Forward Return Expectations[*]

Mihir Gandhi, Niels Joachim Gormsen, and Eben Lazarus

September 2023

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*. While contemporaneous short-term expectations are not countercyclical across all data sources, we find that forward expectations are consistently countercyclical, and excessively so: in bad times, forward expectations are higher than justified by investors' own subsequent short-term return expectations. This excess volatility in forward expectations helps account for excess volatility in prices, inelastic demand for equities, and stylized facts about the equity term structure.

Keywords: asset pricing, expected stock returns, time-varying discount rates

JEL classification: G10, G12, G40

1   Introduction

One of the organizing facts in asset pricing is that expected returns on equities vary over time. A key question is how investors perceive this variation ex ante. One aspect of this question, which has been studied extensively, is how investors perceive the contemporaneous level of the equity premium. Another aspect, which is the focus of this paper, is how investors perceive the term structure dynamics of expected returns. In particular, what do investors, at a given point in time, expect the equity premium to be in the future? Studying such forward expectations of returns is important for understanding stock-price formation — prices depend on expectations about all future equity premia — and it helps strengthen our understanding of the dynamics of expectations.

We study forward return expectations using both option prices and survey data. First, using a global sample of equity index options, we extract the equity premium at multiple horizons as perceived by the "representative agent," by which we mean any investor willing to hold the market portfolio. Once we have long- and short-term expectations, we can construct forward expectations of returns, as detailed below. We provide theoretical results showing that we can identify time-series changes in these expectations, which is our main focus, under much weaker assumptions than those usually used to estimate the level of expected returns at any given time. The option prices thus provide a powerful lab for testing the intertemporal consistency of return expectations in a rich data set covering 20 equity indices over 30 years. We complement the options data with return expectations from surveys of professional forecasters and CFOs. These surveys elicit return expectations at multiple horizons, again allowing us to extract forward expectations, but in a more model-free setting.

Our main empirical finding is that forward expectations are excessively countercyclical: in bad times, investors believe that expected returns will be highly elevated in the future, but these forward 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. This consistency contrasts with results on contemporaneous short-run return expectations; in some surveys, contemporaneous short-run return expectations are not countercyclical, but we consistently find that forward expectations are not just countercyclical but excessively so.

The excess cyclicality in forward expectations represents a new source of excess volatility

in stock prices. One interpretation of our findings is that investors believe that negative asset-pricing shocks will be long-lasting, with risk premia accordingly elevated for a long period after the shock. Soon after, however, investors realize that the bad times will be shorter-lived than initially expected, and accordingly revise their beliefs about expected returns downward. If investors did not initially overestimate the persistence of bad times, and thus future forward rates, cyclical fluctuations in prices would be more modest; we show, for instance, that predictable risk-premium forecast errors can account for nearly half of the decrease in prices during the global financial crisis, and nearly all of the decrease during the Covid crash.

Our findings on forward expectations also help explain a number of facts documented in recent work, related to (i) stylized facts about the equity term structure, (ii) investors' inelastic demand for equities in response to a given change in prices, and (iii) disagreements in previous research related to survey expectations of returns. We explore these implications further below.

1.1 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*. 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:

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

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 forecasts 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 two well-known surveys in the finance literature, namely the Livingston survey of professional forecasters and the Duke CFO survey of financial executives. In the Livingston survey, for instance, we observe expected returns at the 6-month and 12-month horizon. These expectations 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 forecasts errors at many different horizons, in a relatively long sample that spans 20 countries.

We present the main stylized facts about the behavior of forward rates and forecast errors in Table 1. First, as shown in Panel A, forward rates are decent predictors of future realized spot rates. Across the three data sources, future realized spot rates increase by around 0.7 percentage points when ex ante forward rates increase by 1 percentage point. The R^2 value varies between 0.38 and 0.71. In more detailed analysis, we find that forward rates are useful for predicting both transitory and more persistent fluctuations in spot rates. Taken together, the predictive regressions suggest that forward rates have predictive power over future spot rates, but that future spot rates move by significantly less than predicted by forward rates.

As such, investors appear to make predictable forecast errors in expected returns. When forward rates are high, subsequent spot rates are lower than expected by investors ex ante. This result can be seen in Panel B, which reports results of predictive regressions of realized forecasts errors on ex ante forward rates: spot rates are around 0.25 percentage points lower than expected when forward rates are 1 percentage point higher.[1] These results are statistically significant, and the R^2 values indicate meaningful economic predictability. This pattern reflects one of the key results of the paper, namely that there is excess cyclicality in return expectations: in bad times, investors' beliefs about future expected returns (forward

---

[1]There are slight differences in the regression specifications in Panels A and B, as explained further below.

Table 1

Stylized Facts about Forward Rates and Forecasts Errors

This table reports stylized facts about implied equity forward rates measured in three different ways. The first measure is the 6-month, 6-month forward rate estimated in our global sample of option prices (1990-2021); the second measure is the U.S. 6-month, 6-month forward rate from the Livingston survey of professional forecasters (1992-2022); the third measure is the U.S. 1-year, 9-year forward rate from the Duke CFO Survey (2004-2021). The forecast error $\varepsilon_{i,t+1}$, where i indexes the data set, is the difference between the realized spot rate at $t+1$ and the ex ante forward rate. See the remainder of the paper for details.

| | Expectations Measured by: | | |
	Options	Livingston Survey	CFO Survey
Panel A. Predictability in Spot Rates ($\mu_{i,t+1} = \beta_0 + \beta_1 f_{i,t} + e_{i,t+1}$)			
β_1	0.88	0.68	0.63
R^2	0.71	0.38	0.46
Panel B. Predictability of Forecast Errors ($\varepsilon_{i,t+1} = \beta_0 + \beta_1 f_{i,t} + e_{i,t+1}$)			
β_1	-0.34	-0.19	-0.15
R^2	0.06	0.06	0.03
Panel C. Cyclical Variation in Forward Rates and Forecast Errors			
$\rho\left(f_{i,t}, 1/\mathrm{CAPE}_t\right)$	0.04	0.42	0.21
$\rho\left(\varepsilon_{i,t+1}, 1/\mathrm{CAPE}_t\right)$	-0.38	-0.19	-0.38

rates) are elevated by significantly more than justified by their subsequent beliefs.[2]

To illustrate this excess cyclicality, Figure 1 shows forward rates and realized spot rates from the option-based measure during three crises: the 1998 Russian debt crisis, the 2008 global financial crisis, and the early-2020 Covid 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 orange triangles show ex post realized one-month spot rates as of months t, $t+1$, ..., $t+5$. In all cases, the spot equity premium (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. This suggests that while investors understood that the equity premium

---

[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 Rates and Realized Spot Rates in Three Crises

This figure plots the forward curve $f_t^{(0)}$, $f_t^{(1)}$, $\cdots$, $f_t^{(5)}$ (dark blue) and the corresponding ex post spot rates $\mu_t^{(1)}$, $\mu_{t+1}^{(1)}$, $\cdots$, $\mu_{t+5}^{(1)}$ (red) in the U.S. as of three dates: August 31, 1998 for the Russian debt crisis, November 28, 2008 for the financial crisis, and March 31, 2020 for the Covid-19 recession.
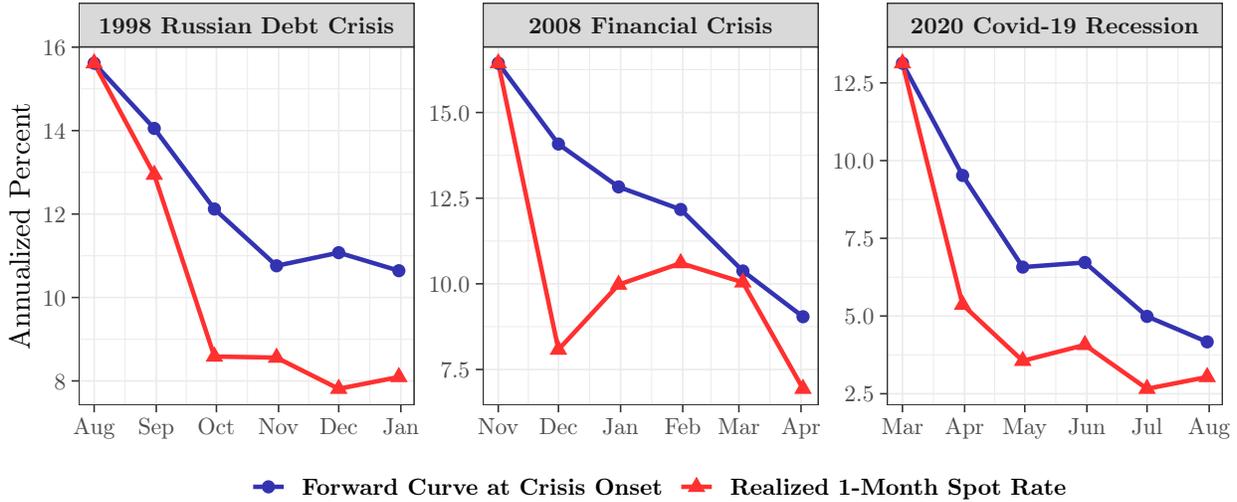
would decrease in the future, they appear to have underestimated the speed of mean reversion. One interpretation of this finding is that the rebound following each crisis was driven by news, like a series of policy shocks, that was unexpected ex ante. Another interpretation is that investors systematically overreact to crises by thinking that future spot rates will be elevated by more and for longer than what one should rationally expect.

While Figure 1 focuses on the option-based measure, we find consistent evidence of excess cyclicality in all three measures of expectations we consider. Moreover, the forward rates share a common, cyclical component across data sources. As shown in Panel C of Table 1, forward rates are higher when the earnings yield (measured as the inverse of the CAPE ratio) is higher, meaning the forward rates are countercyclical. Panel C also shows that the forecasts errors share cyclical variation across the data sources: forecast errors are more negative when the ex ante earnings yield is higher. That is, when expected returns appear objectively high, investors overestimate how high expected returns will be in the future. We similarly find that forward rates measured in option prices are strong predictors of future realized errors in the two surveys. This cross-measure predictability alleviates measurement-error concerns. More generally, the consistency across data sources suggests patterns in expectations that are shared across a broad range of investors.

1.2 Implications and Mechanisms

Implications. Our results on forward return expectations and forecast errors have implications for four literatures in asset pricing.

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 expected returns will be excessively elevated. Without such expectation errors, price fluctuations would be more modest; in our baseline estimation, the expectations errors can account for most of the stock price declines during crises, although it only accounts for 1/10 of the unconditional variation in the price-dividend ratio. Our results thus provide a new interpretation of stock-price variation, and crashes in particular, in recent decades.[3]

Stylized facts about the equity term structure. The literature on the equity term structure studies prices and expected returns for dividend claims with different maturities. The literature finds that the 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 (Golez and Jackwerth 2023, Gormsen 2021). Our results on forecast errors directly speak to these results, which are otherwise difficult to reconcile with standard models. In particular, we study a simple model where the equity term structure is flat, and show that the forecast errors uncovered in our paper result in a term structure of realized returns that is downward sloping on average and countercyclical.

Low price elasticity of demand. A recent literature documents that investors' demand for equities is relatively inelastic with respect to price changes (Gabaix and Koijen 2022). Given that a drop in prices tends to predict a significant increase in expected returns over short horizons, this inelasticity is puzzling. However, our results suggest that investors often mistakenly attribute a significant share of the decrease in prices to increases in expected returns at relatively long horizons, and perceive short-term spot returns to increase only modestly. This structure of expectation errors lowers the elasticity of demand, as a modest increase in spot returns will not lead to a large increase in desired portfolio weights in equities.

The debate on the cyclicality of subjective risk premia. A recent literature studies the cyclical dynamics of return expectations. Greenwood and Shleifer (2014) document procylical variation; Nagel and Xu (2022b) document acyclical variation; and Dahlquist and Ibert (2023) document countercyclical variation, all using somewhat different measures of expected returns and aggregate state variables. (Martin 2017 also finds countercyclicality in option-based

---

[3]We note that our results *account* for a portion of excess volatility without necessarily constituting an *explanation* of that volatility: while we explore theoretical mechanisms that can generate our results, the analysis is to a large degree silent on the underlying sources of forecast errors.

estimates.) We find that the degree of cyclicality depends on the horizon. In the surveys we study, short-run "spot" expectations of returns are on average acyclical but longer-run forward rates are countercyclical. (One interpretation of this finding is that respondents understand present value logic — they understand that future long-run returns must be high during crises when prices are low — but during crises, they believe it will take a while before prices start increasing.) As Nagel and Xu (2022b) generally study shorter-run expectations and Dahlquist and Ibert (2023) study capital market assumptions over longer horizons, our results help reconcile the apparent disagreement. Our focus on forward rates also eliminates any apparent disagreement between option-based measures and survey-based measures: 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 expected returns, which are the key object of interest in determining prices.

Mechanisms. We also consider potential drivers of our results. We first address a set of alternative explanations unrelated to forecast errors in subjective expectations of log returns. For the survey data, we discuss how measurement of simple as opposed to log-returns may influence our results. For the options data, we investigate conditions under which our results can be explained by the behavior of the stochastic discount factor. Our option-based estimates of spot and forward rates contain a risk premium term that may be large, but importantly for our analysis, the effect on forecast errors is likely modest. For forecast errors, we show that the risk premium term on the forward rate largely cancels with the risk premium term on the spot rate, with the remaining risk correction being an order of magnitude smaller than the well-known covariance term in Martin (2017). In order for this remaining risk premium term to rationalize the behavior of forecast errors, the stochastic discount factor must feature a highly volatile and countercyclical price of risk on shocks to the equity premium. The notion that the price of risk increases in bad times is consistent with past work (Campbell and Cochrane 1999), but the pace at which this risk price must change, and the range of values it must take, appear difficult to reconcile with standard models (and with our survey evidence).

1.3 Related Literature

We build on literatures studying the behavior of investor expectations measured both from derivatives prices and in survey data. For the derivatives-based measures, our paper builds most closely on Giglio and Kelly (2018), who document excess volatility in long-maturity relative to short-maturity claim prices across a range of term structures. They focus on the term structure of at-the-money implied variance when considering equity markets. We use a different measure of implied volatility, building on Gao and Martin (2021), and we

show how this measure can be connected to the term structure of expected equity returns and forecast errors.[4] Our option-based estimates are also related to Martin (2017), who studies option-implied spot expected returns under a covariance condition we build on, and Augenblick and Lazarus (2023), who derive volatility bounds for fixed-maturity index options. We differ from these papers in our focus on the term structure of expected returns.[5] We also, in contrast with much of this past work, consider option-based measures jointly with survey-based measures of expectations.

Adam and Nagel (2023) provide a recent review of the literature on survey-based expectations. This literature often finds that investors exhibit less than fully rational expectations in their survey responses, though as above, this behavior varies across settings (in addition to the references in Section 1.2, Boutros et al. 2020, De la O and Myers 2021, Couts, Gonçalves, and Loudis 2023, and Gormsen and Huber 2023 provide further recent evidence). We consider how investors expect future expected returns to evolve, whereas past work has largely focused on contemporaneous expected returns. We also provide novel evidence of the consistency of expectations dynamics across settings when considering these forward expectations. Further, to the degree that there remain differences in the behavior of survey expectations across different groups of investors, one advantage of our options-based analysis is that we directly study the expectations expressed in equilibrium behavior.

Finally, in comparing forward and realized spot rates, our analysis appears similar in spirit to tests of the expectations hypothesis (EH) for the fixed-income term structure (e.g., Fama and Bliss 1987, Campbell and Shiller 1991). Aside from focusing on equity, our methodology is geared towards estimating physical expectations (and particularly forecast errors) for future expected returns; by contrast, the forward and spot rates used in fixed-income EH tests embed risk premia by design, and violations of the EH in that setting are equivalent to bond-return predictability.[6]

---

[4]Our paper thus connects to other work on the term-structure consistency of implied volatility, including Stein (1989) and Mixon (2007). Our result that average forecast errors are close to zero echoes the finding of Dew-Becker, Giglio, Le, and Rodriguez (2017) that it is costless on average to hedge forward variance news.

[5]This brings us slightly closer to Augenblick et al. (2023), who study how updating differs by horizon. For other recent work on measuring expectations from options, see, for example, Polkovnichenko and Zhao (2013) and Chabi-Yo and Loudis (2020). Aït-Sahalia, Karaman, and Mancini (2020) also consider option-based term structures of expected returns using a parametric framework, in contrast to our approach.

[6]Bond-return predictability may, however, arise from the dynamics of physical expectations, as in Cieslak (2018) and Farmer, Nakamura, and Steinsson (2023); our analysis relates more closely to this work. Similarly, d'Arienzo (2020) uses both prices and surveys to study expectations for the fixed-income term structure.

2 Methodology

We begin with a theoretical analysis of the term structure of log equity 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 log equity premium between period 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 portfolio from t to $t+n$. For any horizon n, this expected excess return can be written as the one-period spot rates 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 parts 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

9

rates, forward rates, and forecast errors, respectively. More specifically, in place of (1),

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

Then $\tilde{f}$ and $\tilde{\varepsilon}$ are defined as in (2)–(3), with $\tilde{\mu}$ in place of μ. For 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. This assumption is typical in the literature on derivatives term structures; see Giglio and Kelly (2018) for a discussion.

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).

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

$$\mathbb{E}_t[r_{t,t+n}] = \mathbb{E}_t[r_{t,t+n}]E_t[M_{t,t+n}R_{t,t+n}]$$

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

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 directly observable from option prices, and they label this term the LVIX.[7] 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 this adjustment can be controlled by theory, as discussed in detail 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

---

[7]We use a slightly different definition of LVIX from theirs, as ours captures expected returns rather than excess returns. Separately, one can equivalently write the LVIX in (4) as $\mathscr{L}_t^{(n)} = (R_{t,t+n}^f)^{-1}\mathbb{E}_t^*[R_{t,t+n}r_{t,t+n}]$, where $\mathbb{E}_t^*$ denotes the risk-neutral expectation. Note that the relevant risk-free rate is the rate from t to $t+n$.

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:*

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

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 given this assumption, we can directly identify forecast errors from option prices and thereby 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). And this result does not in fact 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).

If we go beyond log utility, estimates of spot and forward rates are contaminated by the covariance terms in (4). When considering spot rates by themselves, it may be reasonable to assume that $\mathcal{C}_t^{(n)} = \operatorname{cov}_t(M_{t,t+n}R_{t,t+n}, r_{t,t+n}) \leq 0$, 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 one would obtain under the log-utility assumption. To streamline notation, we also define $MR_{t,t+n} = M_{t,t+n}R_{t,t+n}$. We can now show that we can study expected forecast errors up to a single covariance term.

11

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. Then the expected value of the forecast-error proxy satisfies*

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

where $\varepsilon_{t+n}^{(m)}$ is the true forecast error.

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 $\mathcal{C}_t^{(n)} = \text{cov}_t(M_{t,t+n}R_{t,t+n}, r_{t,t+n})$ 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 Section 4.1.5.

The following proposition shows that the above intuition carries through to more general non-log-normal settings, in which case the expected return $\mathbb{E}_{t+n}[r_{t+n,t+n+m}]$ is 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*

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

where $\varepsilon_{t+n}^{(m)}$ is the true forecast error.

The appearance of the LVIX $\mathscr{L}_{t+n}^{(m)}$ in the risk adjustment term in (6) has one benefit relative to (5): the LVIX is directly observable in the data. We can thus quantify the degree to which $\mathscr{L}_{t+n}^{(m)}$ is less volatile than the realized return $r_{t,t+n}$ in (4). In particular, under our main specification (discussed further in Section 3 below), we find that the unconditional volatility of $\mathscr{L}_{t+n}^{(m)}$ 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 forecasts errors using surveys that elicit return expectations at multiple horizons. While measurement of these objects in survey data is more straightforward than in the options data, the exercise is still not fully model-free. First, at a basic level, the interpretation of forecast errors as stemming from intertemporally inconsistent expectations requires that the set of survey respondents stays constant over time (or that the sample is representative period by period).

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.[8] 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 $\mathbb{E}_t[r_{t,t+n}]$ here, we are accordingly assuming that these higher-order cumulants are fixed over time when estimating forecast errors. If this is not the case, then our forecast errors can be thought of as being relevant for expectations of simple returns. These are of separate interest in their own right, but not exactly equivalent to the values measured using options.

---

[8]This reflects both a conceptual desire to understand variation in prices, and a desire for this analysis to remain consistent and comparable with the option-based analysis.

3 Data and Implementation

This section describes the data on options and surveys and discusses practical implementation issues for the option-based and survey-based measures of expectations.

3.1 Option-Based Measures

3.1.1 Options Data

We use a large dataset of option prices largely from OptionMetrics.[9] The dataset contains data on European (in terms of exercise) put and call options for major stock market indexes around the world. In our full sample, we have a total of 20 different indexes from at least 15 different countries and three pan-European indexes, as shown in Appendix Table A1. The sample starts in 1990 for the S&P 500 and runs through 2021. For international indexes, the respective samples start substantially later. In addition to focusing separately on the S&P 500 and the Euro Stoxx 50, our main panel analysis is conducted in a constrained sample of the 10 largest, most liquid, and most dense indexes in our sample, as detailed again in Table A1. Options for these exchanges are available since 2006 at latest, and again run through 2021. In all cases, we sample monthly and apply standard filters. Details on the data, sample selection, and filters are relegated to Appendix B.

3.1.2 Measuring LVIX

We measure $\mathscr{L}_t^{(n)}$ following Breeden and Litzenberger (1978) and Carr and Madan (2001) (see 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). We denote the relevant index at time t as P_t, the time t prices for call and put options on P_{t+n} with strike K as $\text{call}_t^{(n)}(K)$ and $\text{put}_t^{(n)}(K)$, and the time t forward price for the index at $t + n$ as $F_t^{(n)}$. 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 LVIX is a function of put and call prices written on the relevant index (as well as the current and forward price of the index). As in Section 2.2, we differ slightly from Gao and Martin (2021) in adding the risk-free rate in our definition of LVIX.

