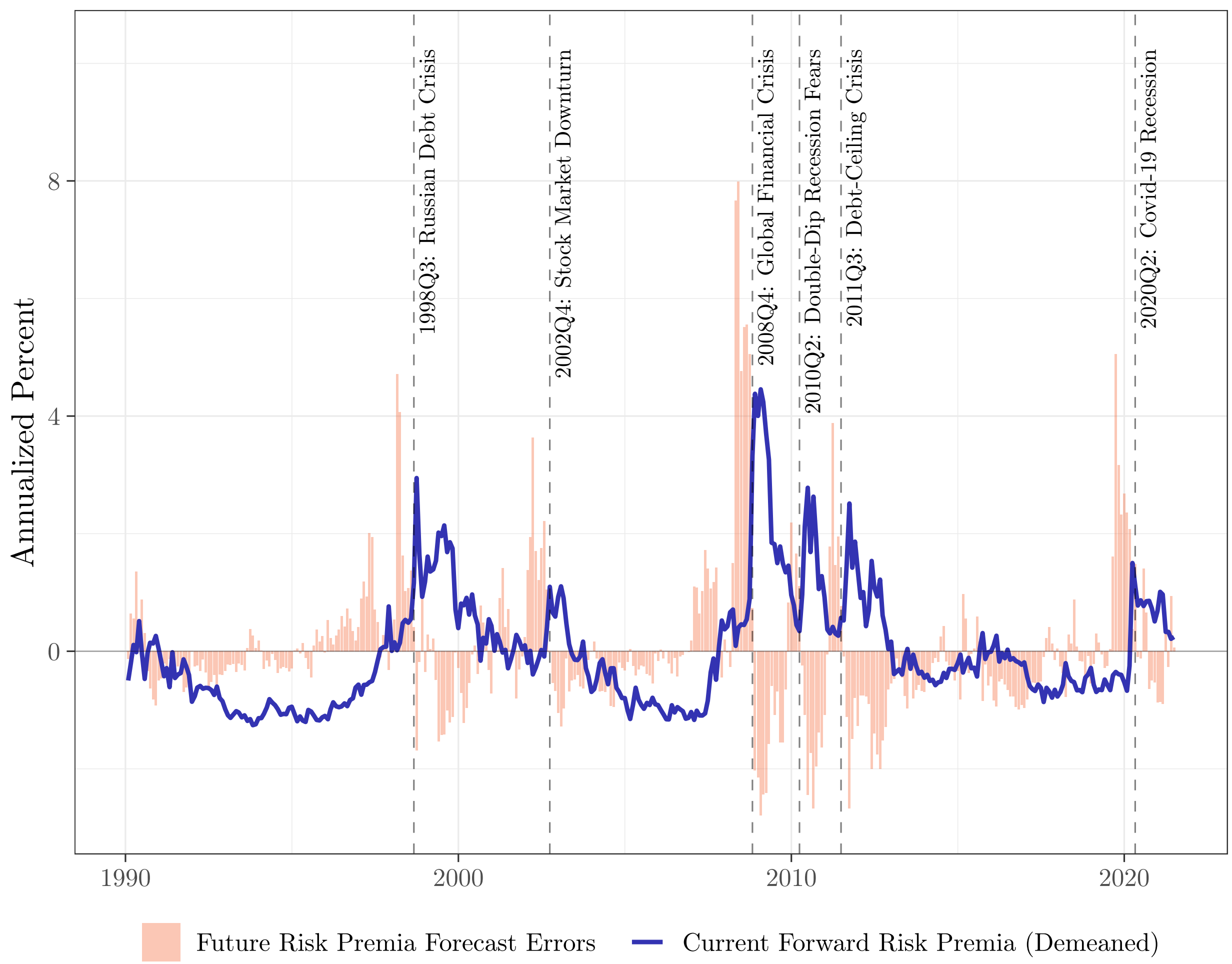

Figure 4
Survey-Based Forward Rates

This figure plots forward rates for Livingston (top panel), CFO (middle panel), and Vanguard (bottom panel) survey expectations. Forward rates are for risk premia. Gray bands are NBER recessions. The Livingston sample is half-yearly from 06/1992 to 12/2021. The CFO sample is quarterly from 12/2001 to 06/2025 and is adjusted for the structural break in 2020. The break adjustment is from a time-series regression of unadjusted forward rates on an indicator variable that equals 1 on and after 09/2020, as in Figure A9. The Vanguard sample is bimonthly from 02/2017 to 12/2024.

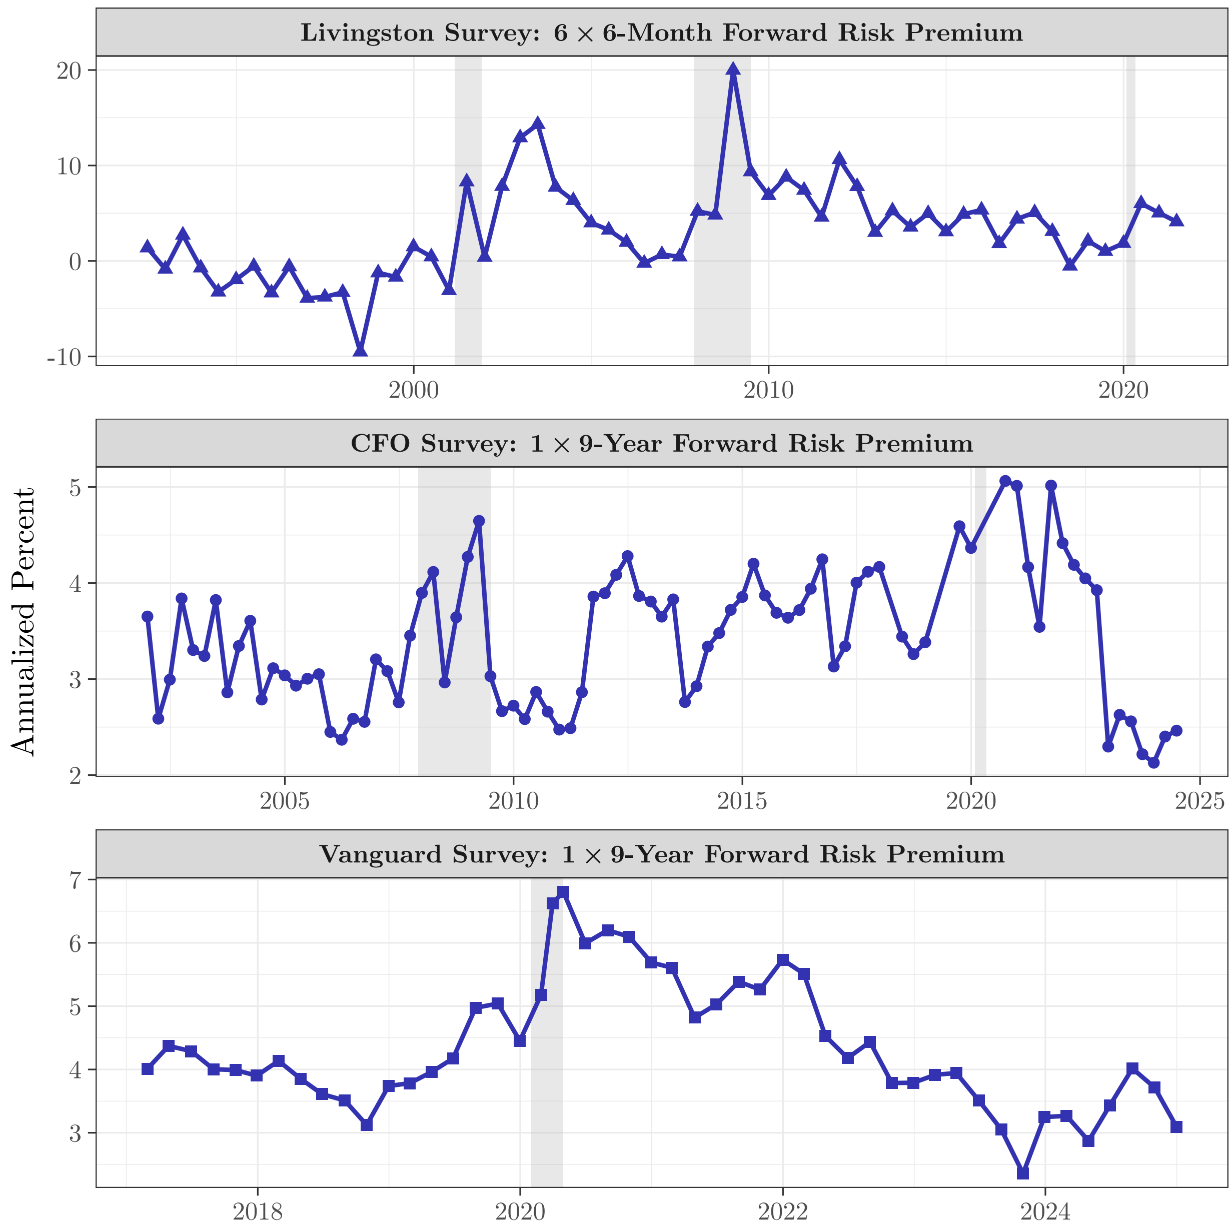

Figure 5
Cyclical Variation in Survey-Based One-Year Spot Rates

This figure plots 1-year spot rates for CFO (top panel) and Vanguard (bottom panel) survey expectations. Spot rates are for risk premia. Gray bands are NBER recessions. The CFO sample is quarterly from 12/2001 to 06/2025. The Vanguard sample is bimonthly from 02/2017 to 12/2024.