---

[9]We augment our sample with data from CBOE Market Data Express to cover the period prior to 1996 in the U.S. data. For risk-free rates used as inputs in calculating risk premia for this analysis, we use the OptionMetrics zero-coupon yield curve for consistency.

3.1.3 Implementation for Option-Based Measures

We run our baseline analysis at the 6-month horizon ($n = 6$); when considering forward rates, our baseline also uses forward maturities of $m = 6$ months. Our choice of horizon reflects a tradeoff: longer horizons are more economically meaningful, but long-maturity options are generally less liquid and have relatively sparse density. 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 magnitude of forecast errors over the entire term structure. We also provide a robustness analysis for our main empirical results at alternative horizons. The 6-month horizon also allows for useful comparison with the survey-based measures, as one of our two 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), and we therefore truncate the integral after extrapolating well past the range of observable contracts. We examine this assumption, along with a number of other assumptions with respect to measurement, in detail in the Appendix.

3.2 Survey-Based Measures

Our analysis uses data from two surveys, the Livingston Survey of professional forecasters and the Duke CFO survey.

3.2.1 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 2022. The survey asks respondents about 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.[10] We use median responses in all cases. We calculate the 6-month spot rate as the annualized log expected price change from month 0 to month 6 (without accounting for dividends), and we similarly calculate the 6-month, 6-month forward rate as the annualized log expected price change from month 6 to month 12.[11] This horizon accordingly aligns

---

[10] For example, for the December 2018 survey, respondents are asked to provide their forecasts for the S&P 500 value as of the end of December 2018, the end of June 2019, and the end of December 2019. S&P forecasts are available beginning in 1990, but the zero-month responses are available only beginning in 1992.

[11] We ignore dividends in order to align this analysis with the option-based estimates. This is 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.

with the main horizon considered for the options data. We compare this forward rate to the realized spot rate as of the following survey to calculate forecast errors.

We also re-conduct this 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, as forecasts of the 6-month rate are not elicited. For the spot rate, we subtract the annualized log T-bill rate as of month 0. For the forward rate, we subtract the annualized forecasted log T-bill rate for the 6-month forecast horizon. We then calculate forecast errors as above.

3.2.2 CFO Survey: Data and Implementation

We obtain data on CFO expectations from the Duke CFO Global Business Outlook survey. We combine data from the website maintained by Duke, which has data for the early years, with new data from a website maintained by the Richmond Fed.[12] The series is quarterly, starting in 2004 Q1 and running until 2023 Q1. There are a few missing quarters as the survey did not appear to ask about expected returns (specifically, Q1 of 2019 and Q1 and Q2 of 2020). We use mean responses as median responses are not available until 2009. The survey asks respondents about expected returns on the S&P 500 at both the 1-year and 10-year horizon. These forecasts allows us to estimate the forward rate for the 9-year expected return as of 1 year from time t, $f_t^{(1 \text{ year, } 9 \text{ year})}$. We annualize this forward rate and compare it to the annualized 10-year spot rate 1 year in the future. (We do not observe the realized 9-year spot rate, but it will likely be very close to the 10-year spot rate in annualized terms.)

We also re-conduct the analysis using equity risk premia. For this analysis, we measure risk premia relative to the 1-year and 10-year yield on U.S. Treasuries.

4 Empirical Results

Sections 4.1 and 4.2 study spot rates, forward rates, and forecast errors based on our options-based and survey-based measures of expectations, respectively. Across the data sources, we find that forward rates are excessively countercyclical relative to realized spot rates.

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

---

[12]The Duke website is https://cfosurvey.fuqua.duke.edu/release/, and the Richmond Fed site is https://www.richmondfed.org/cfosurvey/.

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 both the behavior of 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 the forward expected returns. However, the 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. The forecasting analyses for risk premia are accordingly better behaved. When presenting time-series plots of spot and forward rates, we focus largely on risk premia in order to focus on the relevant cyclical variation rather than the trend in interest rates.

4.1.1 Spot and Forward Rates: Descriptive Statistics and Figures

We begin with an overview of our estimated spot and forward rates. Table A2 in the Appendix provides summary statistics by exchange. For a brief intuitive description of the dynamics of spots and forwards, we focus here on descriptive figures for the U.S. (S&P 500) 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 again procyclical, with forward rates increasing less than one-for-one with contemporaneous spot rates in bad times. The "contemporaneous" qualifier for spot rates is important here, as will be seen below.

The bottom panel of Figure 2 instead compares forward rates with ex post realized spot rates. (The 6×6 forward rate is a predicted value using a shorter-horizon forward as an instrument, as discussed further below.) The gap between the realized spot rate and the forward rate represents the forecast error. Again 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. So while the term structure of expected risk premia is procyclical, forecast errors are countercyclical.

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

17

4.1.2 Forward Rates as Predictors of Future Spot Rates

We first consider Mincer-Zarnowitz (1969) regressions of realized spot rates on forward rates:

$$\mu_{i,t+6}^{(6)} = \beta_0 + \beta_1 \, f_{i,t}^{(6,6)} + e_{i,t+6}, \tag{8}$$

where i now indexes the exchange. We consider both panel regressions and US-specific time series regressions, and in both cases we conduct rolling 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 time-series case, we use heteroskedasticity- and autocorrelation-robust Newey-West standard errors, with lags selected following Lazarus et al. (2018), and fixed-b critical values (see Lazarus et al. 2021).

Table 2 presents the regression 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 consider spot and forward rates for expected returns (i.e., μ and f). The relevant null for forecast predictability features $\beta_0 = 0$, $\beta_1 = 1$. For risk premia, this 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 3.1.3, without any fixed effects. The slope coefficient is around 0.64, and the intercept is statistically significant at 1.05 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. Thus, even though forward rates move less than one-for-one with *contemporaneous* spot rates as in Figure 2, they move *more* than one-for-one with *future* spot rates. 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.

The next columns consider different specifications and samples. Column (2) includes an exchange fixed effect; this is in fact our preferred specification for estimating slope coefficients in the main panel, as it cleanly identifies within-country (or within-exchange) time-series predictability. The slope coefficient in this case is slightly further below 1 than in column (1), and the within-exchange R^2 is slightly below 0.2. Column (3) reports similar results for the S&P 500 in isolation.

In columns (4)–(6) we report results of similar set of regression for expected returns, as opposed to risk premia. These regressions suggest a higher degree of predictability in spot rates, but this increased predictability is driven by the persistent variation in the risk-free rate: in the early parts of the sample, risk-free rates are substantially higher than in the later

part of the sample. Since the spot and forward rates embed this variation in the risk-free rate, the slope coefficients and the R^2 go up substantially.[13] We note that we cannot reject $\beta_0 = 0$, $\beta_1 = 1$, but this failure to reject does not imply that the expectations do not overreact, as will be examined in subsequent regressions.

One might be concerned with attenuation bias for the slope coefficients in Table 2 given possible measurement error in forward rates.[14] We now address this concern by instrumenting variation in the forward rate with another forward rate. This IV specification will account for measurement error idiosyncratic to a given forward rate, though not for measurement error in the overall forward curve. We generally use shorter-maturity forward rates as instruments, as these are likely to be better measured given that options are denser at shorter maturities. In particular, for regression (8), we use the 2×1-month forward rate $f_{i,t}^{(2,1)}$ as an instrument for the 6×6 rate $f_{i,t}^{(6,6)}$.

Table 3 presents the resultant two-stage least squares results.[15] For the risk-premium regressions in column (1) to (3), 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 in column (4) to (6), the results hardly change; this is again consistent with the dominant impact of the secular trend in interest rates in the estimation of slope and intercept in these regressions.

Taken together, 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% relative to their within-exchange mean, this corresponds on average to an increase in future spot rates of only about 0.7%.

4.1.3   The Insignificance of Average Forecast Errors

We next turn our attention to forecast errors, which can be estimated under softer assumptions than the forward and spot rates themselves (see Propositions 2 and 3) We first consider unconditional averages of forecast errors, $\bar{\varepsilon}_{i,t+6}^{(6)}$, across different sets of exchanges and subsamples. These values are reported in Table 4. For risk premia (column (1) to (3)), we find average forecast errors that are statistically insignificant and effectively zero: $\bar{\varepsilon}_{i,t+6}^{(6)}$ is below 20 basis points on an annualized basis. Realized spot rates for risk premia have thus

---

[13]In unreported results, we find that splitting the sample in two substantially lowers the R^2 and slope values in each subsample, consistent with the reported results being driven by persistent risk-free rate variation.

[14]Recall that forwards are estimated from discrete approximations of the integrals in (7), for which extrapolation past the range of observable strikes is necessary. Again see Appendix B for details.

[15]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.

been very slightly (and insignificantly) higher than forward rates on average.[16] This rough magnitude applies to different exchanges, though the U.S. average in columns (2) and (3) is even lower, at roughly 2 to 3 basis points.

For expected returns (column (4) to (6)), we find average forecast errors that 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. Alternatively, it could be driven by a constant term premium on interest rates that causes the pure expectations hypothesis to fail. 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, which we turn to now.

4.1.4 The Predictability of Forecast Errors

Figure 3 plots realized 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: in years following crises, forecast errors tend to be significantly negative, which is to say future realized spot rates are lower than expected ex ante. To the extent that these forecast error are systematic — i.e., they occur systematically following crises — 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 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 $f_{i,t}^{(2,1)}$ here instead of the 6×6 rate, as we would otherwise have the same forward rate on both sides of the regression, leading to a mechanical upward bias in the slope coefficient in the presence of measurement error.[17] All other aspects of the regressions are identical to those in Table 2.

The first three columns of Table 5 reports the results of equation (9). In the main panel without any fixed effects (column (1)), the slope coefficient on the forward rate is -0.13 and

---

[16]In Appendix B, we consider whether the small positive forecast errors might arise due to measurement noise. As discussed there, our LVIX estimate could be biased downwards and more so at long maturities. Nonetheless, in all the robustness tests considered there, we find that measurement error is unlikely to be large enough to generate even the 20 basis points estimated for average errors.

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

is significant at the 5% level.[18] 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 saying that forecast errors are trending over the sample (i.e., positive in the early sample and negative in the later sample, or vice versa). By instead using the more transitory forward rate for risk premia ($\tilde{f}$) as the predictor variable, our predictor variable 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 (column (4) to (6) of Table 5) are in fact our preferred specifications, because we have transitory components on both the left- and right-hand side, and 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 5, the forecast errors for expected returns are highly predictable. The slope coefficients in column (4) to (6) roughly double relative to the regressions in columns (1) to (3). 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 suggests that the patterns observing in columns (1) through (3) are not driven by the risk-free rate component (i.e., a failure of the expectations hypothesis for interest rates). In fact, the risk-free rate slightly masks the true predictability in forecast errors: when using our preferred specification, predictability is substantially higher. This test for forecast-error predictability is evidently more powerful against a rational null than the Mincer-Zarnowitz regressions presented above.

In Appendices B–C, we consider a number of robustness checks with alternative measurement and estimation methodologies and sample cuts. Results are similar in all cases (see Tables A3–A10). These additional results include consistent evidence from Coibion-Gorodnichenko regressions (Table A4), as well as tests in which we expand the sample to all 20 available exchanges, for which we find somewhat stronger error predictability than in the main sample (Table A8).

We conclude that the pattern of excess cyclicality observed in Figure 1 and 3 is systematic

---

[18]The magnitudes presented in the table are not directly comparable to those in Tables 2–3 given the use of $f_{i,t}^{(2,1)}$ as the predictor. Instead, regression (9) can be viewed as an analogue to the reduced form of the IV specification for (8) presented in Table 3.

and statistically significant. Forward rates have exhibited excess cyclicality in this sample, in the sense that 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 physical expectations of future spot rates. Our main source of confirmation for this interpretation is in our analysis of survey data, considered below. Before turning to the survey data, though, we consider the behavior of the risk-premium term in more detail.

4.1.5 Risk-Based Explanations for the Option-Based Results

We now ask what would be required in order for option-based forecast errors to be rationalized by the behavior of the stochastic discount factor alone. Our main conclusion is that we find the required behavior of the SDF difficult to reconcile with standard models. For readers more interested in corroboration of the above patterns in our survey data (which cannot be rationalized through these risk-premium terms), one can skip this self-contained theoretical discussion and resume the remainder of the paper in Section 4.2 without issue.

Overview. When going beyond log utility, our estimates of spot and forward rates partly reflect risk premium adjustments. While most of these cancel out in forecast errors, a risk premium term remains: from Propositions 2–3, expected forecast errors are

$$
\mathbb{E}_t[\widehat{\varepsilon}_{t+n}^{(m)}] = \mathbb{E}_t[\varepsilon_{t+n}^{(m)}] - \varsigma_t, \quad \text{where } \varsigma_t = \begin{cases} \operatorname{cov}_t\big(MR_{t,t+n}, \mathbb{E}_{t+n}\, r_{t+n+m}\big) & (\textit{log-normal case}), \\ \operatorname{cov}_t\big(MR_{t,t+n}, \mathscr{L}_{t+n}^{(m)}\big) & (\textit{general case}). \end{cases}
$$

To simplify analysis, we maintain focus on the $n = m = 1$ case (where one period is 6 months) and suppress the dependence of ς_t on n and m.

As a starting point for the potential role of risk premia, Figure 4 plots the predicted values of the risk-premium forecast errors from the regressions in Table 5. These are empirical estimates of $\mathbb{E}_t[\widehat{\varepsilon}_{t+1}^{(1)}]$, and $-\varsigma_t$ must take on these values in order to rationalize our results through risk premia alone (so that true forecast errors satisfy $\mathbb{E}_t[\varepsilon_{t+1}^{(1)}] = 0$). The main challenge is that the covariance term $-\varsigma_t$ must flip sign over time: the covariance term must be significantly negative in bad times, which means it must be positive in good times to remain close to zero on average. We show below that this requires the price of discount rate risk to be highly volatile — close to zero during good times and high during bad times.[19]

---

[19]As we will show below, it is possible for standard models to generate a ς_t that is sufficiently volatile, but if ς_t does not change sign, the average of $\widehat{\varepsilon}$ will be far from zero, which is inconsistent with the average of $\widehat{\varepsilon}$ being close to zero in the data.

Examining the Risk Premium Term. To better understand the properties required of ς_t in order to match Figure 4, it is useful to formalize its relation to a discount-rate risk premium. As in Proposition 2, assume for now that the SDF and returns are jointly log-normal. Then by Stein's lemma and the fact that $\mathbb{E}_t[MR_{t+1}] = 1$, we have that

$$\text{cov}_t(MR_{t,t+1}, \mathbb{E}_{t+1} \, r_{t+2}) = \text{cov}_t(m_{t+1} + r_{t+1}, \mathbb{E}_{t+1} \, r_{t+2}). \tag{10}$$

For ς_t to flip sign, the covariance in (10) must thus flip sign. To see how this covariance is tied to discount-rate risk, define $P_{\mathbb{E},t}$ as the price of the claim that pays $-\mathbb{E}_{t+1} \, r_{t+2}$ next period, and define its ex ante Sharpe ratio (in log-return terms) as $SR_{\mathbb{E},t}$.[20] This Sharpe ratio captures the price of risk for exposure to increases in expected returns (or discount rates).

Rewriting (10), the sign of ς_t is given by

$$\text{Sign}\,(\varsigma_t) = \text{Sign}\left(SR_{\mathbb{E},t} + \rho_t(r, \mathbb{E}_{t+1} \, r)\sigma_t(r_{t+1}) \right),$$

where $\rho_t(r, \mathbb{E}_{t+1} \, r) = \text{Corr}_t(r_{t+1}, \mathbb{E}_{t+1} \, r_{t+2})$. This correlation is likely negative, as low realized returns are associated with higher expected returns going forward; for simplicity, start by assuming that $\rho_t(r, \mathbb{E}_{t+1} \, r) = -1 \; \forall \, t$.[21] In this case, the above expression shows that $SR_{\mathbb{E},t}$ must be higher than $\sigma_t(r)$ for ς_t to be positive, and lower than $\sigma_t(r)$ for ς_t to be negative. Since ς_t must be positive in bad times and negative in good times, $SR_{\mathbb{E},t}$ must vary more than $\sigma_t(r)$ between good and bad times.

As a benchmark, Figure 5 plots the time variation in $\sigma_t(r)$ at the 6-month horizon from the standpoint of an investor with log utility.[22] The figure shows that volatility varies from close to 15% in good times to more than 50% in bad times. As such, the price of discount-rate risk must vary significantly and countercyclically to generate the needed variation in ς_t.

Moreover, in the case that $\rho_t(r, \mathbb{E}_{t+1} \, r) = -1$, then the value $-\varsigma_t$ is a scaled version of $\text{cov}_t(MR_{t,t+1}, r_{t+1})$. Under the *modified negative correlation condition* (mNCC) used throughout Gao and Martin (2021), this covariance is always negative, so ς_t cannot flip sign. This illustrates that the degree of countercyclical variation in discount-rate risk prices required under $\rho_t(r, \mathbb{E}_{t+1} \, r) = -1$ is quite restrictive, and it is ruled out by a range of standard models considered by Gao and Martin (2021).

In practice, however, we expect $\rho_t(r, \mathbb{E}_{t+1} \, r)$ to be higher than -1, as not all changes in

---

[20]That is, $SR_{\mathbb{E},t} = (\mathbb{E}_t[r_{\mathbb{E},t+1}] - r^f_{t,t+1} + \sigma^2_{\mathbb{E},t}/2)/\sigma_{\mathbb{E},t}$, where $r_{\mathbb{E},t+1}$ is the claim's log return and $\sigma_{\mathbb{E},t}$ is its standard deviation.

[21]This is an extreme case in which all movements in realized returns reflect variation in discount rates (Campbell and Ammer 1993).

[22]This is likely to be a conservative benchmark for the degree of variation in conditional volatility if investors are more risk-averse than implied by log utility.

stock returns are driven by discount rates. A $\rho_t(r, \mathbb{E}_{t+1} r) > -1$ implies a less volatile price of discount-rate risk than just discussed. However, it also implies that the price of discount-rate risk must be even lower during good times than above; for example, in the opposing extreme in which $\rho_t(r, \mathbb{E}_{t+1} r) = 0$ (i.e., with i.i.d. returns), the Sharpe ratio on the discount-rate claim must be *negative* in good times in order for ς_t to flip sign.

The discussion above can also be generalized beyond the log-normal case. For example, assume that the SDF can be written as

$$M_{t,t+1} = \beta \frac{V_W(W_{t+1}, z_{t+1})}{V_W(W_t, z_t)},$$

where V_W is an unconstrained investor's marginal utility of wealth, z_t is a vector of state variables that includes $-\mathscr{L}_{t+1}^{(1)}$, and V_W is weakly decreasing in each entry of z_t. If this agent is fully invested in the market, has relative risk aversion $-WV_{WW}/V_W$ of at least 1, and R_{t+1} and the entries of z_{t+1} are *associated* random variables,[23] then Result 4 of Gao and Martin (2021) can be applied to obtain that ς_t again cannot change sign under the benchmark $\rho_t(r, \widehat{\mu}) = -1$.

Further Intuition Under Power Utility. The above emphasizes that the price of discount-rate risk must vary substantially over time to rationalize our results. To better understand this result, we next reconduct our analysis under various utility functions generating a constant price of risk, and show how such functions will not be able to explain the data. This analysis will help provide intuition on the role of risk aversion in our results.

We calculate spot and forward rates using CRRA utility with γ ranging from 0.5 to 3 (see Appendix C.2 for details). As shown in Appendix Table A5, none of the power utility functions can rationalize the results. The functions with $\gamma > 1$ explain the time variation in the forecast errors, but they cannot explain the average being close to zero. To understand the intuition, consider the results derived above for log-normally distributed variables. When $\gamma > 1$, the covariance term ς_t correctly increases in bad times. However, it is always positive ($\text{cov}_t(m_{t+1} + r_{t+1}, \mathbb{E}_{t+1} r_{t+2}) > 0$ at all times), leading to large average forecast errors. The opposite mechanics are at play for $\gamma < 1$, for which $\text{cov}_t(m_{t+1} + r_{t+1}, \mathbb{E}_{t+1} r_{t+2}) < 0$ for all t.

This illustrates the challenge in explaining both the average and time variation in the covariance term, and thus in explaining the results from options through this term alone.

---

[23]Formally, the elements of the vector $X_{t+1} = (R_{t+1}, z'_{t+1})'$ are associated random variables if $\text{cov}_t(f(X_{t+1}), g(X_{t+1})) \geq 0$ for all nondecreasing functions f and g for which this covariance exists.

4.2 Empirical Results from Survey-Based Measures

We now turn our attention to spot rates, forward rates, and forecast errors estimated using survey expectations. As a preview of the results, Figure 6 plots forward rates along with predicted forecasts errors for both Livingston survey and the CFO survey. As with the option-based measures, we find that (i) forward rates are countercyclical, in that they increase in bad times, and (ii) forecasts errors are predictably negative following crises such as the 2001 recession, the global financial crisis, and the Covid crisis. These results are again consistent with excess cyclicality in forward expectations: during bad times, investors think future spot rates returns will be more elevated than their own beliefs justify ex post. We detail this analysis below for each of the two surveys in turn.

4.2.1 Spot and Forward Rates from the Livingston Survey

As explained in Section 3.2.1, we use Livingston survey of professional forecasters to estimate the 6-month spot rate and the 6-month, 6-month forward rate. As with the options-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. The sample runs from 1990 to 2021.

Column (1) and (2) of Table 6 shows the results of Mincer-Zarnowitz regressions of future spot rates on past 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, though comparable to those for expected returns; on average, around 50% 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 Table 6 report results from the error predictability regressions, where we predict future realized forecast errors based 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 future spot rates that are lower than expected. The R^2 values are between 6.6% and 7.7%, which is close to those obtained from the option-based measures. Given the countercyclical nature of the forward rates (see Figure 6, 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.