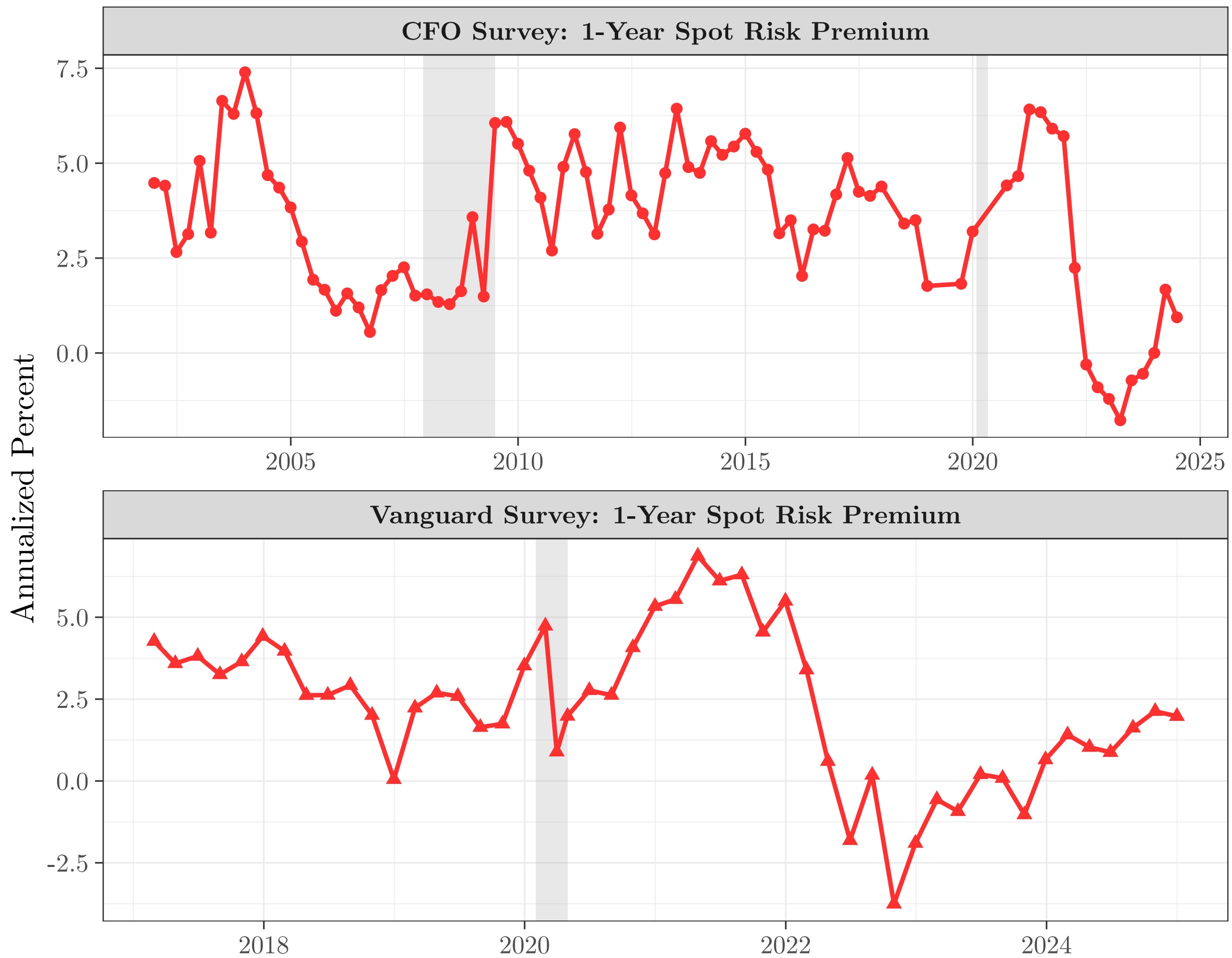

Figure 6
Cyclical Variation in Forward Risk Premia

This figure reports regressions of contemporaneous forward rates on state variables for option-based (first column) and survey-based (second to fourth columns) expectations for risk premia. The bar length reports the slope t-stat, capped at $t = 3.00$. The bar label reports the uncapped t-stat and the statistical significance. On the left-hand side, the option-based horizon is the 6×6-month forward rate, the Livingston horizon is the 6×6-month forward rate, the CFO horizon is the 1×9-year forward rate, and the Vanguard horizon is the 1×9-year forward rate. On the right-hand side, 1/ExCAPE is the excess cyclically-adjusted earnings yield (obtained from Robert Shiller's website), SVIX^2 is the option-implied risk premium from Martin (2017), and NBER is a recession indicator. Each state variable is signed such that higher values correspond to bad times (measured as a positive covariance with the NBER recession indicator), so positive signs in the figure indicate countercyclical forward risk premia. Each t-stat is calculated using Newey-West standard errors with $L = \lceil 1.3 \times T^{1/2} \rceil$ lags, following Lazarus et al. (2018), and significance stars are based on two-sided tests with fixed-b critical values (*** for 1%, ** for 5%, * for 10%). The option-based sample is monthly from 01/1990 to 12/2021. The Livingston sample is half-yearly from 06/1992 to 12/2021. The CFO sample is quarterly from 12/2001 to 06/2025; each regression includes an indicator variable that equals 1 on and after 09/2020. The Vanguard sample is bimonthly from 02/2017 to 12/2024. See Section 4.3 for more details.

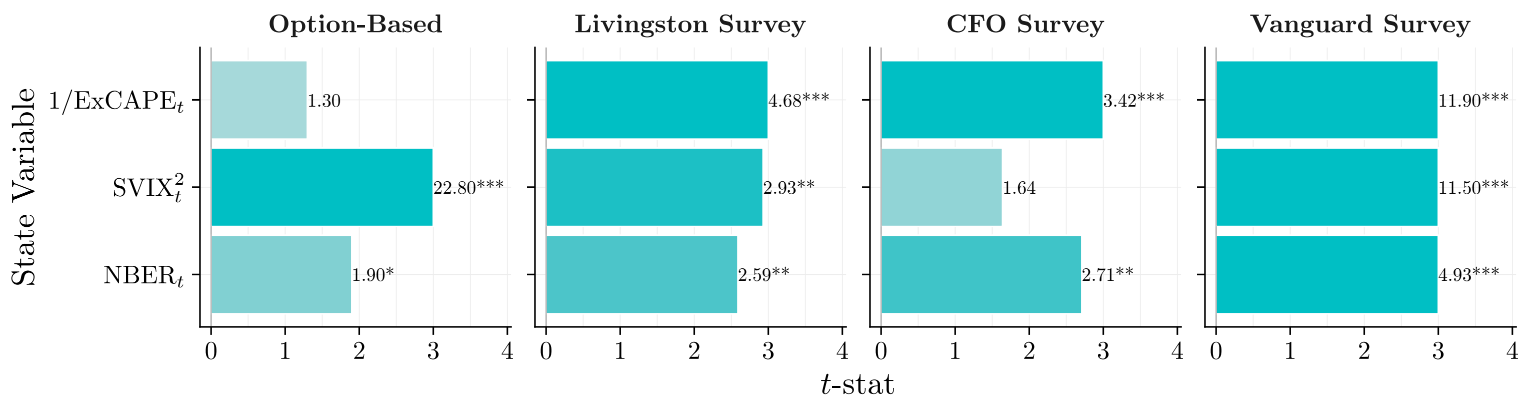

Figure 7
Price Impact of Forecast Errors in Crises

The figure plots the price impact of forecast errors in the financial crisis (top panel) and the Covid-19 recession (bottom panel) for option-based (left axis) and survey-based (right axis) expectations. For option-based expectations, the price impact is the discounted sum of predicted forecast errors from $j = 1$ to $j = \infty$ under the assumption of a flat term structure, as in (12). For survey-based expectations, the price impact is the discounted sum of predicted forecast errors from $j = 1$ to $j = 9$, as in (13). The predicted forecast error is from a time-series regression of future realized forecast errors on current forward rates. Forecast errors and forward rates are for risk premia. The option-based sample is monthly from 01/1990 to 06/2021 for forward rates and from 07/1990 to 12/2021 for realized spot rates. The CFO sample is quarterly from 12/2001 to 12/2018 for forward rates and from 12/2002 to 12/2019 for realized spot rates. The Vanguard sample is bimonthly from 02/2017 to 12/2023 for forward rates and from 02/2018 to 12/2024 for realized spot rates.