4.2.2 Spot and Forward Rates from the Duke CFO Survey

We next turn our attention to the CFO survey. This survey allows us to extend the forecast horizon substantially, as the survey elicits return expectations at both the 1-year and 10-year horizon. The longer horizon is advantageous in that it long-horizon expected returns are more important for prices, allowing us to more easily speak to impacts on prices. As explained in Section 3.2.2, we use the survey to calculate 1-year, 9-year forward rates, which are expected returns over 9 years, starting one year from now. We then compare these to the realized 9-year spot rate, which is approximated from the 10-year spot rate.

Columns (1) and (2) of Table 7 show the results of Mincer-Zarnowitz regressions of future spot rates on forward rates. The forward rates predict future spot rates with slope coefficients around 0.5 on average, slightly below the results from options-based measures and the Livingston measure. The difference in slopes may arise from the fact that the horizons for spot and forward rates (n, m) are different from the options- and Livingston-based measures, which both considered 6-month, 6-month forwards. The R^2 is around 0.3 on average. Columns (3) and (4) further show that the average error is negative; that is, forward rates are systematically higher than future realized spot rates.

Column (5) and (6) of Table 7 report our main results on error predictability. The forecast errors are again predictable by the ex ante forward rates. When forward rates are high, future spot rates are lower than expected. For risk premia, this effect is strong, with an R^2 of 34%. For expected returns, the effect is statistically insignificant, and quantitatively closer to the effects estimated for options and the Livingston survey.

As a final remark on the CFO survey, there is, in this survey, a clear distinction between the cyclical behavior of spot and forward rates. As shown in Figure 7, spot rates appear *procyclical*: short-term expected returns decrease in bad times in a manner that appears extrapolative with respect to past returns (as highlighted by Greenwood and Shleifer 2014), in spite of these being times in which expected returns are likely to be objectively high. The figure, however, also shows that forward rates are strongly countercyclical (and so are 10-year spot rates): long-term expected returns increase in bad times. One interpretation of these findings is that the CFO respondents understand present value logic — they understand that future long-run returns must be high when prices are low — but during crises, they believe prices will continue to decrease for a while before they start increasing. This finding sheds light on the debate about whether spot rates are procyclical (Greenwood and Shleifer 2014), acyclical (Nagel and Xu 2022b) or countercyclical (Dahlquist and Ibert 2023). Interestingly, there is less disagreement about forward rates: across all option-based and survey-based measures, longer-run forward rates are robustly and excessively countercyclical. The apparent disagreement in previous literature is thus 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 expected returns and forward rates are often the key object of interest in determining prices.

4.3 Cyclical Variation in Forecast Errors

Finally, to comprehensively link the forecast errors to stock prices and cyclicality, we now relate the errors back to the valuation ratio of the stock market. For each measure we project the ex post forecast errors onto the ex ante earnings yield for the market portfolio. We use two measures of the earnings yield, one base on the cyclically adjusted price-to-earnings ratio (CAPE) and one based on the excess CAPE, both obtained via Robert Shiller's website.

Table 8 shows that forecast errors are negatively related to the ex ante earnings yield. As such, future spot rates are lower than expected when the earnings yield is relatively high. The slope coefficients are similar in magnitude across sources, varying between -0.3 and -0.9. The slopes are slightly smaller for the CFO data than for the other sources, but this difference may be driven by the longer horizon for the forward rates using this measure. The fact that the earnings yield predicts forward rates and forecast errors for all sources alleviates measurement error concerns and points at a commonality in the expectations across sources. Taken together, the results suggest that when expected returns are objectively high, investors think expected returns in the future will be high, and more so than their own subsequent beliefs justify. In the upcoming section, we explore additional implications of this excess cyclicality for asset pricing.

5 Implications for Asset Pricing

This section studies additional implications of the excess cyclicality in forward rates documented in the previous sections.

5.1 A New Source of Excess Volatility

If investors think that bad asset-pricing shocks lead to 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 effect using errors from the option-based measure, and then conduct a similar exercise using errors from CFO survey. Both sources point to a quantitatively similar impact of predictable forecast errors on prices.

Excess volatility from the option-based measure. We focus on the component of excess volatility coming from expectations about risk premia, as opposed to expected returns, and note that this leads to 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 as expected equity premia:

$$
\begin{aligned}
p_t - d_t &= k + \sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t[\Delta d_{t+j+1}] \ - \sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t[r_{t+j,t+j+1}] \\
&= k + \underbrace{\sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t[\Delta d_{t+j+1}]}_{CF_t} \ - \underbrace{\sum_{j=0}^{\infty} \rho^j \, \tilde{f}_t^{(j,1)}}_{\mathcal{F}_t} \ - \underbrace{\sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t[r_{t+j,t+j+1}^f]}_{RF_t},
\end{aligned}
\tag{11}
$$

where k is a constant and $0 < \rho < 1$. We further split the forward-rate term $\mathcal{F}_t$ into

$$
\underbrace{\sum_{j=0}^{\infty} \rho^j \, \tilde{f}_t^{(j,1)}}_{\mathcal{F}_t} = \underbrace{\sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t[\tilde{\mu}_{t+j}^{(1)}]}_{\mathcal{M}_t} \ - \underbrace{\sum_{j=0}^{\infty} \rho^j \, \mathbb{E}_t[\tilde{\varepsilon}_{t+j}^{(1)}]}_{\mathcal{E}_t}.
\tag{12}
$$

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. The length of one period in (11)–(12) is arbitrary, so we set it to 6 months to correspond to our baseline horizon.

In our main sample, we estimate forecast errors for horizons only up to one year. To obtain estimates for longer horizons, we re-estimate spot and forward rates 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.[24] For each horizon, we predict forecast errors as in (9) by regressing realized forecast errors on shorter-horizon forward rates. We then estimate a decay function $\phi^{(n,m)}$ for forecast errors, such that $\mathbb{E}_t[\tilde{\varepsilon}_{t+n}^{(m)}] = \phi^{(n,m)} \, \mathbb{E}_t[\tilde{\varepsilon}_{t+1}^{(1)}]$. Details are in Appendix C. As reported in Table A6, the decay parameters are close to or greater than 1 at all horizons. We therefore use the assumption that predictable forecast errors have a flat term structure, $\mathbb{E}_t[\tilde{\varepsilon}_{t+j+1}^{(1)}] = \mathbb{E}_t[\tilde{\varepsilon}_{t+j}^{(1)}]$ for all j in (12), to obtain a back-of-the-envelope quantification of $\mathcal{E}_t$ for the U.S. sample:

$$
\mathcal{E}_t = \frac{\rho}{1 - \rho} \, \mathbb{E}_t[\tilde{\varepsilon}_{t+1}^{(1)}],
\tag{13}
$$

where again one period is 6 months. We then compare variation in $\mathcal{E}_t$ 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 (2022a).[25]

---

[24]This is similar to the case for the dividend futures market (Binsbergen and Koijen 2017).

[25]We estimate ρ for the 6-month (half-yearly) horizon in (13) as $\rho = (1 + \exp(\overline{d - p}))^{-1/2}$, and we use

Figure 8 presents the results, with both discounted forecast errors $\mathcal{E}_t$ (red) and the repurchase-adjusted price-dividend ratio $p_t - d_t$ (blue) shown as log differences from their full-sample means. The two series do not always comove positively. During the Russian debt crisis in the late 1990s, for example, forward rates were quite elevated (leading to negative $\mathcal{E}_t$), but the stock market's dot-com boom largely continued apace (leading to high $p_t - d_t$). But during large stock-market declines (low $p_t - d_t$), predictable forecast errors generally play a significant role in the decomposition in (11)–(12). During the depths of the 2008 financial crisis, for example, when the price-dividend ratio reached more than 70 log points (about 50 percentage points) below its mean, discounted forecast errors accounted for more than half of this decline. A meaningful component of the stock-price collapse therefore consisted of elevated expected future equity premia, which, in a predictable manner, declined faster than suggested by the forward curve. However, the while the forecast errors can explain a meaningful amount of price variation during crises, they explain a more modest amount of the overall variation in prices. A regression of forecast errors on the log divide-price ratio 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.

Excess volatility from the CFO survey. One can use the errors in from the long-horizon forward rates in the Duke-CFO survey to quantify the impact of forecast errors on prices more directly. As shown in Figure 6, the predicted errors for the 9-year spot rates are around -1.5% during the global financial crises and the Covid crises. Since these are annualized numbers, the finding implies that predictable forecast errors can roughly account for $1.5\% \times \sum_{j=1}^{9} \rho^j = 12.21\%$ of the decrease in prices observed during these crises (for $\rho = 0.98$). This estimate is arguably a lower bound, as it is plausible that investors make similar mistakes at horizons beyond 9 years. Overall, the estimate from the survey-based measures is again consistent with the estimates from the option-based measure.

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.

---

predicted forecast errors $\mathbb{E}_t[\tilde{\varepsilon}_{t+1}^{(1)}]$ expressed on a non-annualized basis.

In general, expected-return forecast errors should have a larger impact on equity claims with longer maturity, because these claims have longer duration, making them more sensitive to changes in expected returns. To illustrate this effect, we calculate the impact of forecast errors on realized returns to dividend claims with different maturity under two assumptions: first, we assume that expected returns to all equity claims are the same (i.e., the equity term structure is flat); second, we assume, as in the previous section, a flat term structure of forecast errors (in this case, $\varepsilon_{t+1}^{(j)} = \varepsilon_{t+1}^{(j+1)}$ for all j). Under these assumptions, log returns on dividend claims with different maturity (ν) are given by

$$r_{t+1}^{(\nu)} = \mu_t^{(1)} + \Delta E_{t,t+1} d_{t+\nu} - (\nu - 1) \times \varepsilon_{t+1}, \tag{14}$$

where $\Delta E_{t,t+1}$ denotes the change in expectations between period t and $t+1$ (i.e., ex post forecast errors) and where we have suppressed the superscript from the forecast error because of the assumption that these errors are the same across all horizons (and claims). Equation (14) shows that mistakes about forward rates influence realized returns on equity claims, and more so for longer maturity claims.

Predictability in forecast errors thus influences the slope of the equity term structure. If 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, meaning the equity term structure will be upward sloping. Similarly, when forecast errors are predicted to be positive (investors have incorrectly low expectations about future spot rates), the equity term structure will be downward sloping.

To visualize the implications of our forecast errors for the equity term structure, we calculate the impact of forecast errors on realized returns to dividend strips with different maturities based on equation (14). We focus on forecast errors for risk premia and calculate both average returns and conditional returns in good and bad times. We follow Gormsen (2021) and define good and bad times based on the ex ante valuation ratio.[26]

Figure 9 plots these term structures. The figure shows that the term structure of realized returns implied by the forecast errors is slightly negative on average and countercyclical. The estimates happen to be numerically close to the estimates for realized returns on dividend strips in Gormsen (2021), but we interpret the results qualitatively rather than quantitatively, given that our estimates reay fairly heavily on modeling assumptions (and represent log returns as opposed to arithmetic returns). Qualitatively, the figure shows how the estimated forecast errors can, in principle, produce a term structure that is downward sloping on

---

[26]We define bad times as the 20% lowest of values for the CAPE ratio.

average ([Binsbergen, Brandt, and Koijen 2012](#)) and countercyclical ([Gormsen 2021](#), [Golez and Jackwerth 2023](#)).[27] As such, our forecast errors help account for the dynamics of the equity term structure.

5.3   Impact of Forecast Errors on Demand Elasticities

Empirically, there is a strong relation between changes in prices and expected returns over the next period. [Nagel and Xu (2022b)](#), for instance, shows that a 1 standard deviation increase in the dividend yield increases expected returns on the market by 6 percentage points over the subsequent year. This large impact of prices on future expected returns should lead investors to allocate substantially towards equities when prices move. But our results suggest that investors often mistakenly attribute a significant share of the decrease in prices to increases in expected returns at relatively long horizons, and perceive short-term spot returns to increase only modestly. This structure of expectation errors lowers the elasticity of demand, as a modest increase in spot returns will not lead to a large increase in desired portfolio weights in equities. It can accordingly help make sense of the apparently puzzling inelasticity in investor demand for equities ([Gabaix and Koijen 2022](#)). We leave potential quantification of this effect to future work.

6   A Model of Expectation Errors

Finally, to provide a more positive potential explanation for our results, this section presents a simple model of expectations errors in the term structure of option-based spot rates that results in excess cyclicality in forward rates. The models builds on the framework proposed by [Bordalo, Gennaioli, and Shleifer (2018)](#), but it is also qualitatively similar to the natural expectations model proposed by [Fuster, Laibson, and Mendel (2010)](#) and adapted by [Giglio and Kelly (2018)](#). We then calibrate the model and compare it to the data under the assumption that the representative investor has log utility over the market return. The model is reduced-form in the sense that we directly model investors' expectations about an endogenous variable, rather than taking a stand on the underlying foundations for these expectations.

---

[27]If we considered forecast errors for for expected returns instead of risk premia, the term structure would be slightly upward sloping on average but still countercyclical.

6.1 Model Setup

We assume that the 3-month spot rate follows an AR(3) process under the objective measure:

$$\mu_t^{(3)} = \left(1 - \sum_{j=1}^{3} \phi_j\right) \bar{\mu} + \phi_1 \, \mu_{t-1}^{(3)} + \phi_2 \, \mu_{t-2}^{(3)} + \phi_3 \, \mu_{t-3}^{(3)} + e_t,$$

where $\bar{\mu} = \mathbb{E}[\mu_t^{(3)}]$ and $e_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_e^2)$.[28]

If expectations were rational, investors would use the objective AR(3) dynamics to make iterated forecasts of future 3-month spot rates:

$$\mathbb{E}_t\left[\mu_{t+n}^{(3)}\right] = \left(1 - \sum_{j=1}^{3} \phi_j\right) \bar{\mu} + \sum_{j=1}^{3} \phi_j \, \mathbb{E}_t\left[\mu_{t+n-j}^{(3)}\right] \text{ for } n > 0. \tag{15}$$

In such a rational world, these objective forecasts of future 3-month spot rates would define the current term structure of quarterly option-implied spot rates recursively using (2):

$$\mu_t^{(n)} = \begin{cases} \mu_t^{(3)} & \text{for } n = 3 \\ \mu_t^{(n-3)} + \mathbb{E}_t\left[\mu_{t+n-3}^{(3)}\right] & \text{for } n \in \{6, 9, 12\}. \end{cases} \tag{16}$$

Rather than assuming rationality as in (15), however, we assume that investors potentially overreact to objective news about $\mu_t^{(n)}$ in forming their subjective forecasts. We use $\mathbb{E}_t^{\theta}[\cdot]$ to refer to these investors' subjective expectations, as the parameter θ indexes the degree of excess sensitivity to spot-rate news as follows:

$$\mathbb{E}_t^{\theta}\left[\mu_{t+n}^{(3)}\right] = \mathbb{E}_t\left[\mu_{t+n}^{(3)}\right] + \theta \underbrace{\left(\mathbb{E}_t\left[\mu_{t+n}^{(3)}\right] - \mathbb{E}_{t-3}\left[\mu_{t+n}^{(3)}\right]\right)}_{\text{spot-rate news}}. \tag{17}$$

In the language of Bordalo et al. (2019), θ is the "diagnosticity" parameter. This representation nests rationality when $\theta = 0$. When $\theta > 0$, investors are excessively sensitive to news.[29] After bad news, for example, spot rates increase, the news term is positive, and investors' subjective expectations overreact by θ times news.

---

[28]By focusing on the ex ante equity premium as the stochastic process of interest, our model departs somewhat from the diagnostic model considered by, e.g., Bordalo et al. (2019), as discussed in Section 6.3. We also generalize slightly by assuming an AR(3) process.

[29]While overreaction is defined relative to objective news (which investors do not directly observe), the news term in (17) is proportional to the objective innovation e_t (which they do observe). Since news is defined over a 3-month period in the objective dynamics, we assume that the relevant lagged expectations in (17) are 3-month lagged expectations. In principle, this could also be estimated from the data as in Bordalo et al. (2019).

The actual term structure of quarterly option-implied spot rates $\mu_t^{(n)}$ is defined by (16), but with the objective expectation $\mathbb{E}_t[\mu_{t+n}^{(3)}]$ replaced by the subjective expectation $\mathbb{E}_t^\theta[\mu_{t+n}^{(3)}]$. As a consequence, long-horizon spot rates — and therefore forward rates — increase after bad news even more than they would under rationality.

6.2 Calibration and Results

To understand whether the model can match our empirical findings, we turn to Monte Carlo simulations. We first estimate the objective dynamics of the 3-month spot rate by country under the assumption of log utility. Table A11 reports these parameters. We then simulate 10,000 artificial samples of length equivalent to that in the data. In each artificial dataset, we run the same regressions as conducted in Section 4 to evaluate how the regression slopes and average errors vary with deviations from rational expectations.

While we consider a range of values for the key parameter θ, one particular focus is the value estimated by Bordalo, Gennaioli, and Shleifer (2018) and Bordalo et al. (2019), $\theta = 0.91$. They estimate this value in different settings than the one considered here, allowing us to test whether our empirical results are consistent with a fully externally calibrated θ. Given that this value of θ is close to 1, the magnitude of forecast errors is roughly comparable to news.

Figure 10 plots regression slopes and average errors by the degree of overreaction. The model is qualitatively consistent with the data. As θ increases, the Mincer-Zarnowitz slope β_1 decreases toward zero and the error-predictability and Coibion-Gorodnichenko slopes become more negative (see the Appendix for the empirical versions of these Coibion-Gorodnichenko regressions). Given (17), subjective expectations are on average unbiased, so the model cannot produce any non-zero average error. There is thus no relationship between the average error and θ. At the calibrated value $\theta = 0.91$, the model and data are roughly in line, as the empirical estimates are close to the model-implied confidence intervals in all cases; slightly lower values of θ would align the two even more closely. This simple model of expectation errors thus appears capable of explaining the excess sensitivity of forward rates we document empirically.

However, while the model is able to match overreaction reasonably well despite the stylized setup and few parameters, it does seem to fall short of matching the data on certain additional dimensions. In particular, it appears to underestimate the degree of rational variation in forward rates observed in the data: as seen in Figure A8, the model's regression R^2 values are somewhat below those observed in the data on average. For example, the Mincer-Zarnowitz R^2 is 20% in the data, but only 10% at the calibrated $\theta = 0.91$.[30] Nonetheless, the model

---

[30]Lower values of θ again improve this aspect of the model's match to the data, though this is not the case for the forecast-error predictability R^2 values.

still replicates the same qualitative patterns in R^2 values as one would expect.

6.3 Interpretation and Discussion

While the above model of expectation errors is capable of matching our empirical findings, the cyclicality of expectation errors distinguishes our setting from certain frameworks featuring return extrapolation, as discussed in Section 4.1.4. In the model, forward rates overreact to news about spot rates: when the objective equity premium is high, forward rates increase, and forecast errors are predictably negative. This echoes our empirical findings. Those findings would not be matched by specifying a model of overreaction of forward rates to realized returns: the estimated equity premium is high in bad times, when realized returns are *low*.

If the cyclicality of our model's expectation errors were to be flipped — if forward rates were too low in bad times, generating positive errors — then these errors would not be capable of explaining any of the variation in the price-dividend ratio shown in Figure 8. Forecast errors instead would exert a *dampening* force on stock prices, making it more difficult to explain their observed volatility. We discuss and formalize this idea in more detail in Appendix D, which introduces a "trilemma" for expectation errors: it is difficult to simultaneously make sense of (i) volatile expectation errors, (ii) countercyclical expectation errors (e.g., extrapolation of past returns), and (iii) volatile prices.[31] Our framework discards (ii), leaving (i) and (iii) on the table.

7   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 empirical analysis suggests that the term structure of investor expectations is well-behaved along multiple dimensions: 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, in fact, excessively countercyclical: in bad times, investors' expectations about future expected returns are more elevated than their own subsequent beliefs justify.

The finding of excessively countercyclical forward rates has broad implications for asset pricing. First, mistakes about long-run expected returns have a meaningful impact on stock

---

[31]This discussion builds on Campbell (2017), who introduces related trilemmas for present value and portfolio choice.

prices, and the excess cyclicality in forward rates therefore leads to substantial excess volatility. Across data sources, we find that the excess cyclicality can account for half the drop in expected returns during the global financial crisis and most of the drop during the Covid crises. The crises in recent time thus appear to have been much more short-lived than investors expected ex ante.

The excess cyclicality also has implications for other areas in asset pricing. For instance, the forecast errors induced by excess cyclicality can account for the dynamics of the equity term structure studied by Binsbergen, Brandt, and Koijen (2012): the structure of the forecast errors is such that the equity term structure is slightly downward sloping on average and countercyclical. The forecast errors also tend to reduce the price-elasticity of demand.

An exciting feature of these empirical results is that they are robust across option-based and survey-based measures. Whereas the dynamics of short-run expected returns appears to vary across data sources, we find no disagreement between option-based measures and survey-based measures when considering forward rates: in all cases, longer-run forward rates are robustly and excessively countercyclical, with strikingly similar magnitudes across measures. We therefore provide a unified body of evidence on the dynamics of forward return expectations.

Tables and Figures

Table 2
Mincer-Zarnowitz Regressions from Option-Based Expectations

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

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

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

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

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 lags selected following Lazarus et al. (2018) and fixed-b p-values. This sample is from January 1990 to June 2021.