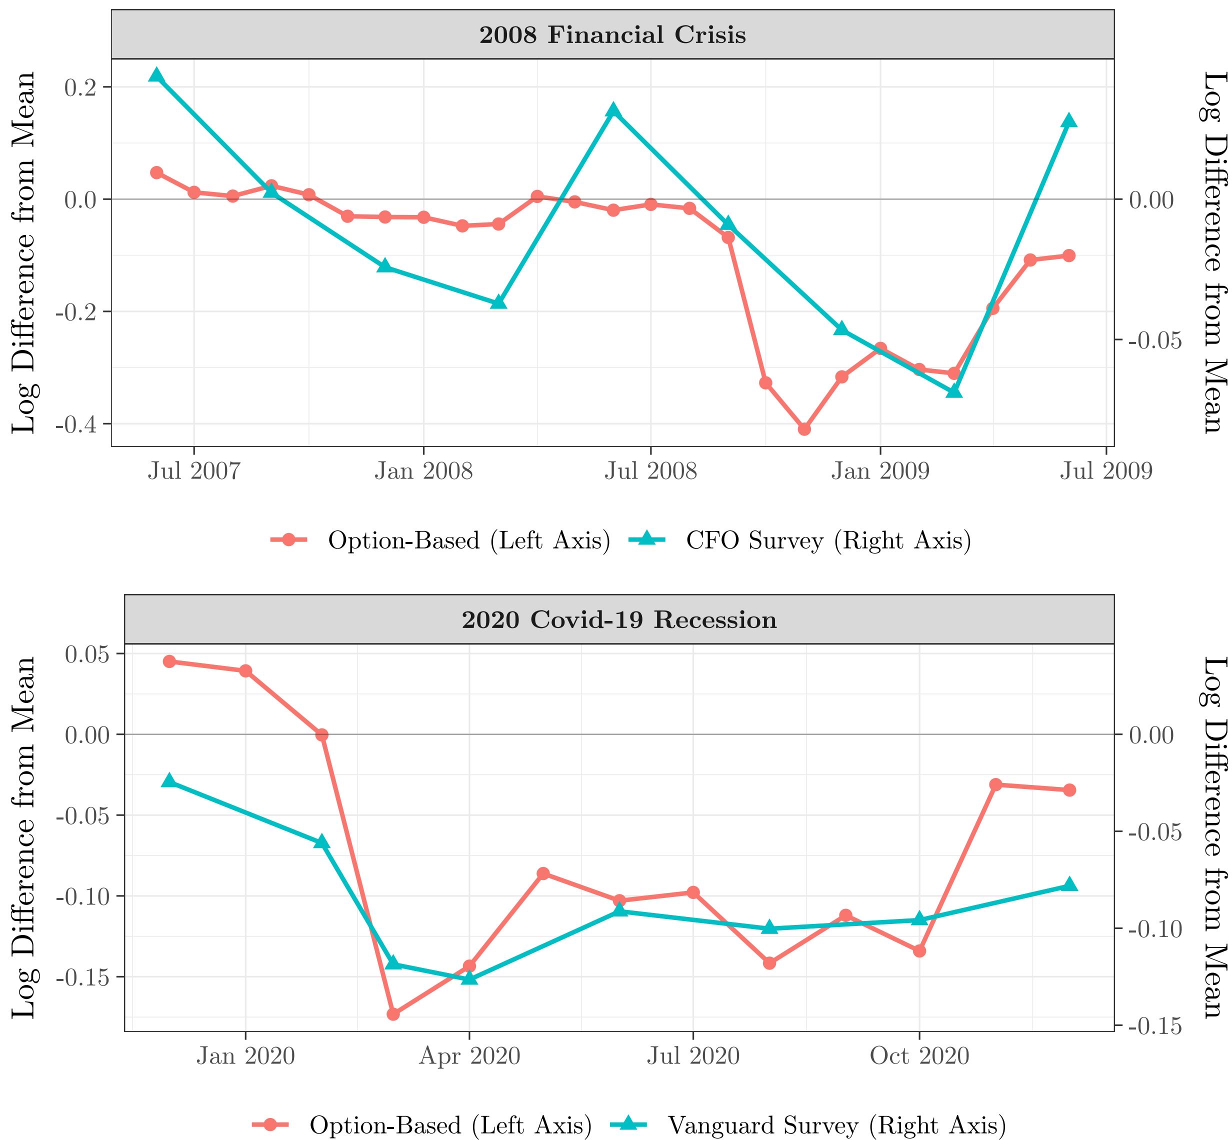

References

Adam, Klaus, Albert Marcet, and Johannes Beutel. 2017. "Stock Price Booms and Expected Capital Gains." *American Economic Review* 107 (8):2352–2408.

Andrews, Donald W. K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica* 59 (3):817–858.

Binsbergen, Jules H. van and Ralph S.J. Koijen. 2017. "The term structure of returns: Facts and theory." *Journal of Financial Economics* 124 (1):1–21.

Binsbergen, Jules van, Michael Brandt, and Ralph Koijen. 2012. "On the Timing and Pricing of Dividends." *American Economic Review* 102 (4):1596–1618.

Breeden, Douglas T. and Robert H. Litzenberger. 1978. "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business* 51 (4):621–651.

Campbell, John Y. and Robert J. Shiller. 1988. "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors." *Review of Financial Studies* 1 (3):195–228.

Carr, Peter and Dilip Madan. 1998. "Towards a Theory of Volatility Trading." In *Volatility: New Estimation Techniques for Pricing Derivatives*. Cambridge University Press, 417–427.

Dahlquist, Magnus and Markus Ibert. 2024. "Equity Return Expectations and Portfolios: Evidence from Large Asset Managers." *Review of Financial Studies* 37 (6):1887–1928.

———. 2026. "Institutions' return expectations across assets and time." *Journal of Financial Economics* 175:104188.

De la O, Ricardo and Sean Myers. 2021. "Subjective Cash Flow and Discount Rate Expectations." *Journal of Finance* 76 (3):1339–1387.

Diebold, Francis X., Canlin Li, and Vivian Z. Yue. 2008. "Global yield curve dynamics and interactions: A dynamic Nelson–Siegel approach." *Journal of Econometrics* 146 (2):351–363.

Gabaix, Xavier and Ralph S.J. Koijen. 2022. "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis." Working Paper 28967, National Bureau of Economic Research.

Gao, Can and Ian W.R. Martin. 2021. "Volatility, Valuation Ratios, and Bubbles: An Empirical Measure of Market Sentiment." *Journal of Finance* 76 (6):3211–3254.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus. 2020. "Five facts about beliefs and portfolios." *American Economic Review* 111 (5):1481–1522.

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

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2021. "The joint dynamics of investor beliefs and trading during the COVID-19 crash." *Proceedings of the National Academy of Sciences* 118 (4):e2010316118.

Golez, Benjamin and Jens Jackwerth. 2024. "Holding Period Effects in Dividend Strip Returns." *Review of Financial Studies* 37 (10):3188–3215.

Gormsen, Niels Joachim. 2021. "Time Variation of the Equity Term Structure." *Journal of Finance* 76 (4):1959–1999.

Gormsen, Niels Joachim and Eben Lazarus. 2026. "Interest Rates and Equity Valuations." Working Paper 34814, National Bureau of Economic Research.

Graham, John, Brent Meyer, Nicholas Parker, and Sonya Ravindranath Waddell. 2020. "Introducing The CFO Survey: A collaboration between Duke University's Fuqua School of Business and the Federal Reserve Banks of Richmond and Atlanta." URL https://www.richmondfed.org/research/national_economy/cfo_survey/research_and_commentary/2020/20200515_introducing_cfo_survey.

Greenwood, Robin and Andrei Shleifer. 2014. "Expectations of Returns and Expected Returns." *Review of Financial Studies* 27 (3):714–746.

Hirshleifer, David, Jun Li, and Jianfeng Yu. 2015. "Asset pricing in production economies with extrapolative expectations." *Journal of Monetary Economics* 76:87–106.

Lazarus, Eben, Daniel J. Lewis, and James H. Stock. 2021. "The Size-Power Tradeoff in HAR Inference." *Econometrica* 89 (5):2497–2516.

Lazarus, Eben, Daniel J. Lewis, James H. Stock, and Mark W. Watson. 2018. "HAR Inference: Recommendations for Practice." *Journal of Business & Economic Statistics* 36 (4):541–559.

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

Ludvigson, Sydney C. and Serena Ng. 2009. "Macro Factors in Bond Risk Premia." *Review of Financial Studies* 22 (12):5027–5067.

Martin, Ian. 2017. "What is the Expected Return on the Market?" *Quarterly Journal of Economics* 132 (1):367–433.

Martin, Ian W.R. and Dimitris Papadimitriou. 2022. "Sentiment and Speculation in a Market with Heterogeneous Beliefs." *American Economic Review* 112 (8):2465–2517.

Mincer, Jacob A. and Victor Zarnowitz. 1969. "The Evaluation of Economic Forecasts." In *Economic Forecasts and Expectations*. National Bureau of Economic Research, 3–46.

Nagel, Stefan and Zhengyang Xu. 2022. "Asset Pricing with Fading Memory." *Review of Financial Studies* 35 (5):2190–2245.

———. 2023. "Dynamics of subjective risk premia." *Journal of Financial Economics* 150 (2):103713.

Newey, Whitney K. and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3):703–708.

Shiller, Robert J., Laurence Black, and Farouk Jivraj. 2020. "CAPE and the COVID-19 Pandemic Effect." Technical report, Barclays Research and SSRN. URL https://ssrn.com/abstract=3714737.

Vanguard. 2025. "Investor Pulse: Both optimism and uncertainty about 2025." URL https://corporate.vanguard.com/content/corporatesite/us/en/corp/articles/investor-pulse-mix-optimism-uncertainty-about-2025.html.