	Risk Premia $\tilde{\mu}_{i,t+6}^{(6)} = \beta_0 + \beta_1\,\tilde{f}_{i,t}^{(6,6)} + e_{i,t+6}$			Expected Returns $\mu_{i,t+6}^{(6)} = \beta_0 + \beta_1\,f_{i,t}^{(6,6)} + e_{i,t+6}$		
	Main	Main	U.S. Only	Main	Main	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)	(0.072)
Intercept	1.04***		0.74***	0.095		0.27
	(0.17)		(0.28)	(0.22)		(0.34)
p-value: $\beta_1 = 1$	<0.001	<0.001	0.003	0.157	0.181	0.111
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
Cluster	Ex/Date	Ex/Date	-	Ex/Date	Ex/Date	-
Adjusted R^2	0.20	0.22	0.22	0.63	0.63	0.71
Within R^2	-	0.15	-	-	0.59	-

Table 3
Instrumented Mincer-Zarnowitz Regressions from Option-Based Expectations

This table reports instrumented 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 is the future expectation of the 6-month equity premium and the forward rate is the current expectation of the same risk premium. The instrument is the current expectation of the 1-month equity premium in 2 months:

$$\tilde{f}_{i,t}^{(2,1)} = \mathbb{E}_t\left[r_{i,t+2}^{(1)}\right]$$

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 lags selected following Lazarus et al. (2018) and fixed-b p-values. This sample is from January 1990 to June 2021.

	Risk Premia $\tilde{\mu}_{i,t+6}^{(6)} = \beta_0 + \beta_1\,\tilde{f}_{i,t}^{(6,6)} + e_{i,t+6}$			Expected Returns $\mu_{i,t+6}^{(6)} = \beta_0 + \beta_1\,f_{i,t}^{(6,6)} + e_{i,t+6}$		
	Main	Main	U.S. Only	Main	Main	U.S. Only
$\tilde{f}_t^{(6,6)}$	0.77***	0.70***	0.73***			
	(0.071)	(0.074)	(0.062)			
$f_t^{(6,6)}$				0.93***	0.93***	0.89***
				(0.060)	(0.072)	(0.029)
Intercept	0.73***		0.59***	0.021		0.21
	(0.19)		(0.13)	(0.24)		(0.15)
p-value: $\beta_1 = 1$	0.001	0.003	0.018	0.278	0.331	0.225
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
Cluster	Ex/Date	Ex/Date	-	Ex/Date	Ex/Date	-
Adjusted R^2	0.19	-	0.22	0.63	-	0.71
Within R^2	-	0.14	-	-	0.59	-

Table 4
Average Forecast Errors from Option-Based Expectations

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 lags selected following Lazarus et al. (2018) and fixed-b p-values. This sample is from January 1990 to June 2021.

	Risk Premia $\tilde{\varepsilon}_{i,t+6}^{(6)} = \tilde{\mu}_{i,t+6}^{(6)} - \tilde{f}_{i,t}^{(6,6)}$			Expected Returns $\varepsilon_{i,t+6}^{(6)} = \mu_{i,t+6}^{(6)} - f_{i,t}^{(6,6)}$		
	Main	U.S. Only	U.S. Post-2006	Main	U.S. Only	U.S. Post-2006
Intercept	0.17	0.021	0.029	-0.28**	-0.40**	-0.24
	(0.11)	(0.15)	(0.26)	(0.10)	(0.17)	(0.19)
Observations	2227	378	174	2227	378	174
Standard Errors	Cluster	Newey-West	Newey-West	Cluster	Newey-West	Newey-West
Cluster	Ex/Date	-	-	Ex/Date	-	-

Table 5
Predictability of Forecast Errors from Option-Based Expectations

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 lags selected following Lazarus et al. (2018) and fixed-b p-values. This sample is from January 1990 to June 2021.

	Risk Premia $\tilde{\varepsilon}_{i,t+6}^{(6)} = \beta_0 + \beta_1\, \tilde{f}_{i,t}^{(2,1)} + e_{i,t+6}$			Expected Returns $\varepsilon_{i,t+6}^{(6)} = \beta_0 + \beta_1\, \tilde{f}_{i,t}^{(2,1)} + e_{i,t+6}$		
	Main	Main	U.S. Only	Main	Main	U.S. Only
$\tilde{f}_t^{(2,1)}$	-0.13**	-0.16***	-0.17**	-0.25***	-0.29***	-0.20**
	(0.047)	(0.047)	(0.066)	(0.050)	(0.047)	(0.090)
Intercept	0.51***		0.39*	0.39**		0.039
	(0.15)		(0.23)	(0.17)		(0.29)
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
Cluster	Ex/Date	Ex/Date	-	Ex/Date	Ex/Date	-
Adjusted R^2	0.02	0.03	0.03	0.08	0.11	0.04
Within R^2	-	0.03	-	-	0.10	-

Table 6
Livingston Survey Regressions

This table reports regressions for Livingston survey expectations. See Section 3 for more details. Columns (1)–(2) are Mincer-Zarnowitz regressions of future realized spot rates on forward rates, as in Table 2. Columns (3)–(4) are average forecast errors, as in Table 4. Columns (5)–(6) are error-predictability regressions, as in Table 5. The units are annualized percentage points. Each regression reports Newey-West standard errors with lags selected following Lazarus et al. (2018) and fixed-b p-values. The sample is half-yearly from June 1992 to June 2021.

	Mincer-Zarnowitz		Average Error		Error Predictability	
	Risk Premia $\tilde{\mu}_{t+6}^{(6)}$	Expected Returns $\mu_{t+6}^{(6)}$	Risk Premia $\tilde{\varepsilon}_{t+6}^{(6)}$	Expected Returns $\varepsilon_{t+6}^{(6)}$	Risk Premia $\tilde{\varepsilon}_{t+6}^{(6)}$	Expected Returns $\varepsilon_{t+6}^{(6)}$
$\tilde{f}_t^{(6,6)}$	0.81***				-0.19**	-0.19***
	(0.077)				(0.077)	(0.073)
$f_t^{(6,6)}$		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-value: $\beta_1 = 1$	0.052	0.002	-	-	-	-
Observations	59	59	59	59	59	59
Adjusted R^2	0.56	0.38	-	-	0.05	0.06

Table 7
CFO Survey Regressions

This table reports regressions for CFO survey expectations. See Section 3 for more details. Columns (1)–(2) are Mincer-Zarnowitz regressions of future realized spot rates on forward rates, as in Table 2. Columns (3)–(4) are average forecast errors, as in Table 4. Columns (5)–(6) are error-predictability regressions, as in Table 5. The units are annualized percentage points. Each regression reports Newey-West standard errors with lags selected following Lazarus et al. (2018) and fixed-b p-values. The sample is quarterly from March 2004 to March 2022.

|	Mincer-Zarnowitz		Average Error		Error Predictability	
	Risk Premia $\tilde{\mu}_{t+1\mathrm{y}}^{(9\mathrm{y})}$	Expected Returns $\mu_{t+1\mathrm{y}}^{(9\mathrm{y})}$	Risk Premia $\tilde{\varepsilon}_{t+1\mathrm{y}}^{(9\mathrm{y})}$	Expected Returns $\varepsilon_{t+1\mathrm{y}}^{(9\mathrm{y})}$	Risk Premia $\tilde{\varepsilon}_{t+1\mathrm{y}}^{(9\mathrm{y})}$	Expected Returns $\varepsilon_{t+1\mathrm{y}}^{(9\mathrm{y})}$
$\tilde{f}_t^{(1\mathrm{y},9\mathrm{y})}$	0.42***				-0.58***	-0.15
	(0.15)				(0.15)	(0.11)
$f_t^{(1\mathrm{y},9\mathrm{y})}$		0.63***				
		(0.11)				
Intercept	2.26***	2.39***	-0.021	-0.21*	2.26***	0.37
	(0.58)	(0.71)	(0.14)	(0.12)	(0.58)	(0.39)
p-value: $\beta_1 = 1$	0.004	0.012	-	-	-	-
Observations	68	68	68	68	68	68
Adjusted R^2	0.20	0.46	-	-	0.34	0.03

Table 8
Cyclical Variation in Forecast Errors

This table reports cyclicality regressions for option-based (left panel) and survey-based (middle and right panels) expectations. Each panel reports predictability regressions of future realized forecast errors on the cyclically adjusted price-earnings yield and the excess yield (ExCAPE). The units are annualized percentage points. Each regression reports Newey-West standard errors with lags selected following Lazarus et al. (2018) and fixed-b p-values. The sample is monthly from January 1990 to June 2021 for option-based expectations, half-yearly from June 1992 to June 2021 for Livingston survey expectations, and quarterly from March 2004 to March 2022 for CFO survey expectations.

| | Expected Returns | | | | | |
| | Option-Based | | Livingston Survey | | CFO Survey | |
	$\varepsilon_{t+6}^{(6)}$	$\varepsilon_{t+6}^{(6)}$	$\varepsilon_{t+6}^{(6)}$	$\varepsilon_{t+6}^{(6)}$	$\varepsilon_{t+1y}^{(9y)}$	$\varepsilon_{t+1y}^{(9y)}$
$1/\text{CAPE}_t$	-0.51***		-0.74**		-0.27***	
	(0.18)		(0.32)		(0.097)	
$1/\text{ExCAPE}_t$		-0.58***		-0.87**		-0.30***
		(0.22)		(0.36)		(0.10)
Intercept	1.68**	1.71**	3.01**	3.17**	0.87**	0.88**
	(0.81)	(0.85)	(1.27)	(1.29)	(0.35)	(0.34)
Observations	378	378	59	59	68	68
Adjusted R^2	0.14	0.12	0.02	0.02	0.13	0.13

Figure 2
Option-Based Spot and Forward Rates

This 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: see Table 3 for more details. Spot and forward rates are for equity risk premia and are from option-based expectations. Gray bands are NBER recessions. The sample is from January 1990 to June 2021 in the U.S.
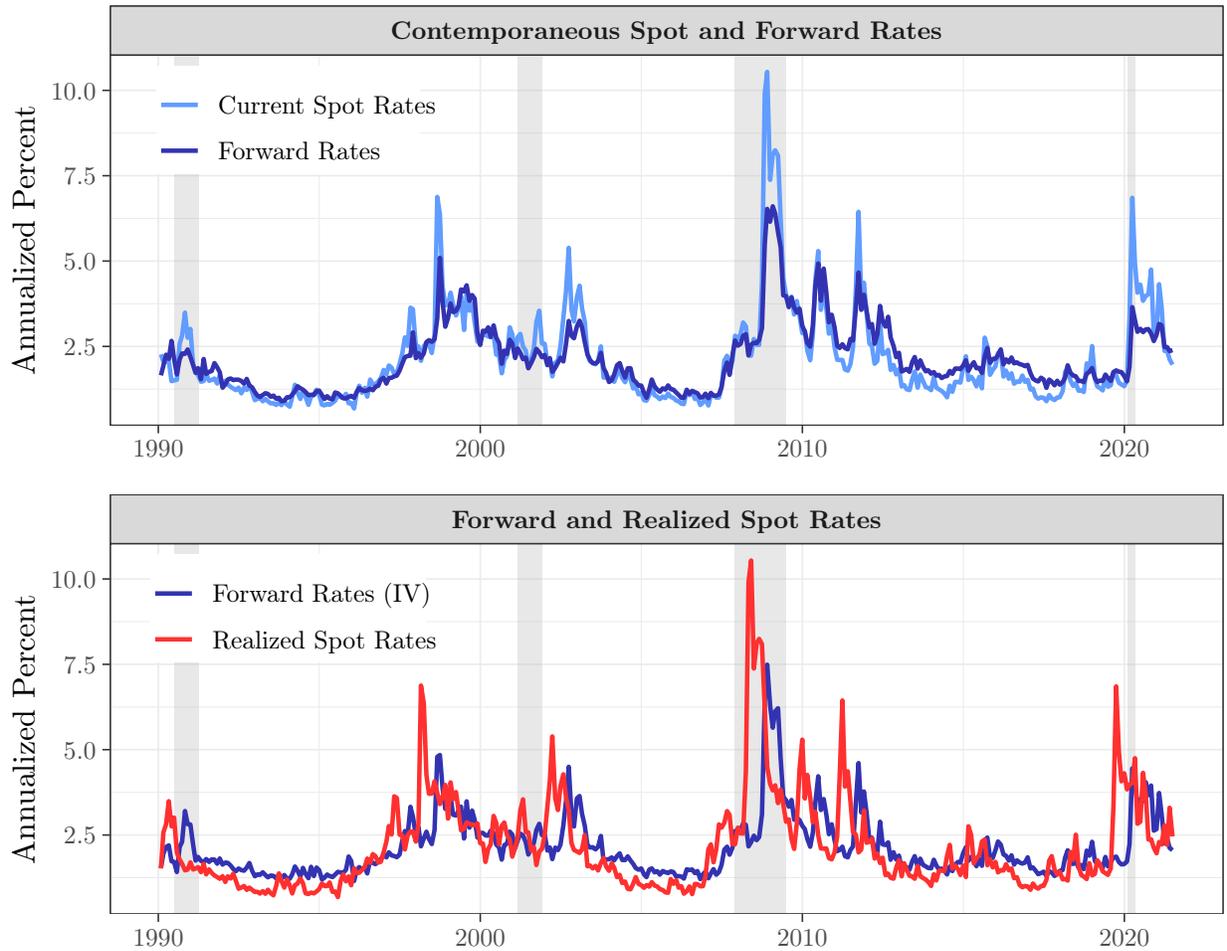
Figure 3
Option-Based Forecast Errors and Lagged Forward Rates

This figure plots forecast errors $\tilde{\varepsilon}_{t+6}^{(6)}$ (orange bars) and lagged 6×6-month forward rates $\tilde{f}_t^{(6,6)}$ (dark blue line). Forecast errors and forward rates are for equity risk premia and are from option-based expectations. The sample is from January 1990 to June 2021 in the U.S.
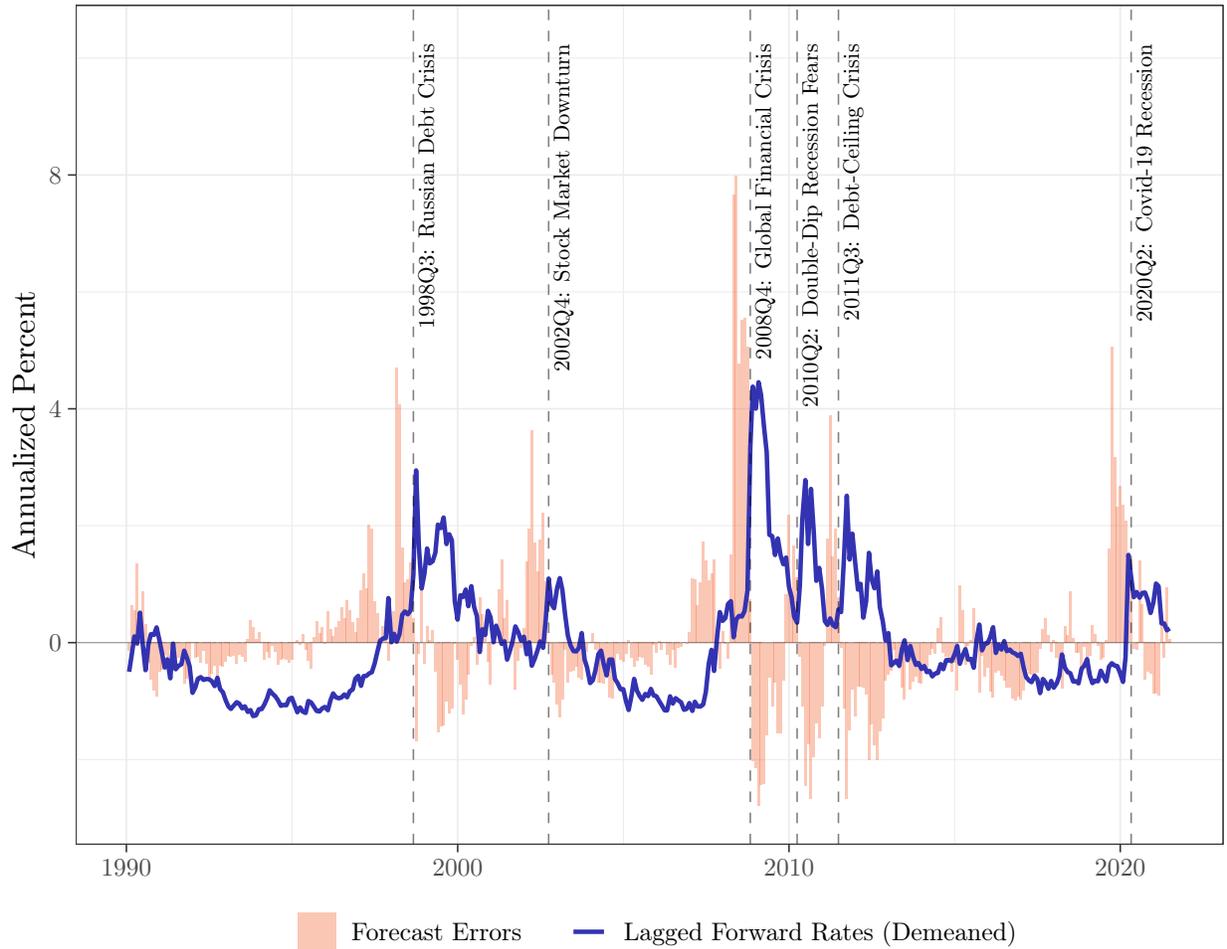
Figure 4
Option-Based Forward Rates and Predicted Forecast Errors

This figure plots instrumented 6×6-month forward rates $\tilde{f}_t^{(6,6)}$ (top panel) and the corresponding predicted forecast errors $\varepsilon_{t+6}^{(6)}$ (bottom panel). The instrument is the 2×1-month forward rate: see Table 3 for more details. The predicted forecast error is from a time-series regression of future realized forecast errors on current forward rates: see Table 5 for more details. Forward rates and forecast errors and are for equity risk premia and are from option-based expectations. Gray bands are NBER recessions. The sample is from January 1990 to June 2021 in the U.S.
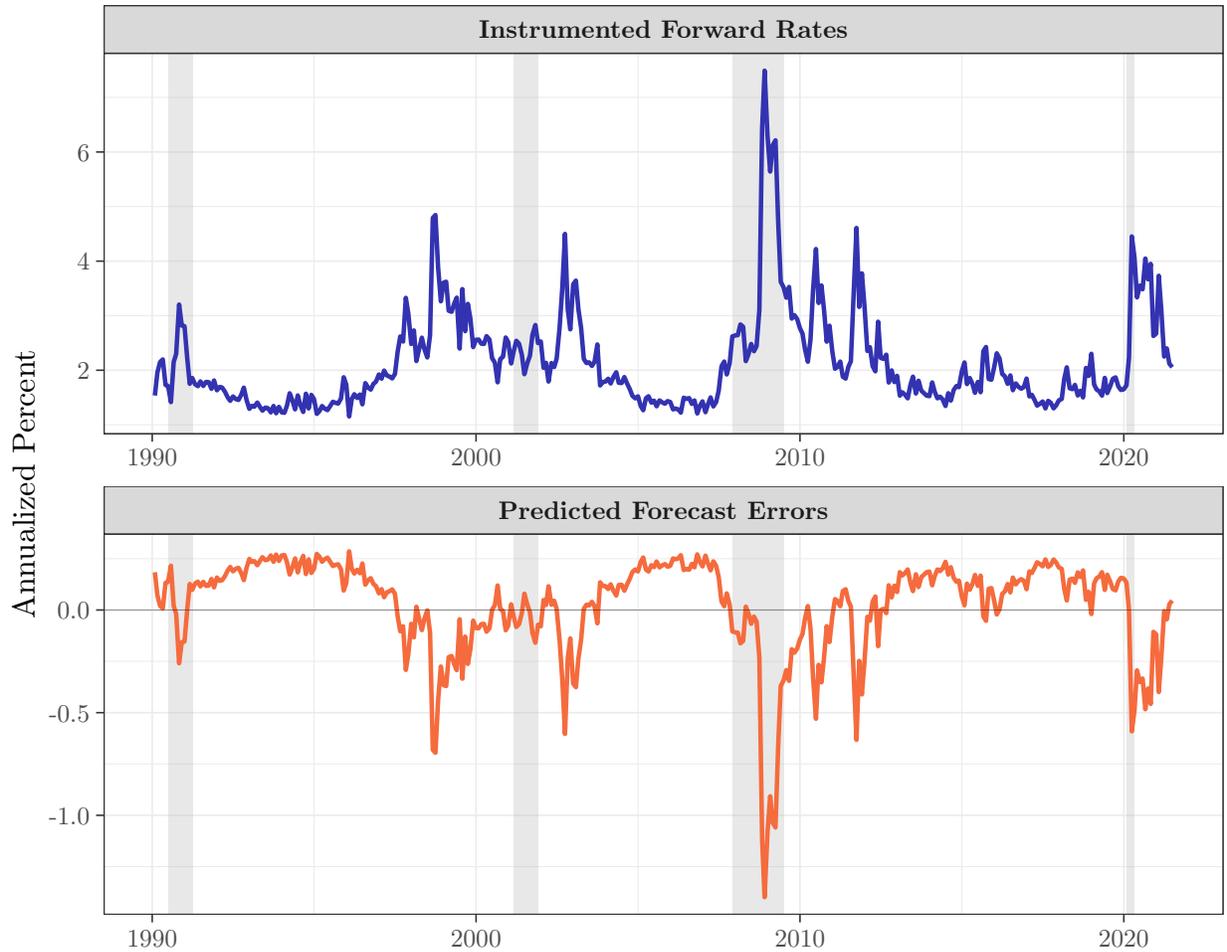
Figure 5
Option-Based Conditional Volatility of the Market Return

This figure plots the conditional volatility of the 6-month market return $\sigma_t\left(\ln R_{t,t+6}\right)$ from the standpoint of an unconstrained log utility investor fully invested in the market. See Appendix C.2 for more details. Gray bands are NBER recessions. The sample is from January 1990 to June 2021 in the U.S.

Figure 6
Survey-Based Forward Rates and Predicted Forecast Errors

This figure plots forward rates (top) and the corresponding predicted forecast errors (bottom) for Livingston (left) and CFO (right) survey expectations. See Section 3 for more details. The predicted forecast error is from a time-series regression of future realized forecast errors on current forward rates, as in Figure 4. Forward rates and forecast errors and are for equity risk premia. Gray bands are NBER recessions. The sample is half-yearly from June 1992 to June 2021 for the Livingston survey and quarterly from March 2004 to March 2022 for the CFO survey.


PANEL A. LIVINGSTON SURVEY — PANEL B. CFO SURVEY
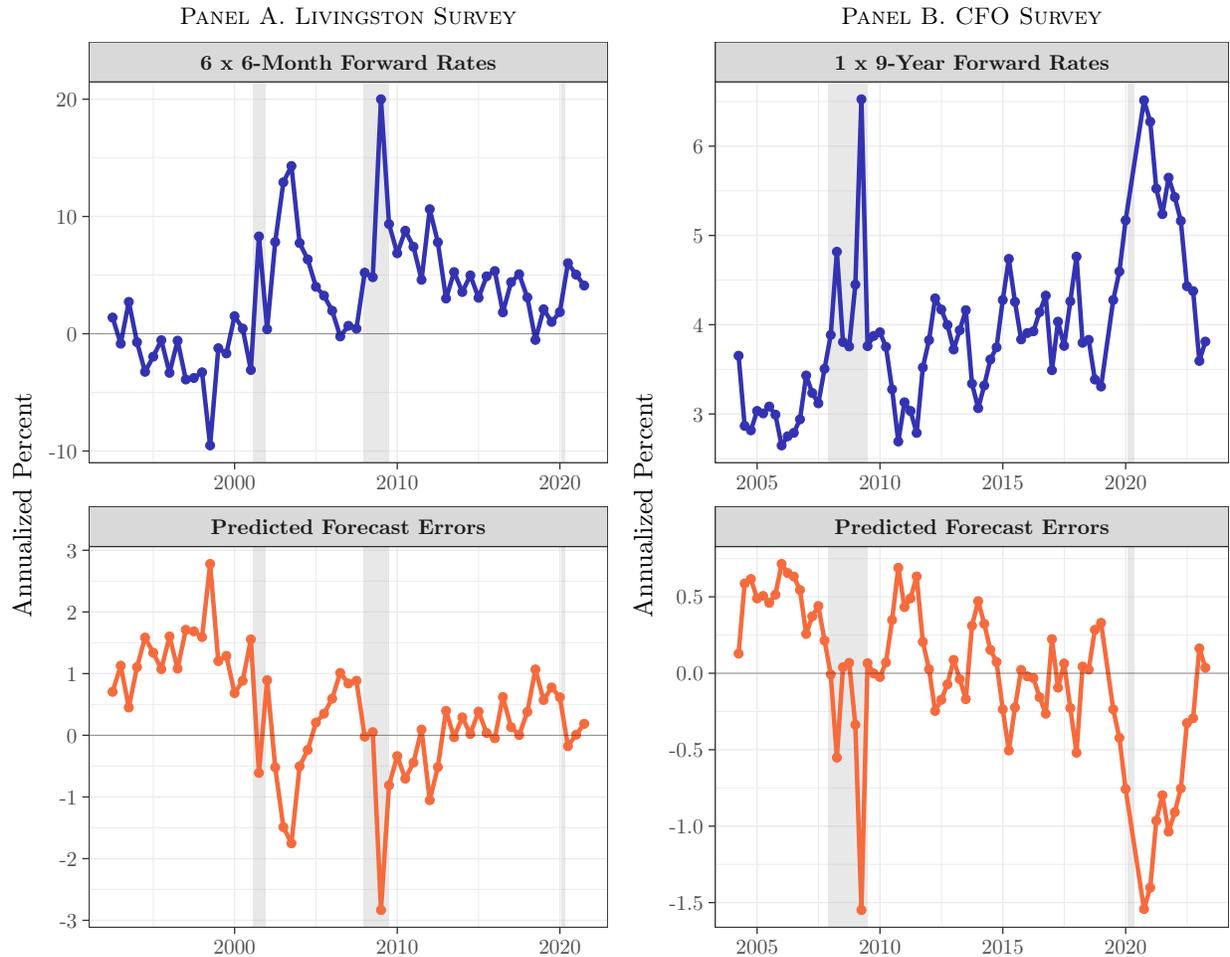
Figure 7
Spot and Forward Rates in the CFO Survey

This figure plots the spot rate and forward rate based on data from the CFO Survey. The spot and forward rates are for expected returns; see text for details on definitions and estimation. Gray bands are NBER recessions. The sample is quarterly from March 2004 to March 2022.
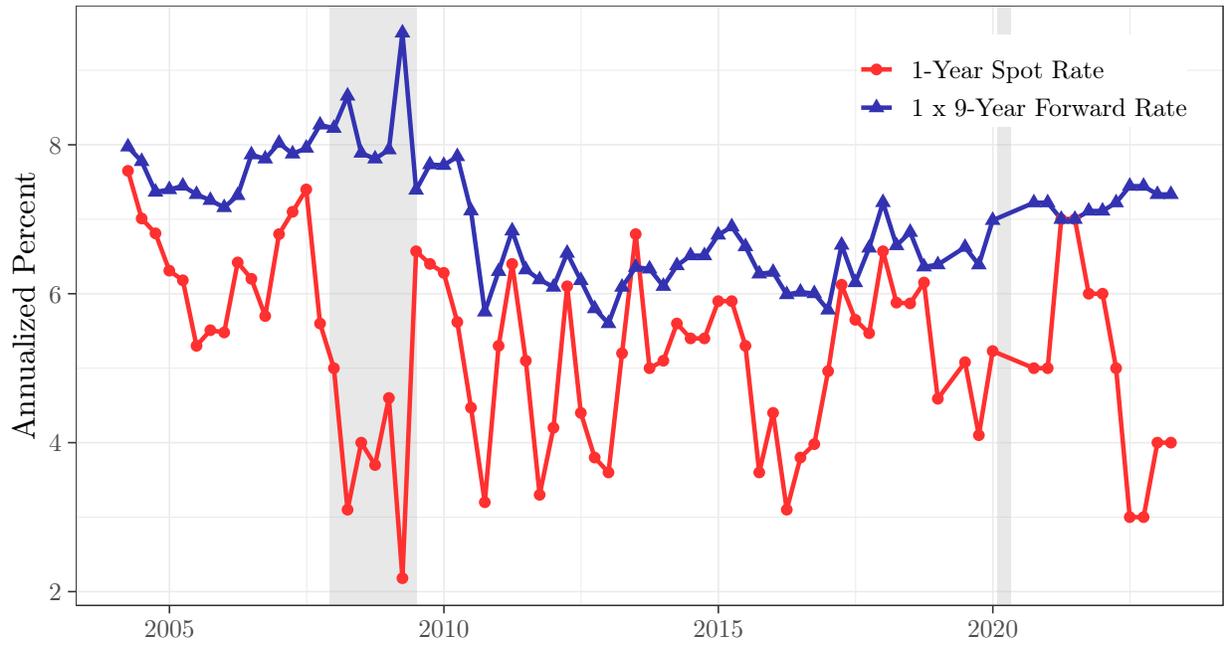
Figure 8
Discounted Forecast Errors and the Price-Dividend Ratio

This figure plots discounted predicted forecast errors over all future horizons $\mathcal{E}_t$ (red) and the log repurchase-adjusted price-dividend ratio $p_t - d_t$ (blue) in the United States. Discounted forecast errors $\mathcal{E}_t$ are calculated as in (13) using the 6×6 predicted forecast error (non-annualized), with ρ calculated using the full-sample average price-dividend ratio. The monthly repurchase-adjusted $p_t - d_t$ for the CRSP value-weighted stock index is obtained from Nagel and Xu (2022a) via Zhengyang Xu's website. Both series are plotted as log differences from their full-sample means. Gray bands are NBER recessions. The sample is from January 1990 to June 2021 in the U.S.
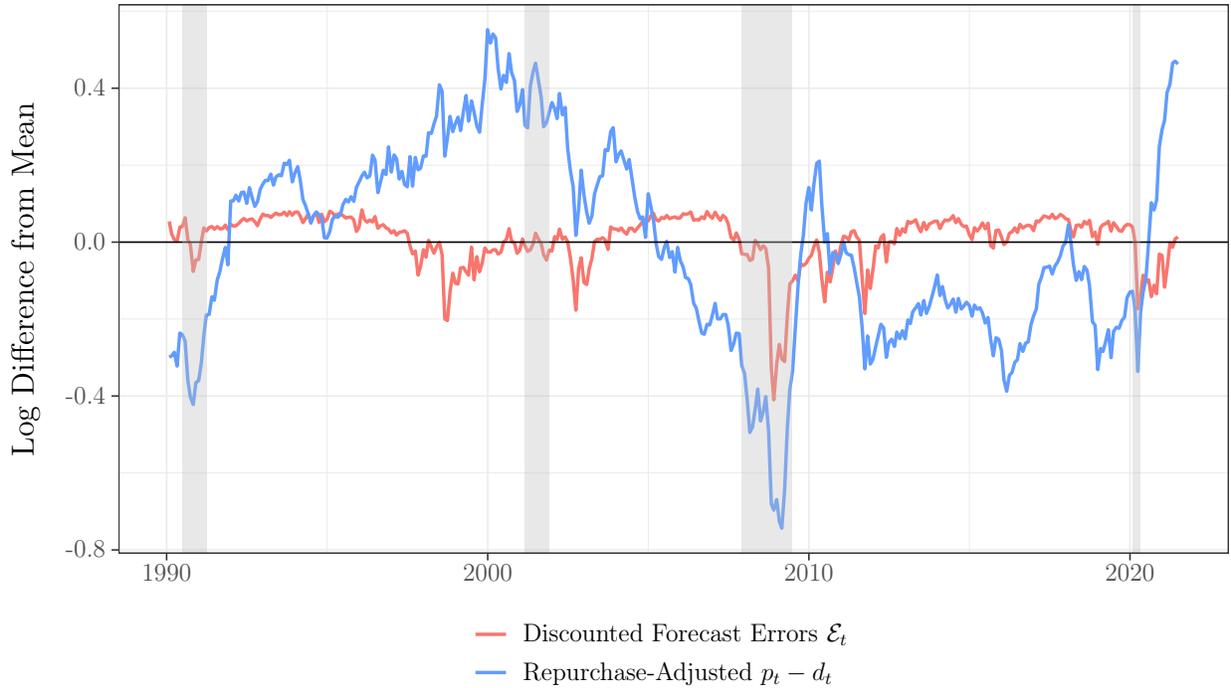
Figure 9
The Effect of Forecast Errors on Realized Returns for the Equity Term Structure

This figure shows the impact of predicable variation in forecast errors on realized returns to dividend claims with different maturity. We calculate realized returns for each equity claim based on the assumption that spot rates are the same for all dividend claims and given by the spot rates from Proposition 1. Difference in realized returns arise from differences in exposure to forecast errors (duration) across claims. We divide the sample into good and bad times based on the ex ante CAPE ratio, with good times being the 80% of the months where the CAPE ratio is highest. See text for additional details.
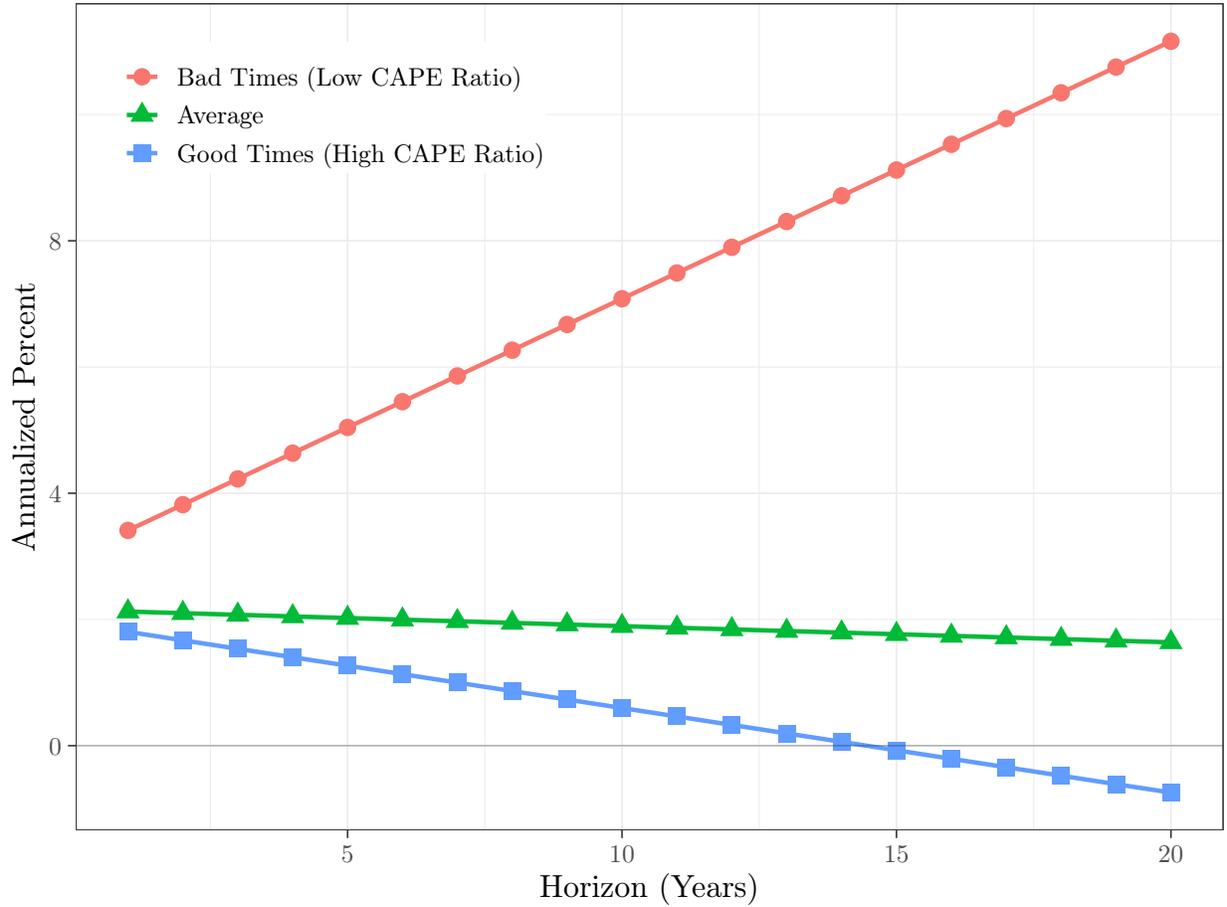
Figure 10
Model Calibration: Regression Slopes and Average Forecast Errors

This figure reports regression slopes and average forecast errors in the calibrated model of expectation errors. The model is calibrated from the standpoint of an unconstrained log utility investor fully invested in the market. Table A11 reports the objective parameters. The solid lines are model-implied population moments in a single long sample. The shaded regions are model-implied 95% confidence bands in 10,000 short samples. The blue circles are model-implied moments under rational expectations with $\theta = 0$. The red squares are model-implied moments under diagnostic expectations with $\theta = 0.91$ from Bordalo, Gennaioli, and Shleifer (2018). The green triangles are moments in the data. The sample is the longest available for each exchange in the main sample.
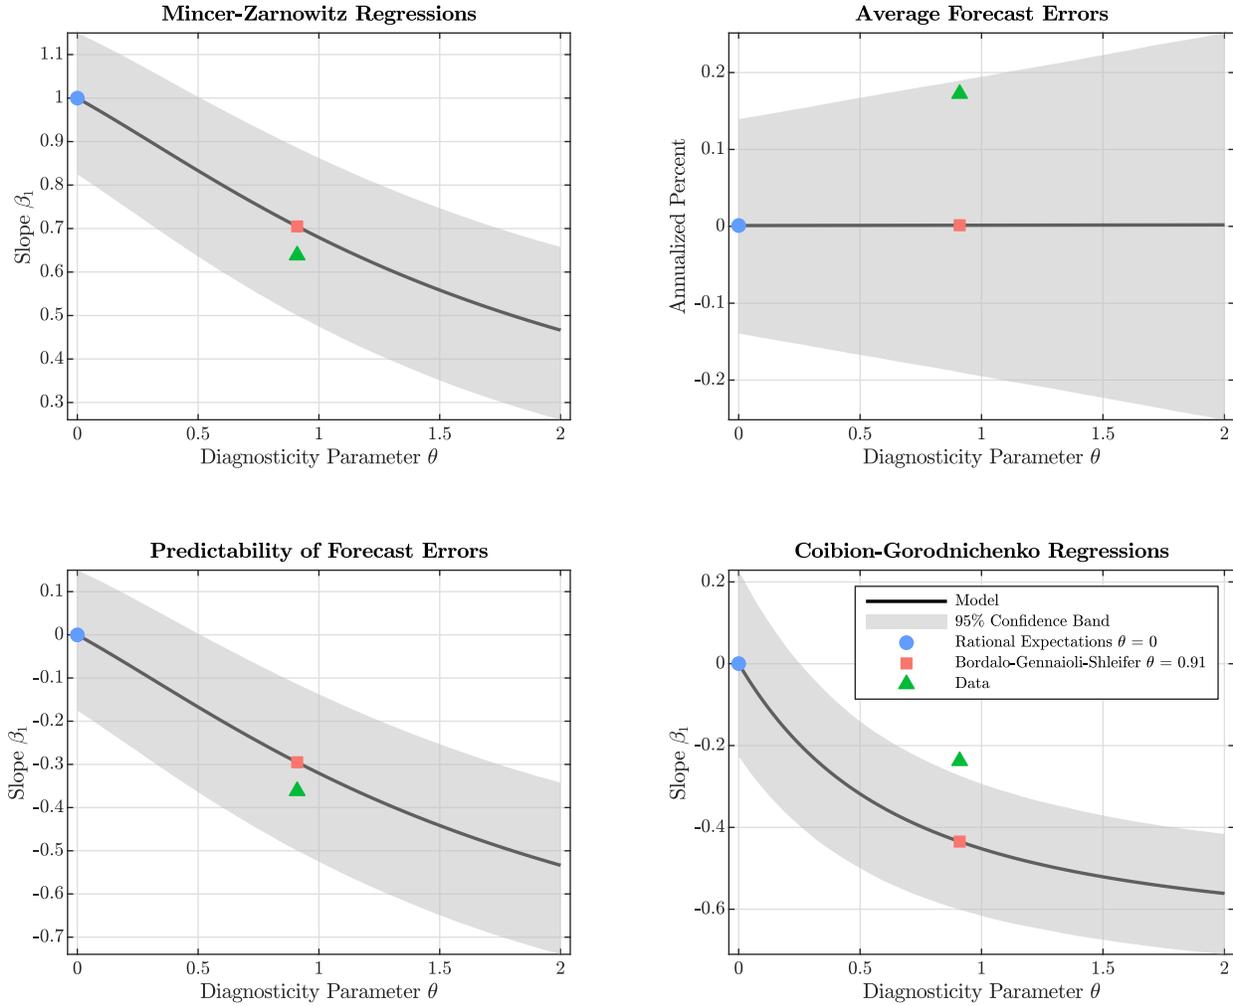
References

Adam, Klaus and Stefan Nagel. 2023. "Expectations Data in Asset Pricing." In *Handbook of Economic Expectations*, edited by R. Bachmann, G. Topa, and W. van der Klaauw. Elsevier, 477–506.

Aït-Sahalia, Yacine, Mustafa Karaman, and Loriano Mancini. 2020. "The term structure of equity and variance risk premia." *Journal of Econometrics* 219 (2):204–230.

Augenblick, Ned and Eben Lazarus. 2023. "A New Test of Excess Movement in Asset Prices." Working paper, MIT.

Augenblick, Ned, Eben Lazarus, and Michael Thaler. 2023. "Overinference from Weak Signals and Underinference from Strong Signals." Working Paper 2109.09871, arXiv.

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.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer. 2019. "Diagnostic Expectations and Stock Returns." *Journal of Finance* 74 (6):2839–2874.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2018. "Diagnostic Expectations and Credit Cycles." *Journal of Finance* 73 (1):199–227.

Boutros, Michael, Itzhak Ben-David, John Graham, Campbell Harvey, and John Payne. 2020. "The Persistence of Miscalibration." Working Paper 28010, National Bureau of Economic Research.

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. 2017. "Discussion of Lawrence Jin and Pengfei Sui, 'Asset Pricing with Return Extrapolation'." NBER Behavioral Finance Meeting.

Campbell, John Y. and John Ammer. 1993. "What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns." *Journal of Finance* 48 (1):3–37.

Campbell, John Y. and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107 (2):205–251.

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.

———. 1991. "Yield Spreads and Interest Rate Movements: A Bird's Eye View." *Review of Economic Studies* 58 (3):495–514.

Carr, Peter and Dilip Madan. 2001. "Optimal positioning in derivative securities." *Quantitative Finance* 1 (1):19.

Chabi-Yo, Fousseni and Johnathan Loudis. 2020. "The conditional expected market return." *Journal of Financial Economics* 137 (3):752–786.

Cieslak, Anna. 2018. "Short-Rate Expectations and Unexpected Returns in Treasury Bonds." *Review of Financial Studies* 31 (9):3265–3306.

Coibion, Olivier and Yuriy Gorodnichenko. 2015. "Information rigidity and the expectations formation process: A simple framework and new facts." *American Economic Review* 105 (8):2644–78.

Couts, Spencer J., Andrei S. Gonçalves, and Johnathan Loudis. 2023. "The Subjective Risk and Return Expectations of Institutional Investors." Working Paper 2023-03, Fisher College of Business.

Dahlquist, Magnus and Markus Ibert. 2023. "Equity Return Expectations and Portfolios: Evidence from Large Asset Managers." Working paper, Swedish House of Finance.

d'Arienzo, Daniele. 2020. "Maturity Increasing Overreaction and Bond Market Puzzles." Working paper, Nova School of Business and Economics.

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

Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez. 2017. "The price of variance risk." *Journal of Financial Economics* 123 (2):225–250.

Fama, Eugene F. and Robert R. Bliss. 1987. "The Information in Long-Maturity Forward Rates." *American Economic Review* 77 (4):680–692.

Farmer, Leland, Emi Nakamura, and Jón Steinsson. 2023. "Learning about the Long Run." Working Paper 29495, National Bureau of Economic Research.

Fuster, Andreas, David Laibson, and Brock Mendel. 2010. "Natural Expectations and Macroeconomic Fluctuations." *Journal of Economic Perspectives* 24 (4):67–84.

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

Golez, Benjamin and Jens Carsten Jackwerth. 2023. "Holding Period Effects in Dividend Strip Returns." Working paper, University of Notre Dame.

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

Gormsen, Niels Joachim and Kilian Huber. 2023. "Corporate Discount Rates." Working Paper 31329, National Bureau of Economic Research.

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

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.

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.

Mixon, Scott. 2007. "The implied volatility term structure of stock index options." *Journal of Empirical Finance* 14 (3):333–354.

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

———. 2022b. "Dynamics of Subjective Risk Premia." Working Paper 29803, National Bureau of Economic Research.

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

Polkovnichenko, Valery and Feng Zhao. 2013. "Probability weighting functions implied in options prices." *Journal of Financial Economics* 107 (3):580–609.

Stein, Jeremy. 1989. "Overreactions in the Options Market." *Journal of Finance* 44 (4):1011–1023.

Internet Appendix:

Forward Return Expectations[*]

Mihir Gandhi, Niels Joachim Gormsen, and Eben Lazarus

September 2023

Contents

─────────────────

[*]Contact: mihir.a.gandhi@gmail.com; niels.gormsen@chicagobooth.edu; lazarus@berkeley.edu.

A Proofs

Proof of Proposition 1. Given that $M_{t,t+n}R_{t,t+n} = 1$ by assumption, $\mathcal{C}_t^{(n)} = 0$ in (4). The stated results then follow immediately. $\square$

Proof of Proposition 2. First consider $n = m = 1$, and write

$$\varepsilon_{t+1}^{(1)} = \mu_{t+1}^{(1)} - f_t^{(1)} = \mathcal{L}_{t+1}^{(1)} - \mathcal{L}_t^{(2)} + \mathcal{L}_t^{(1)} + \operatorname{cov}_t(MR_{t,t+2}, r_{t,t+2})$$
$$- \operatorname{cov}_t(MR_{t,t+1}, r_{t,t+1}) - \operatorname{cov}_{t+1}(MR_{t+1,t+2}, r_{t+1,t+2}). \tag{A1}$$

Consider the first covariance term. Given the joint log-normality of the SDF and returns (and the normality of $r_{t,t+n}$), Stein's lemma gives that

$$\operatorname{cov}_t(MR_{t,t+2}, r_{t,t+2}) = \operatorname{cov}_t(MR_{t,t+1}MR_{t+1,t+2}, r_{t,t+2})$$
$$= \operatorname{cov}_t(mr_{t,t+1} + mr_{t+1,t+2}, r_{t,t+2})\,\mathbb{E}_t[MR_{t,t+2}]$$
$$= \operatorname{cov}_t(mr_{t,t+1} + mr_{t+1,t+2}, r_{t,t+1} + r_{t+1,t+2}),$$

where $mr_{t,t+n} = \ln(MR_{t,t+n})$, and where the last line uses that $\mathbb{E}_t[MR_{t,t+2}] = 1$. Having separated the two MR terms, apply Stein's lemma again to obtain

$$\operatorname{cov}_t(MR_{t,t+2}, r_{t,t+2}) = \operatorname{cov}_t(mr_{t,t+2}, r_{t,t+1}) + \operatorname{cov}_t(mr_{t,t+1}, r_{t+1,t+2}) + \operatorname{cov}_t(mr_{t+1,t+2}, r_{t,t+2})$$
$$= \operatorname{cov}_t(MR_{t,t+2}, r_{t,t+1}) + \operatorname{cov}_t(MR_{t,t+1}, r_{t+1,t+2})$$
$$+ \operatorname{cov}_t(MR_{t+1,t+2}, r_{t+1,t+2}). \tag{A2}$$

For the first two terms in (A2), by the law of total covariance and using that $\mathbb{E}_{t+1}[MR_{t+1,t+2}] = 1$,

$$\operatorname{cov}_t(MR_{t,t+2}, r_{t,t+1}) = \mathbb{E}_t\left[MR_{t,t+1}r_{t,t+1}\operatorname{cov}_{t+1}(MR_{t+1,t+2}, 1)\right]$$
$$+ \operatorname{cov}_t(MR_{t,t+1}\,\mathbb{E}_{t+1}[MR_{t+1,t+2}], r_{t,t+1})$$
$$= \operatorname{cov}_t(MR_{t,t+1}, r_{t,t+1}), \tag{A3}$$
$$\operatorname{cov}_t(MR_{t,t+1}, r_{t+1,t+2}) = \mathbb{E}_t[MR_{t,t+1}\operatorname{cov}_{t+1}(1, r_{t+1,t+2})] + \operatorname{cov}_t(MR_{t,t+1}, \mathbb{E}_{t+1}[r_{t+1,t+2}])$$
$$= \operatorname{cov}_t(MR_{t,t+1}, \mathbb{E}_{t+1}[r_{t+1,t+2}]). \tag{A4}$$

Turning now to the last term in (A1), the law of total covariance can similarly be applied to obtain that as of time t,

$$\mathbb{E}_t[\operatorname{cov}_{t+1}(MR_{t+1,t+2}, r_{t+1,t+2})] = \operatorname{cov}_t(MR_{t+1,t+2}, r_{t+1,t+2}). \tag{A5}$$

Taking expectations in (A1), substituting in results (A2)–(A5), and applying the definition of $\widehat{\varepsilon}_{t+1}^{(1)}$, we obtain:

$$\mathbb{E}_t\left[\varepsilon_{t+1}^{(1)}\right] = \mathbb{E}_t\left[\widehat{\varepsilon}_{t+1}^{(1)}\right] + \operatorname{cov}_t(MR_{t,t+1}, \mathbb{E}_{t+1}[r_{t+1,t+2}]). \tag{A6}$$

Rearranging to solve for $\mathbb{E}_t[\widehat{\varepsilon}_{t+1}^{(1)}]$ yields the stated result for the $n = m = 1$ case. While this case is convenient for straightforward derivations, note that all the above steps apply when using $t + n$ in place of $t + 1$ and using $t + n + m$ in place of $t + 2$, so the stated result holds for general n, m. $\square$

Proof of Proposition 3. Starting again with (A1) and expanding the first covariance term,

$$\text{cov}_t(MR_{t,t+2}, r_{t,t+2}) = \text{cov}_t(MR_{t,t+1}MR_{t+1,t+2}, r_{t,t+1}) + \text{cov}_t(MR_{t,t+1}MR_{t+1,t+2}, r_{t+1,t+2}). \quad \text{(A7)}$$

We consider each of the two terms on the right side of (A7) in turn, and in both cases apply the law of total covariance. For the first term, as in (A3),

$$\text{cov}_t(MR_{t,t+1}MR_{t+1,t+2}, r_{t,t+1}) = \text{cov}_t(MR_{t,t+1}, r_{t,t+1}). \quad \text{(A8)}$$

For the second term,

$$\begin{aligned}
\text{cov}_t(MR_{t,t+1}MR_{t+1,t+2}, r_{t,t+2}) &= \mathbb{E}_t\big[MR_{t,t+1}\,\text{cov}_{t+1}(MR_{t+1,t+2}, r_{t+1,t+2})\big] \\
&\quad + \text{cov}_t(MR_{t,t+1}, \mathbb{E}_{t+1}[r_{t+1,t+2}]). \quad \text{(A9)}
\end{aligned}$$

Using (A8) and (A9) in (A1), applying the definition of $\widehat{\varepsilon}_{t+1}^{(1)}$, and taking expectations,

$$\begin{aligned}
\mathbb{E}_t\Big[\varepsilon_{t+1}^{(1)}\Big] &= \mathbb{E}_t\Big[\widehat{\varepsilon}_{t+1}^{(1)}\Big] + \mathbb{E}_t[(MR_{t,t+1} - 1)\,\text{cov}_{t+1}(MR_{t+1,t+2}, r_{t+1,t+2})] \\
&\quad + \text{cov}_t(MR_{t,t+1}, \mathbb{E}_{t+1}[r_{t+1,t+2}]) \\
&= \mathbb{E}_t\Big[\widehat{\varepsilon}_{t+1}^{(1)}\Big] + \text{cov}_t(MR_{t,t+1}, \mathbb{E}_{t+1}[r_{t+1,t+2}] + \text{cov}_{t+1}(MR_{t+1,t+2}, r_{t+1,t+2})). \quad \text{(A10)}
\end{aligned}$$

Note from (4) that $\mathscr{L}_{t+1}^{(1)} = \mathbb{E}_{t+1}[r_{t+1,t+2}] + \text{cov}_{t+1}(MR_{t+1,t+2}, r_{t+1,t+2})$. Using this in (A10),

$$\mathbb{E}_t\Big[\widehat{\varepsilon}_{t+1}^{(1)}\Big] = \mathbb{E}_t\Big[\varepsilon_{t+1}^{(1)}\Big] - \text{cov}_t\Big(MR_{t,t+1}, \mathscr{L}_{t+1}^{(1)}\Big).$$

The above steps again apply when using $t + n$ in place of $t + 1$ and using $t + n + m$ in place of $t + 2$, completing the proof. $\square$

Proof of Lemma A1. To compute the risk-neutral expectation of $H[P_T] = R^\alpha (\ln R)^\beta$, we apply standard spanning theorems (Bakshi and Madan 2000, Carr and Madan 2001). We have

$$\begin{aligned}
\frac{1}{R_f}\mathbb{E}_t^*\Big[R^\alpha (\ln R)^\beta\Big] &= \left(H[\bar{P}] - \bar{P}H_P[\bar{P}]\right)\frac{1}{R_f} + H_P[\bar{P}]\,P_t \\
&\quad + \int_0^{\bar{P}} H_{PP}[K]\,\text{put}_t^{(n)}(K)dK + \int_{\bar{P}}^\infty H_{PP}[K]\,\text{call}_t^{(n)}(K)dK.
\end{aligned}$$

The result follows by setting $\bar{P} = F_t^{(n)}$ and simplifying. $\square$

Proof of Proposition A1. (A13) is immediate from Martin (2017) Result 8. (A14) and (A15) follow from Lemma A1 by setting the appropriate α and β and simplifying. $\square$

B   Measurement Details

B.1   Data

United States Data. For the 1996 to 2021 period, we obtain end-of-day option prices, index prices, projected dividend yields, and risk-free rates from OptionMetrics. To maximize the sample size, we use options with both AM and PM settlement. We use the bid/ask midpoint as the option price in

the main analysis. We linearly interpolate the risk-free rate curve to match option maturities. If either the dividend yield or risk-free rate is missing, we use the last non-missing observation.

For the 1990 to 1995 period, we obtain intraday option quotes from CBOE Market Data Express, as in Kelly, Pástor, and Veronesi (2016) and Culp, Nozawa, and Veronesi (2018). We obtain end-of-day index prices/returns from CRSP and estimate dividend yields from lagged one-year cum/ex-dividend index returns. We obtain Treasury bill rates and constant maturity Treasury yields from FRED to construct risk-free rates, as in Culp, Nozawa, and Veronesi.

Unlike OptionMetrics, CBOE provides intraday quotes. To construct end-of-day prices, we first apply filters to the intraday data and then use the last available quote. We drop quotes with the missing codes of 998 or 999. We drop quotes with negative bid-ask spreads. We correct erroneously recorded quotes – quotes with strike price less than 100 – by multiplying the strike/option price by 10. We drop end-of-day quotes that increase and then decrease fourfold (or vice versa), following similar filters in Andersen, Bondarenko, and Gonzalez-Perez (2015) and Duarte, Jones, and Wang (2022). We interpret these large reversals as probable data errors. To validate these filters, we compare data from CBOE and OptionMetrics in 1996. We match approximately 99.3% of option prices in OptionMetrics, suggesting these filters are not unreasonable.

We apply standard filters to the end-of-day data (Constantinides, Jackwerth, and Savov 2013). (1) We drop options with special settlement. (2) To eliminate duplicate quotes, we select the quote with highest open interest. (3) We drop options with fewer than seven days-to-maturity. (4) We drop options with price less than 0.01. (5) We drop options with zero bid prices or negative bid-ask spreads. (6) We drop options that violate static no-arbitrage bounds:

$$\text{put}_t^{(n)}(K) \le Ke^{-r\tau} \qquad \text{call}_t^{(n)}(K) \le P_t.$$

(7) We drop options for which the Black-Scholes implied volatility computation does not converge and options with implied volatility less than 5% or greater than 100%.

International Data. We again obtain end-of-day option prices, index prices, projected dividend yields, and risk-free rates from OptionMetrics. Unlike the United States, most option prices are either end-of-day settlement prices or last traded prices. Only a small fraction are from either bid/ask prices. The index price is time synchronized to the option price. If the index price is missing, we obtain the end-of-day price from Compustat Global. Risk-free rates are from currency-matched LIBOR curves. Dividend yields are from put-call parity and so are maturity-specific. As before with risk-free rates, we linearly interpolate the dividend yield curve to match option maturities. We apply the same filters to the end-of-day data as with the United States, except for filters that require bid/ask prices.

Table A1 describes the international sample. The European sample begins in January 2002 and ends in September 2021. The Asian sample begins in January 2004 and ends in April 2021. Our international sample closely follows Kelly, Pástor, and Veronesi (2016) and Dew-Becker and Giglio (2021), but we also use pan-European Stoxx indexes. These indexes represent a substantive addition to the sample. At long maturities, the Euro Stoxx 50 is arguably the most liquid options market in the world, as is the case with the dividend futures market (Binsbergen and Koijen 2017, Binsbergen et al. 2013).

Main Sample. As the international data is not equally robust across exchanges, we select the most reliable exchanges for the main analysis. We select the main sample by elimination. We drop Netherlands and Japan because they do not have sufficiently dense options, as seen in Panel A of Figure A1. We drop Finland, the Stoxx Europe 50, and the Stoxx Europe 600 because they do not

have reliable open interest data, as seen in Panel B of Figure A1. We drop Belgium, Korea, and Taiwan because they do not consistently have long-maturity options, as seen in Panel D of Figure A2. We drop China and Sweden because they do not have sufficiently deep out-of-the-money options, as seen in Figure A2. This leaves the 10 exchanges in the last column of Table A1 for the main sample. As a robustness check, we examine the full sample of 20 exchanges in Table A8 and Table A10.

B.2   Baseline Measures

Methodology. On each date and separately for puts/calls,

1. We convert option prices to implied volatilities via Black-Scholes. Here we follow an extensive literature on option-implied risk-neutral densities that finds interpolation more conducive in the space of implied volatilities, not option prices (Figlewski 2010, Malz 2014).

2. We fit a Delaunay triangulation to implied volatilities. The grid consists of strike prices between $\underline{K} = 0.10 \times P_t$ and $\overline{K} = 2.00 \times P_t$ with $\Delta K = 0.001 \times P_t$ and maturities $\tau = 30$, 60, 91, 122, 152, 182, 273, 365 days. The triangulation extrapolates as necessary with the nearest implied volatility in moneyness and time-to-maturity space.

3. We convert the triangulation of implied volatilities back to option prices via Black-Scholes. We then use the implied triangulation of option prices to evaluate the LVIX integral in (7) via Gaussian quadrature.

4. With the LVIX in hand, we can immediately compute spot rates, forward rates, and forecast errors under log utility via Proposition 1, as shown in Figure A3. Figure A4 plots contemporaneous 6-month spot rates and 6×6-month forward rates in the full sample, analogous to Figure 2 in the United States.

5. We occasionally find negative forward rates. Gao and Martin (2021) argue that negative forward rates are unlikely theoretically and likely represent data errors. We follow Gao and Martin and drop such observations, but our results are not quantitatively sensitive to this choice.

Discussion. Three empirical challenges in the computation of option-implied moments – discretization, truncation, and interpolation bias – motivate our baseline methodology (Carr and Wu 2009, Jiang and Tian 2007). We discuss each in turn. First, discretization bias arises because (7) requires numerical integration. To minimize this bias, we integrate on a fine grid of interpolated option prices in step 3. Second, truncation bias arises because (7) requires integration over an infinite range of strike prices in theory. In practice, we truncate the integral. To minimize this bias, we extrapolate and integrate over strike prices well beyond the range of observable option prices in step 2. Finally, interpolation bias arises because (7) usually requires options with unavailable maturities. To address this bias, we interpolate the option surface at target maturities in step 2.

Measurement Error. To better understand the role of measurement error, Figure A5 examines spot/forward rates in simulations. We first compute option prices from a parametric model. Since we know the true data generating process, we then quantify how varying integration bounds affects the integral relative to the true value. The thought experiment follows a similar exercise in Jiang and Tian (2007) for the VIX.

We first truncate the integral (7) without extrapolation, as in Table A3. In Panel A, we consider a Black-Scholes model. We find a large truncation bias in bad times. In bad times, volatility is

high, deep out-of-the-money options are expensive, and so the bias is large. In contrast, in good times, volatility is low, deep out-of-the-money options are cheap, and so the bias is small. The bias is especially large for 12-month spot rates and forward rates because longer-maturity option prices have more time value. In Panel B, we consider a stochastic volatility model with jumps (SVJ). We again find an uncomfortably large truncation bias. Relative to Black-Scholes, the bias is larger when volatility is low, but smaller when volatility is high because volatility mean-reverts in SVJ.

We next truncate the integral after extrapolating beyond the range of observable strikes, as in the baseline analysis. We continue with a SVJ model in Panel C. Relative to Panel B, we find that extrapolation reduces truncation bias across the board.

We emphasize, however, that this exercise only motivates extrapolation in our baseline integration scheme and our use of shorter-maturity forward rate as instruments/predictors, as in Table 3 and Table 5. We make no claim that measurement error is unconditionally small. By construction, these simulations address only truncation bias. There is no scope for either discretization or interpolation bias, as we simulate option prices on a counterfactually dense grid. We think these biases may be non-trivial at times and especially so when options are less dense.

B.3 Alternative Measures

Table A3 reports robustness checks where we use alternative choices to measure spot rates, forward rates, and forecast errors. As we discuss below, the main results are largely robust to these choices.

Integration Bounds. Panel A repeats the analysis with alternative integration bounds. The first four rows consider static bounds without extrapolation. As an example, the first row evaluates the integral in (7) between strike prices $\underline{K} = 0.65 \times P_t$ and $\overline{K} = 1.35 \times P_t$ at each maturity. The fifth row uses observable option prices between strike prices $\underline{K} = 0.10 \times P_t$ and $\overline{K} = 2.00 \times P_t$, again without extrapolation. The bounds in the first five rows naturally vary both by time and maturity with the availability of option prices. The sixth row considers static bounds with extrapolation, following a similar robustness check in Gormsen and Jensen (2022):

$$\left[\underline{K}^{(n)}, \overline{K}^{(n)}\right] = \begin{cases} [0.75, 1, 25] \times P_t & n \in \{1, 2\} \\ [0.55, 1.45] \times P_t & n \in \{3, 4, 5\} \\ [0.35, 1.65] \times P_t & n \in \{6, 9\} \\ [0.20, 1.80] \times P_t & n \in \{12\}. \end{cases}$$

These bounds vary by maturity, but not by time. The seventh row considers dynamic bounds with extrapolation, again following a similar robustness check in Gormsen and Jensen:

$$\underline{K}^{(n)} = \max\left\{0.10, 1.00 - 5\sigma_t^{(n)}\sqrt{\tau}\right\} \times P_t \qquad \overline{K}^{(n)} = \min\left\{2.00, 1.00 + 5\sigma_t^{(n)}\sqrt{\tau}\right\} \times P_t,$$

where $\sigma_t^{(n)}$, the price of the volatility contract in Bakshi, Kapadia, and Madan (2003), proxies for the risk-neutral volatility of the market return:

$$\begin{aligned} \left(\sigma_t^{(n)}\sqrt{\tau}\right)^2 &= \frac{1}{R_{t,t+n}^f}\mathbb{E}_t^*\left[(\ln R_{t,t+n})^2\right] \\ &= \int_0^{P_t} \frac{2\left(1 + \ln\left[\frac{P_t}{K}\right]\right)}{K^2}\, \mathrm{put}_t^{(n)}(K)dK + \int_{P_t}^{\infty} \frac{2\left(1 + \ln\left[\frac{P_t}{K}\right]\right)}{K^2}\, \mathrm{call}_t^{(n)}(K)dK. \end{aligned} \qquad \text{(A11)}$$

These bounds vary by both time and maturity with volatility. The eighth row considers the baseline integration bounds, as discussed in the main text and Appendix B.2.

In sum, this exercise illustrates the significant effect truncation/extrapolation have on the regression estimates. With shallow bounds, forecast errors are relatively large on average but less predictable. With deep bounds, forecast errors are relatively small on average but more predictable.

Liquidity Filters. Panel B repeats the analysis with alternative liquidity filters. The first row considers an outlier filter, following similar filters in Constantinides, Jackwerth, and Savov (2013) and Beason and Schreindorfer (2022). On each date and separately for puts/calls, we first fit a quadratic function to implied volatilities in terms of moneyness K/P and time-to-maturity. To minimize the effect of deep out-of-the-money, short/long-maturity options, we only use options with maturity $14 \leq \tau \leq 365$ days and moneyness $0.65 \leq K/P \leq 1.35$. We then drop influential observations via Cook's Distance. The second row considers an open interest filter. We drop options with zero open interest. We do not have open interest data before 1996. The third row combines the outlier and open interest filters. In all, this exercise is consistent with the baseline results, suggesting that option illiquidity does not explain our findings.

Volatility Surface. The first row in Panel C repeats the analysis with the interpolated volatility surface from OptionMetrics. OptionMetrics provides interpolated Black-Scholes implied volatilities on a constant moneyness/maturity grid. The literature often uses this surface for options with American exercise because OptionMetrics reports an equivalent, European exercise, implied volatility (Kelly, Lustig, and Van Nieuwerburgh 2016, Martin and Wagner 2019). We instead simply use it as a robustness check on our own Delaunay triangulation of the volatility surface. In short, this exercise is consistent with the baseline results, although the average forecast error is somewhat smaller.

SVI Surface. The second row in Panel C repeats the analysis with the stochastic volatility inspired (SVI) surface from Jim Gatheral at Merrill Lynch (Gatheral 2011, Gatheral and Jacquier 2011, 2014). Our implementation of the SVI surface closely follows Berger, Dew-Becker, and Giglio (2020) and Beason and Schreindorfer (2022). We parameterize squared Black-Scholes implied volatilities with the function

$$\sigma_{BS}^2 (t, \kappa, \tau) = a + b \left(\rho (\kappa - m) + \sqrt{(\kappa - m)^2 + \sigma^2} \right), \tag{A12}$$

where κ is standardized forward moneyness

$$\kappa = \frac{\ln K - \ln F_t^{(n)}}{\sigma_t^{(n)} \sqrt{\tau}},$$

$\sigma_t^{(n)}$ proxies for the risk-neutral volatility of the market return as in (A11), and each parameter is a linear function of time-to-maturity (e.g., $a = a_0 + a_1 \tau$). On each date, we estimate parameters $\theta = (a_0, a_1, b_0, b_1, \rho_0, \rho_1, m_0, m_1, \sigma_0, \sigma_1)$ that minimize the implied volatility RMSE between the surface (A12) and the data, subject to standard no-arbitrage constraints: option prices are nonnegative and monotonic/convex in K (Aït-Sahalia and Duarte 2003). We check these constraints on a grid with moneyness between $-20 \leq \kappa \leq 0.50$ for puts, between $-0.50 \leq \kappa \leq 10$ for calls, and maturities $\tau = 30, 60, 91, 122, 152, 182, 273, 365$ days. We estimate the surface with outlier-filtered, as discussed in Appendix B.3, out-of-the-money puts/calls: puts with $\kappa \leq 0$ and calls with $\kappa \geq 0$. We estimate the surface separately for puts/calls and separately for short/long-maturity options: $14 \leq \tau \leq 122$ days and $122 < \tau \leq 365$ days, respectively.

Bid/Ask Prices. Panel D repeats the analysis with bid/ask prices, following similar robustness checks in Martin (2017) and Gao and Martin (2021). We only have bid/ask prices in the United States. The first row reports the baseline results with the bid-ask midpoint. The second row repeats the analysis with bid prices, the third ask prices. In sum, this exercise is consistent with the baseline results, although the Coibion-Gorodnichenko regression slope is somewhat smaller with ask prices.

C Additional Empirical Results and Robustness Checks

C.1 Coibion-Gorodnichenko Regressions

Methodology. Here, we study the predictability of forecast errors using forward-rate *changes* (rather than levels, as in the main text), using the method in Coibion and Gorodnichenko (2015). This allows us to address more specifically whether forward rates exhibit excess sensitivity to news. These regressions take the following form:

$$\tilde{\varepsilon}_{i,t+n}^{m} = \beta_0 + \beta_1 \left(\tilde{f}_{i,t}^{(n,m)} - \tilde{f}_{i,t-h}^{(n+h,m)} \right) + e_{i,t+m},$$

for both option-based risk premia and expected returns. We consider multiple horizons for the forecast revision and forecast error.

Results. Table A4 presents results. We find, consistent with the main analysis, meaningful forecast-error predictability with significantly negative coefficients at multiple (though not all) horizons. The long-horizon risk-premium revisions generate more consistent significance in predictability, consistent with the fact that forward rates are persistent (so short-horizon revisions contain quite a bit of noise unrelated to the predictable component of forecast errors).

Measurement Error. In Coibion-Gorodnichenko regressions, we use forecast revisions to predict forecast errors. Since forecast revisions/errors involve the same forward rate, measurement error may produce spurious evidence of predictability. To better understand the role of measurement error, we again turn to simulations. We quantify how much correlated measurement error would be necessary to produce the Coibion-Gorodnichenko regression slopes in the data.

We assume we observe forecast revisions and forecast errors with noise:

$$\widetilde{x} = x\sigma_x + v\sigma_v \text{ and } \widetilde{y} = y\sigma_y - v\sigma_v,$$

respectively, with $\sigma_{xy} = \sigma_{xv} = \sigma_{yv} = 0$ and $x, y, v \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$. We vary σ_v^2 exogenously. In each simulation draw, we set the variance of the truth (σ_x^2 and σ_y^2) such that the observed variance ($\sigma_{\widetilde{x}}^2$ and $\sigma_{\widetilde{y}}^2$) equals that in the data. As σ_v^2 varies, these weights ensure all variation in slopes comes from variation in noise and none from variation in observed variances. Any evidence of predictability – any non-zero slope – is spurious because $\sigma_{xy} = 0$.

Figure A6 reports the results from these simulations. To produce the Coibion-Gorodnichenko regression slopes in the data, we require σ_v be about 40 basis points or more than one-quarter the volatility of forecast errors in the data. This, at least to us, seems implausibly large. We conclude that correlated measurement error cannot fully explain forecast-error predictability in Coibion-Gorodnichenko regressions, although we cannot fully rule out some bias due to measurement error.

C.2 Power Utility Regressions

Methodology. This section derives the power utility analogue to the LVIX. To do so, we apply results from Martin (2017) and Gao and Martin (2021). We omit time subscripts throughout to minimize clutter.

LEMMA A1 (Spanning $R^\alpha (\ln R)^\beta$). *For any α and β,*

$$\frac{1}{R_f} \mathbb{E}_t^* \left[R^\alpha (\ln R)^\beta \right] = R_f^\alpha (\ln R_f)^\beta + \int_0^{F_t^{(n)}} \omega(\alpha, \beta) \, \mathrm{put}_t^{(n)}(K) dK + \int_{F_t^{(n)}}^\infty \omega(\alpha, \beta) \, \mathrm{call}_t^{(n)}(K) dK,$$

where

$$\omega(\alpha, \beta) = -\frac{\alpha (1-\alpha) m^\beta + \beta (1-2\alpha) m^{\beta-1} + \beta (1-\beta) m^{\beta-2}}{P_t^2} \left(\frac{K}{P_t} \right)^{\alpha-2}$$

and $m = \ln K - \ln P_t$.

As is well-known, under certain regularity conditions, we can compute the price of any function of the index price via a replicating portfolio of bonds, stocks, and options. We simply apply this result to the function $R^\alpha (\ln R)^\beta$, which is useful for expectations under power utility below.

PROPOSITION A1 (Expected Equity Premium with Power Utility). *From the standpoint of an unconstrained power utility investor fully invested in the market,*

$$\mathbb{E}_t [\ln R] - \ln R_f = \frac{\mathbb{E}_t^* [R^\gamma \ln R]}{\mathbb{E}_t^* [R^\gamma]} - \ln R_f, \tag{A13}$$

where

$$\frac{1}{R_f} \mathbb{E}_t^* [R^\gamma \ln R] = R_f^\gamma \ln R_f + \int_0^{F_t^{(n)}} \omega(\gamma, 1) \, \mathrm{put}_t^{(n)}(K) dK + \int_{F_t^{(n)}}^\infty \omega(\gamma, 1) \, \mathrm{call}_t^{(n)}(K) dK \tag{A14}$$

and

$$\frac{1}{R_f} \mathbb{E}_t^* [R^\gamma] = R_f^\gamma + \int_0^{F_t^{(n)}} \omega(\gamma, 0) \, \mathrm{put}_t^{(n)}(K) dK + \int_{F_t^{(n)}}^\infty \omega(\gamma, 0) \, \mathrm{call}_t^{(n)}(K) dK \tag{A15}$$

and γ is the investor's risk aversion.

The LVIX uses a special case of (A13) with $\gamma = 1$, and so the mechanics under power utility are similar, if only messier, to that under log utility. However, there is one caveat: as risk aversion γ increases, the weights $\omega(\gamma, 0)$ and $\omega(\gamma, 1)$ on deep out-of-the-money call options become untenably large. Unfortunately, these options are largely unobservable. As such, we can only realistically measure expectations for a $\gamma \leq 3$ investor in practice.

Armed with the expected equity premium from the standpoint of a power utility investor, we can compute spot rates, forward rates, and forecast errors in the usual way.

Results. Table A5 presents results for option-based risk-premium and expected-return forecast errors when re-estimated under the assumption of an unconstrained investor with power utility, rather than log utility, for different values of constant relative risk aversion. Figure A7 provides a visual representation of these results. See Section 4.1.5 in the text for a discussion of the results.

C.3 Long-Horizon Forecast Errors

For the forecast-error quantification in Section 5.1, we re-estimate spot rates, forward rates, and forecast errors at longer horizons (up to $m + n = 8$ years) for the Euro Stoxx 50. The sample runs from September 2005 through September 2014 (beyond which we cannot yet observe realized forecast errors). The combinations of m and n (in months) can be seen in Table A6. For each such combination, we predict forecast errors as in (9) using a regression of realized forecast errors on shorter-horizon forward rates; we use the $n - 12 \times 12$ forward rate (with horizons again now in months) for $n \geq 24$, and for $n = 12$ we use the 6×6 rate. After obtaining these predicted forecast errors, we calculate a decay parameter for each date's forecast errors, $\phi_t^{(n,m)}$, as the ratio of estimated $\mathbb{E}_t[\varepsilon_{t+n}^{(m)}]$ to $\mathbb{E}_t[\varepsilon_{t+12}^{(12)}]$ for each available $m, n > 12$. This decay specification builds on the one used by De la O and Myers (2021, eq. (13)) for expected returns, but we estimate it directly for each date t (whereas they use a full-sample regression for one horizon). The entries of Table A6 report the median decay parameter over all t for each combination of m and n. In all cases the estimates are close to or above 1. Assuming that predictable forecast errors are permanent at all horizons might be thought of as providing an estimate of their maximal possible effect. That said, when we estimate the decay parameter in the U.S. (at shorter horizons, unreported), we in fact generally obtain estimates greater than 1, suggesting that setting $\phi^{(n,m)} = 1$ may, if anything, be slightly conservative in the U.S. sample.

C.4 Monetary Policy Shocks

Methodology. While a full accounting of the drivers of forecast errors is beyond the scope of the paper, here we consider one leading candidate: unexpected monetary policy shocks. We separately consider forecast errors for the risk-free rate and for the option-based risk premium (which together add up to the expected-return forecast errors). For the risk-free rate, the forecast error is the difference between the forecast error on the option-based expected return and the option-based risk premium. For each of the two sets of realized forecast errors, we then run time-series regressions — separately for the U.S. and the Eurozone — of forecast errors on contemporaneous monetary policy shocks:

$$\tilde{\varepsilon}_{t+6}^{(6)} = \beta_0 + \beta_1 \left(\sum_{h=1}^{6} \mathrm{MPS}_{t+h} \right) + e_{t+6},$$

where MPS_{t+h} is a monetary policy shock measure for month $t + h$. The summed measure is contemporaneous to the 6-month forecast error, and the spot-rate horizon is 6 months. For MPS_{t+h} in the U.S., we consider fed funds rate changes and surprises from Bernanke and Kuttner (2005, BK), obtained from Ken Kuttner's website; target shocks and path shocks from Gürkaynak, Sack, and Swanson (2005, GSS), obtained via Acosta (2023) from Miguel Acosta's website; and fed funds rate shocks and policy news shocks from Nakamura and Steinsson (2018, NS), also obtained from Miguel Acosta's website. For the Eurozone, we consider target shocks, forward guidance shocks, and risk premium shocks from Leombroni et al. (2021, LVVW), obtained from Andrea Vedolin's website. In all cases, we report the R^2 from the above regression, as well as the correlation between the two variables.

Results. Table A7 presents results. We find strong, positive correlations in most cases between the risk-free rate forecast error and the monetary policy shocks, particularly the shocks related to *future* policy expectations like the GSS path shock or the NS policy news shock (as should be expected given that these are forecast errors for expected future short-term rates). We find more limited explanatory power (R^2) for equity risk-premium forecast errors, indicating that these are not largely

driven by monetary policy shocks. That said, the shocks still explain some of the forecast errors, and there are meaningful negative correlations: monetary policy tightening that was unexpected ex ante tends to correlate with future risk-premium spot rates being lower than predicted according to the forward rates, consistent with improving economic conditions. This is also potentially consistent with the monetary policy shocks partially driving the decreases in spot rates, along the lines of the information effect discussed by Nakamura and Steinsson (2018) (see also Golez and Matthies 2023).

C.5 Additional Robustness Checks

Alternative Samples. Table A8 considers how our results extend to the full sample with all 20 available exchanges, as opposed to the 10 developed-market exchanges used in our main panel. The table presents results from a concise set of key regressions from Tables 2–5 in the main text. The main results hold here. The Mincer-Zarnowitz regression coefficients are similar to (in fact slightly below) the results in the main sample; the average forecast error is nearly identical to that in the main sample; and the error-predictability and Coibion-Gorodnichenko results are slightly stronger in the extended sample than in the main sample.

Alternative Horizons. Table A9 examines alternative horizons for Mincer-Zarnowitz, average error, and error-predictability regressions. The baseline analysis in Section 4 considers the 6-month spot rate in 6 months ($n = m = 6$ in Panel E). This exercise illustrates the effect of the horizon on the regression estimates. Holding $n + m$ fixed, forecast errors are relatively small on average and less predictable with small n; forecast errors are relatively large on average and more predictable with large n.

Additional Robustness Checks. Table A10 examines additional robustness checks. Panel A reports the baseline results. Panel B winsorizes spot rates, forward rates, forecast errors, and forecast revisions at the 2.5% level by exchange. Panel C is the trimming analogue to Panel B. Panel D repeats the analysis in balanced panels. Panel E repeats the analysis in subsamples. This exercise is generally consistent with the baseline results, although the Coibion-Gorodnichenko regression slopes are sensitive to winsorization/trimming and the forecast errors are somewhat less predictable in the later subsample.

D Additional Model Discussion: A Trilemma for Expectation Errors

This appendix continues the discussion in Section 6.3 on how different moments of the data are tied together by the cyclicality of forecast errors. We begin with the Campbell-Shiller price-dividend decomposition in (11). Assume that the expectations $\mathbb{E}_t[\cdot]$ in that decomposition refer to agents' subjective beliefs, and $p_t - d_t$ is the observed log price-dividend ratio. Now consider an alternative economy in which all agents have rational expectations. For arbitrary equilibrium variable x_t in the observed data, denote the corresponding variable in the alternative RE economy by x_t^{RE}. Define the wedge between these two variables to be $\widetilde{x}_t = x_t - x_t^{RE}$. For example, $\widetilde{p_t - d_t}$ is the wedge between the observed price-dividend ratio and the one that would be observed in the alternative economy with RE. Up to a constant, it satisfies

$$\widetilde{p_t - d_t} = \widetilde{CF}_t - \widetilde{\mathcal{F}}_t - \widetilde{RF}_t. \tag{A16}$$

Assume for simplicity that $\widetilde{RF}_t = 0$. The following variance decomposition for the price-dividend wedge therefore holds:

$$\text{var}\left(\widetilde{p_t - d_t}\right) = \text{var}\left(\widetilde{CF}_t\right) + \text{var}\left(\widetilde{\mathcal{F}}_t\right) - 2\,\text{cov}\left(\widetilde{CF}_t, \widetilde{\mathcal{F}}_t\right). \tag{A17}$$

Alternatively, one can also use the following decomposition given (A16):

$$\text{var}\left(\widetilde{p_t - d_t}\right) = \text{cov}\left(\widetilde{p_t - d_t}, \widetilde{CF}_t\right) - \text{cov}\left(\widetilde{p_t - d_t}, \widetilde{\mathcal{F}}_t\right). \tag{A18}$$

The wedges $\widetilde{CF}_t$ and $\widetilde{\mathcal{F}}_t$ can be understood as expectation errors along the lines considered in Section 6: if subjective expectations are too high relative to RE, then the wedge will be positive (and forecast errors, defined as realized − forecast, are likely to be negative). According to either of the decompositions in (A17)–(A18), therefore, one must choose from at most two of the following three features of any model of expectation errors:

1. Volatile expectation errors for returns (and/or fundamentals)
2. Volatile price-dividend ratio relative to a rational benchmark
3. Countercyclical return expectation errors (positive return expectation errors in bad times)

For example, if excessively positive cash-flow and return forecast revisions occur in good times (after positive news), then $\text{cov}\left(\widetilde{CF}_t, \widetilde{\mathcal{F}}_t\right) > 0$ in (A17). Alternatively, in the version expressed in (A18), positive comovement between price-dividend and forward-rate wedges similarly detracts from a model's ability to generate volatile $\widetilde{p_t - d_t}$. This form of overreaction to *realized outcomes* (cash flows and/or returns) may be intuitively appealing, but it limits a model's ability to speak to variation in the price-dividend ratio through expectation errors alone.[1]

Our empirical results, and our model of expectation errors, instead suggest overreaction of forward rates to *spot rates*, rather than realized returns. Unlike realized returns, we find that spot and forward rates *increase* in bad times. The negative covariance between fundamental news and return expectation errors in principle allows for a volatile price-dividend ratio.

---

[1] For example, Nagel and Xu (2022) obtain a price-dividend ratio volatility about 50% lower than that observed in the data (see their Table 5). Similarly, De la O and Myers (2021) report that in the model of Barberis et al. (2015), "movements in dividend change expectations are almost completely negated by movements in price change expectations. This leads to low variation in the price-dividend difference" (p. 1370); Campbell (2017) provides a related discussion of the Barberis et al. (2015) results.

Appendix Tables and Figures

Table A1
Option Sample

This table reports the region, the abbreviation, the underlying index, the sample period, and the sample length in months for each exchange. The last column indicates whether the exchange is in the main sample. See Appendix B.1 for more details.

Region	Abbrv	Index	Start	End	Length	Main
North America						
United States	USA	S&P 500	199001	202112	384	Y
Europe						
Belgium	BEL	BEL 20	200201	202109	228	
Switzerland	CHE	SMI	200201	202109	237	Y
Germany	DEU	DAX	200201	202109	237	Y
Spain	ESP	IBEX 35	200610	202109	180	Y
Finland	FIN	OMXH25	200201	202109	237	
France	FRA	CAC 40	200304	202109	222	Y
United Kingdom	GBR	FTSE 100	200201	202109	237	Y
Italy	ITA	FTSE MIB	200610	202109	180	Y
Netherlands	NLD	AEX	200201	202109	219	
Sweden	SWE	OMXS30	200705	202109	173	
Pan-Europe						
Euro Stoxx 50	SX5E	SX5E	200201	202109	237	Y
Stoxx Europe 50	SX5P	SX5P	200201	202109	237	
Stoxx Europe 600	SXXP	SXXP	200509	202109	193	
Asia-Pacific						
Australia	AUS	ASX 200	200401	202104	208	Y
China	CHN	HSCEI	200601	202104	184	
Hong Kong	HKG	Hang Seng	200601	202104	184	Y
Japan	JPN	Nikkei 225	200405	202104	204	
Korea	KOR	KOSPI 200	200407	202104	202	
Taiwan	TWN	TAIEX	200510	202104	187	

Table A2
Summary Statistics

This table reports summary statistics for option-based risk premia (left panel) and expected returns (right panel). The units are annualized percentage points. The horizon is the 6-month spot rate, 6 months from now. The sample is the longest available for each exchange in the full sample.

| | RISK PREMIA | | | | EXPECTED RETURNS | | | |
| | Spot $\tilde{\mu}_{t+6}^{(6)}$ | | Forward $\tilde{f}_t^{6,6}$ | | Spot $\tilde{\mu}_{t+6}^{(6)}$ | | Forward $f_t^{6,6}$ | |
	Mean	StDev	Mean	StDev	Mean	StDev	Mean	StDev
North America								
USA	2.17	1.40	2.15	0.99	5.04	2.41	5.44	2.32
Europe								
BEL	2.42	2.07	2.56	2.22	3.59	2.82	4.13	2.79
CHE	1.96	1.34	1.88	0.80	2.26	1.85	2.60	1.50
DEU	2.85	1.86	2.63	1.07	4.10	2.74	4.27	2.12
ESP	3.37	1.83	2.89	1.17	4.19	2.73	4.19	2.19
FIN	2.79	2.21	2.38	1.85	4.03	2.85	4.02	2.65
FRA	2.57	1.56	2.44	1.04	3.71	2.42	3.97	2.07
GBR	2.21	1.51	2.12	1.00	4.33	2.59	4.76	2.13
ITA	3.60	1.67	3.17	1.35	4.39	2.40	4.42	2.23
NLD	2.85	2.18	2.43	1.26	3.93	3.20	3.95	2.40
SWE	2.79	1.80	2.44	1.41	3.79	2.83	3.77	2.68
Pan-Europe								
SX5E	2.89	1.83	2.65	1.08	4.14	2.69	4.29	2.10
SX5P	2.23	1.62	2.11	1.48	3.47	2.64	3.75	2.51
SXXP	2.14	1.63	2.05	1.53	3.12	2.65	3.44	2.57
Asia-Pacific								
AUS	1.97	1.39	1.84	1.20	5.65	2.65	5.95	2.39
CHN	4.24	4.01	3.90	3.29	5.61	4.42	5.92	3.64
HKG	2.98	2.56	2.85	1.93	4.35	2.98	4.86	2.26
JPN	2.93	2.20	2.60	1.57	3.21	2.35	3.15	1.76
KOR	2.23	2.16	2.14	1.85	5.13	3.02	5.40	2.95
TWN	2.39	1.90	2.36	1.45	3.37	2.11	3.84	1.70

Table A3
Alternative Measures

This table reports regression estimates with alternative measures for option-based risk premia (Panel A) and expected returns (Panel B). See Appendix B.3 for more details. The first three subpanels report estimates from panel regressions in the main sample. The last subpanel reports estimates from time-series regressions in the U.S. sample. Mincer-Zarnowitz regressions test $H_0 : \beta_1 = 1$, as in Table 2. The average forecast error tests $H_0 : \bar{\varepsilon}_t = 0$, as in Table 4. Error-predictability regressions test $H_0 : \beta_1 = 0$, as in Table 5. The horizon is the 6-month spot rate, 6 months from now. The units are annualized percentage points. Panel regressions, in the main sample, include exchange fixed effects, compute a within R^2, and 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 lags selected following Lazarus et al. (2018) and fixed-b p-values. This sample is from January 1990 to June 2021.

| | PANEL A. RISK PREMIA $\tilde{\mu}_{t+6}^{(6)}$ | | | | | | | | | | | |
| | Mincer-Zarnowitz | | | | Average Error | | | Error Predictability | | | | |
	β_1	$se(\beta_1)$	p-val	R^2	$\tilde{\varepsilon}_t$	$se(\tilde{\varepsilon}_t)$	p-val	β_1	$se(\beta_1)$	p-val	R^2	N
Alternative Integration Bounds												
Truncation: $0.65 \le K/P \le 1.35$	0.78	0.091	**	0.18	0.34	0.094	***	-0.019	0.039		0.00	2139
Truncation: $0.55 \le K/P \le 1.45$	0.68	0.074	***	0.17	0.26	0.097	**	-0.082	0.041	*	0.01	2137
Truncation: $0.45 \le K/P \le 1.55$	0.63	0.065	***	0.16	0.22	0.099	*	-0.12	0.042	**	0.02	2140
Truncation: $0.35 \le K/P \le 1.65$	0.59	0.059	***	0.15	0.20	0.10	*	-0.15	0.044	***	0.03	2138
Observable Moneyness	0.56	0.053	***	0.15	0.19	0.10	*	-0.17	0.043	***	0.04	2140
Extrapolation: Static Moneyness	0.56	0.056	***	0.15	0.14	0.10		-0.15	0.041	***	0.04	2244
Extrapolation: Dynamic Moneyness	0.56	0.055	***	0.15	0.16	0.11		-0.17	0.046	***	0.03	2241
Extrapolation: Baseline	0.56	0.055	***	0.15	0.17	0.11		-0.16	0.047	***	0.03	2227
Alternative Liquidity Filters												
Outliers	0.60	0.055	***	0.17	0.18	0.098		-0.20	0.050	***	0.05	2241
Open Interest: After 199601	0.52	0.057	***	0.14	0.16	0.11		-0.18	0.045	***	0.04	2033
Outliers and Open Interest: After 199601	0.56	0.051	***	0.16	0.19	0.097	*	-0.21	0.050	***	0.06	2040
Alternative Surfaces												
Volatility Surface: After 199601	0.57	0.051	***	0.17	0.087	0.095		-0.21	0.053	***	0.06	2163
SVI Surface: U.S. and SX5E	0.59	0.047	*	0.17	0.051	0.11		-0.19	0.039		0.04	609
Alternative Prices												
Bid-Ask Midpoint: U.S. Only	0.67	0.096	***	0.22	0.021	0.15		-0.17	0.066	**	0.03	378
Bid Prices: U.S. Only	0.64	0.099	***	0.21	0.034	0.14		-0.15	0.077	**	0.03	378
Ask Prices: U.S. Only	0.66	0.089	***	0.23	0.0051	0.15		-0.18	0.060	***	0.04	378

(Continued on the next page)

Table A3
Alternative Measures (Continued)

| | Panel B. Expected Returns $\mu_{t+6}^{(6)}$ | | | | | | | | | | | |
| | Mincer-Zarnowitz | | | | Average Error | | | Error Predictability | | | | |
	β_1	$se(\beta_1)$	p-val	R^2	ε_t	$se(\varepsilon_t)$	p-val	β_1	$se(\beta_1)$	p-val	R^2	N
Alternative Integration Bounds												
Truncation: $0.65 \leq K/P \leq 1.35$	1.02	0.053		0.71	-0.12	0.094		-0.18	0.043	***	0.04	2139
Truncation: $0.55 \leq K/P \leq 1.45$	0.99	0.054		0.67	-0.19	0.097	*	-0.23	0.044	***	0.07	2137
Truncation: $0.45 \leq K/P \leq 1.55$	0.96	0.055		0.64	-0.23	0.099	**	-0.26	0.045	***	0.09	2140
Truncation: $0.35 \leq K/P \leq 1.65$	0.94	0.056		0.63	-0.25	0.10	**	-0.29	0.046	***	0.10	2138
Observable Moneyness	0.93	0.058		0.61	-0.26	0.10	**	-0.31	0.046	***	0.12	2140
Extrapolation: Static Moneyness	0.92	0.060		0.60	-0.31	0.100	**	-0.26	0.039	***	0.12	2244
Extrapolation: Dynamic Moneyness	0.91	0.061		0.59	-0.29	0.10	**	-0.30	0.047	***	0.11	2241
Extrapolation: Baseline	0.91	0.061		0.59	-0.28	0.10	**	-0.29	0.047	***	0.10	2227
Alternative Liquidity Filters												
Outliers	0.92	0.062		0.62	-0.27	0.098	**	-0.34	0.050	***	0.15	2241
Open Interest: After 199601	0.91	0.062		0.59	-0.26	0.10	**	-0.31	0.043	***	0.12	2033
Outliers and Open Interest: After 199601	0.92	0.064		0.62	-0.24	0.096	**	-0.35	0.048	***	0.16	2040
Alternative Surfaces												
Volatility Surface: After 199601	0.93	0.065		0.63	-0.34	0.098	***	-0.36	0.052	***	0.16	2163
SVI Surface: U.S. and SX5E	0.89	0.048		0.65	-0.36	0.11		-0.28	0.049		0.10	609
Alternative Prices												
Bid-Ask Midpoint: U.S. Only	0.88	0.072		0.71	-0.40	0.17	**	-0.20	0.090	**	0.04	378
Bid Prices: U.S. Only	0.87	0.070	*	0.73	-0.38	0.17	**	-0.18	0.11	*	0.04	378
Ask Prices: U.S. Only	0.87	0.076		0.69	-0.41	0.17	**	-0.21	0.077	***	0.05	378

(Continued from the previous page)

Table A4
Coibion-Gorodnichenko Regressions

This table reports Coibion-Gorodnichenko regressions of future realized forecast errors on current forecast revisions for option-based risk premia (left panel) and expected returns (right panel):

$$\tilde{\varepsilon}_{i,t+n}^{m} = \beta_0 + \beta_1 \left(\tilde{f}_{i,t}^{(n,m)} - \tilde{f}_{i,t-h}^{(n+h,m)} \right) + e_{i,t+n}$$

For risk premia, the realized spot rate is the future expectation of the m-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. The units are annualized percentage points. All regressions include exchange fixed effects and compute a within R^2. Standard errors are clustered by exchange and date. The sample is the longest available for each exchange in the main sample.

				RISK PREMIA $\tilde{\mu}_{t+n}^{(m)}$				EXPECTED RETURNS $\mu_{t+n}^{(m)}$				
				β_1	$se(\beta_1)$	p-val	R^2	β_1	$se(\beta_1)$	p-val	R^2	N
Short-Horizon Forecast Revisions												
$h = $ 1-Month Revision	$n+m=4$	$n=3$	$m=1$	-0.20	0.11	*	0.01	-0.23	0.13	*	0.01	2214
$h = $ 1-Month Revision	$n+m=5$	$n=4$	$m=1$	-0.27	0.11	**	0.01	-0.27	0.13	*	0.01	2217
$h = $ 2-Month Revision	$n+m=4$	$n=3$	$m=1$	-0.22	0.081	**	0.01	-0.23	0.11	*	0.01	2205
$h = $ 3-Month Revision	$n+m=6$	$n=3$	$m=3$	-0.055	0.056		0.00	-0.067	0.094		0.00	2215
$h = $ 3-Month Revision	$n+m=9$	$n=6$	$m=3$	-0.24	0.12	*	0.01	0.033	0.11		0.00	2186
Long-Horizon Forecast Revisions												
$h = $ 4-Month Revision	$n+m=5$	$n=2$	$m=3$	-0.081	0.053		0.00	-0.086	0.079		0.01	2223
$h = $ 5-Month Revision	$n+m=4$	$n=1$	$m=3$	-0.051	0.053		0.00	-0.050	0.064		0.00	2213
$h = $ 6-Month Revision	$n+m=6$	$n=3$	$m=3$	-0.13	0.055	**	0.01	-0.031	0.063		0.00	2183
$h = $ 7-Month Revision	$n+m=5$	$n=2$	$m=3$	-0.12	0.040	**	0.02	-0.065	0.053		0.00	2191
$h = $ 8-Month Revision	$n+m=4$	$n=1$	$m=3$	-0.082	0.034	**	0.01	-0.056	0.044		0.01	2181

Table A5
Power Utility Regressions

This table reports regression estimates from the standpoint of an unconstrained power utility investor fully invested in the market for option-based risk premia (Panel A) and expected returns (Panel B). See Section 4.1.5 and Appendix C.2 for more details. Mincer-Zarnowitz regressions test $H_0 : \beta_1 = 1$, as in Table 2. The average forecast error tests $H_0 : \bar{\varepsilon}_t = 0$, as in Table 4. Error-predictability regressions test $H_0 : \beta_1 = 0$, as in Table 5. The horizon is the 6-month spot rate, 6 months from now. The units are annualized percentage points. Panel regressions, in the main sample, include exchange fixed effects, compute a within R^2, and 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 lags selected following Lazarus et al. (2018) and fixed-b p-values. This sample is from January 1990 to June 2021.

	Mincer-Zarnowitz				Average Error			Error Predictability				
	β_1	$se(\beta_1)$	p-val	R^2	$\tilde{\varepsilon}_t$	$se(\tilde{\varepsilon}_t)$	p-val	β_1	$se(\beta_1)$	p-val	R^2	N
PANEL A. RISK PREMIA $\tilde{\mu}_{t+6}^{(6)}$												
Main Sample												
$\gamma = 0.75$	0.51	0.059	***	0.13	0.045	0.057		-0.19	0.045	***	0.04	2224
$\gamma = 1.00$	0.56	0.055	***	0.15	0.17	0.11		-0.16	0.046	***	0.03	2220
$\gamma = 1.25$	0.60	0.057	***	0.15	0.36	0.15	**	-0.13	0.050	**	0.02	2215
$\gamma = 1.50$	0.62	0.062	***	0.16	0.59	0.20	**	-0.10	0.054	*	0.01	2206
$\gamma = 2.00$	0.64	0.075	***	0.15	1.14	0.29	***	-0.044	0.057		0.00	2181
U.S. Sample												
$\gamma = 0.75$	0.61	0.094	***	0.21	-0.022	0.083		-0.20	0.069	***	0.05	378
$\gamma = 1.00$	0.67	0.096	***	0.22	0.021	0.15		-0.17	0.066	**	0.03	378
$\gamma = 1.25$	0.72	0.098	**	0.23	0.12	0.20		-0.13	0.064	**	0.02	378
$\gamma = 1.50$	0.76	0.100	**	0.24	0.25	0.26		-0.099	0.062		0.01	378
$\gamma = 2.00$	0.85	0.10		0.26	0.57	0.35		-0.048	0.059		0.00	378
$\gamma = 2.50$	0.91	0.10		0.27	0.94	0.44	**	-0.0087	0.057		-0.00	378
$\gamma = 3.00$	0.97	0.10		0.28	1.32	0.52	**	0.022	0.056		-0.00	378
PANEL B. EXPECTED RETURNS $\mu_{t+6}^{(6)}$												
	β_1	$se(\beta_1)$	p-val	R^2	ε_t	$se(\varepsilon_t)$	p-val	β_1	$se(\beta_1)$	p-val	R^2	N
Main Sample												
$\gamma = 0.75$	0.93	0.042		0.75	-0.40	0.061	***	-0.42	0.056	***	0.17	2224
$\gamma = 1.00$	0.91	0.061		0.59	-0.28	0.10	**	-0.30	0.046	***	0.11	2220
$\gamma = 1.25$	0.90	0.075		0.50	-0.084	0.15		-0.22	0.048	***	0.07	2215
$\gamma = 1.50$	0.90	0.088		0.43	0.15	0.19		-0.17	0.052	***	0.04	2206
$\gamma = 2.00$	0.88	0.11		0.36	0.69	0.28	**	-0.093	0.057		0.01	2181
U.S. Sample												
$\gamma = 0.75$	0.89	0.061	*	0.84	-0.44	0.14	***	-0.24	0.13	*	0.04	378
$\gamma = 1.00$	0.88	0.072		0.71	-0.40	0.17	**	-0.20	0.090	**	0.04	378
$\gamma = 1.25$	0.87	0.082		0.61	-0.30	0.21		-0.15	0.077	**	0.03	378
$\gamma = 1.50$	0.87	0.089		0.53	-0.17	0.25		-0.12	0.070	*	0.02	378
$\gamma = 2.00$	0.90	0.096		0.45	0.15	0.34		-0.066	0.063		0.00	378
$\gamma = 2.50$	0.93	0.098		0.41	0.52	0.42		-0.024	0.059		-0.00	378
$\gamma = 3.00$	0.96	0.098		0.39	0.90	0.50	*	0.0077	0.056		-0.00	378

Table A6
Long-Horizon Forecast Errors

This table reports estimates of the predicted forecast error decay for the Euro Stoxx 50. The decay $\phi_t^{(n,m)}$ parameter follows the expected decay specification from De la O and Myers (2021):

$$\mathbb{E}_t\left[\varepsilon_{t+n}^{(m)}\right] = \phi_t^{(n,m)}\,\mathbb{E}_t\left[\varepsilon_{t+12}^{(12)}\right]$$

The estimate is the median by horizon:

$$\phi^{(n,m)} = median\left\{\left|\phi_t^{(n,m)}\right|\right\}$$

The predicted forecast error is from a time-series regression of future realized forecast errors on current forward rates: the predictor is the $n - 12 \times 12$-month forward rate for $n \geq 24$ and the 6×6-month forward rate for $n = 12$. The sample is from 09/2005 to 09/2014.

	n-months						
m-months	12	24	36	48	60	72	84
12	1.00	2.13	1.88	0.73	0.89	1.23	1.05
24	1.00	2.13	1.94	0.90	0.88	1.37	
36	1.00	2.09	1.98	0.95	0.98		
48	1.00	2.13	1.99	1.08			
60	1.00	2.09	1.91				
72	1.00	2.07					
84	1.00						

Table A7
Forecast Errors and Monetary Policy Shocks

The table reports the comovement between forecast errors and monetary policy shocks in the U.S. (Panel A) and the Eurozone (Panel B). See Appendix C.4 for more details. For option-based risk premia, the R^2 is from a time-series regression of contemporaneous forecast errors on monetary policy shocks

$$\tilde{\varepsilon}_{t+6}^{(6)} = \beta_0 + \beta_1 \left(\sum_{h=1}^{6} \mathrm{MPS}_{t+h} \right) + e_{t+6}$$

and the correlation is between the same two variables. These expectations are analogously defined for the risk-free rate. The horizon is the 6-month spot rate, 6 months from now. The units are annualized percentage points. U.S. forecast errors are for the S&P 500, Eurozone the Euro Stoxx 50. BK monetary policy shocks are obtained from Ken Kuttner's website, GSS and NS are obtained via Acosta (2023) from Miguel Acosta's website, and LVVW are obtained from Andrea Vedolin's website.

				PANEL A. U.S. MONETARY POLICY SHOCKS			
				Correlation ρ		Explained Variation R^2	
	Start	End	Length	Risk-Free	Risk Premia	Risk-Free	Risk Premia
Bernanke-Kuttner: BK							
Fed Funds Rate Change	199402	201812	299	0.64	-0.31	0.40	0.09
Fed Funds Rate Surprise	199402	201812	299	0.12	-0.26	0.01	0.07
Gurkaynak-Sack-Swanson: GSS							
Target Shock	199905	202106	266	0.21	-0.18	0.05	0.03
Path Shock	199905	202106	266	0.51	-0.25	0.26	0.06
Nakamura-Steinsson: NS							
Fed Funds Rate Shock	199905	202106	266	0.17	-0.18	0.03	0.03
Policy News Shock	199905	202106	266	0.56	-0.32	0.32	0.10
				PANEL B. EUROZONE MONETARY POLICY SHOCKS			
				Correlation ρ		Explained Variation R^2	
	Start	End	Length	Risk-Free	Risk Premia	Risk-Free	Risk Premia
Leombroni-Vedolin-Venter-Whelan: LVVW							
Target Shock	200201	202006	222	-0.30	0.15	0.09	0.02
Forward Guidance Shock	200201	202006	222	0.23	-0.36	0.05	0.13
Risk Premium Shock	200201	202006	222	0.53	-0.28	0.28	0.08

Table A8
Alternative Samples

This table reports regression estimates in alternative samples for option-based risk premia (Panel A) and expected returns (Panel B). See Appendix C.5 and Table A1 for more details. Mincer-Zarnowitz regressions test $H_0 : \beta_1 = 1$, as in Table 2. The average forecast error tests $H_0 : \bar{\varepsilon}_t = 0$, as in Table 4. Error-predictability regressions test $H_0 : \beta_1 = 0$, as in Table 5. The horizon is the 6-month spot rate, 6 months from now. The units are annualized percentage points. All regressions include exchange fixed effects and compute a within R^2. Standard errors are clustered by exchange and date. The sample is the longest available for each exchange.

| | PANEL A. RISK PREMIA $\tilde{\mu}_{t+6}^{(6)}$ | | | | | | | | | | | |
| | Mincer-Zarnowitz | | | | Average Error | | | Error Predictability | | | | |
	β_1	$se(\beta_1)$	p-val	R^2	$\tilde{\varepsilon}_t$	$se(\tilde{\varepsilon}_t)$	p-val	β_1	$se(\beta_1)$	p-val	R^2	N
Baseline												
Main: 10 Indexes	0.56	0.055	***	0.15	0.17	0.11		-0.16	0.047	***	0.03	2227
Full: 20 Indexes	0.48	0.053	***	0.14	0.19	0.12		-0.26	0.049	***	0.08	4199
Eurozone												
Main: 5 Indexes	0.53	0.080	***	0.12	0.28	0.13	*	-0.15	0.053	*	0.03	1019
Full: 8 Indexes	0.45	0.10	***	0.11	0.26	0.13	*	-0.22	0.076	**	0.05	1665
Europe												
Main: 7 Indexes	0.55	0.071	***	0.13	0.22	0.12		-0.15	0.049	**	0.03	1474
Full: 13 Indexes	0.46	0.075	***	0.13	0.21	0.12	*	-0.22	0.058	***	0.06	2696
Asia-Pacific												
Main: 2 Indexes	0.54	0.093		0.18	0.13	0.11		-0.20	0.12		0.04	375
Full: 6 Indexes	0.48	0.057	***	0.14	0.18	0.16		-0.33	0.064	***	0.12	1125
Excludes U.S.												
Main: 9 Indexes	0.55	0.056	***	0.14	0.20	0.11		-0.16	0.049	**	0.03	1849
Full: 19 Indexes	0.47	0.053	***	0.13	0.20	0.12		-0.27	0.051	***	0.08	3821

| | PANEL B. EXPECTED RETURNS $\mu_{t+6}^{(6)}$ | | | | | | | | | | | |
| | Mincer-Zarnowitz | | | | Average Error | | | Error Predictability | | | | |
	β_1	$se(\beta_1)$	p-val	R^2	ε_t	$se(\varepsilon_t)$	p-val	β_1	$se(\beta_1)$	p-val	R^2	N
Baseline												
Main: 10 Indexes	0.91	0.061		0.59	-0.28	0.10	**	-0.29	0.047	***	0.10	2227
Full: 20 Indexes	0.79	0.089	**	0.47	-0.25	0.12	*	-0.37	0.058	***	0.15	4199
Eurozone												
Main: 5 Indexes	0.92	0.081		0.57	-0.13	0.12		-0.30	0.056	***	0.11	1019
Full: 8 Indexes	0.83	0.11		0.49	-0.15	0.13		-0.35	0.080	***	0.14	1665
Europe												
Main: 7 Indexes	0.94	0.077		0.59	-0.21	0.11		-0.30	0.052	***	0.12	1474
Full: 13 Indexes	0.84	0.089	*	0.53	-0.20	0.12		-0.35	0.064	***	0.15	2696
Asia-Pacific												
Main: 2 Indexes	0.86	0.092		0.51	-0.40	0.14		-0.31	0.10		0.10	375
Full: 6 Indexes	0.67	0.12	**	0.32	-0.29	0.17		-0.42	0.072	***	0.17	1125
Excludes U.S.												
Main: 9 Indexes	0.92	0.074		0.57	-0.25	0.11	*	-0.30	0.048	***	0.11	1849
Full: 19 Indexes	0.78	0.096	**	0.45	-0.23	0.12	*	-0.38	0.059	***	0.16	3821

Table A9
Alternative Horizons

This table reports regression estimates at alternative horizons for option-based risk premia (Panel A) and expected returns (Panel B). See Appendix C.5 for more details. Mincer-Zarnowitz regressions test $H_0 : \beta_1 = 1$, as in Table 2. The average forecast error tests $H_0 : \bar{\varepsilon}_t = 0$, as in Table 4. Error-predictability regressions test $H_0 : \beta_1 = 0$, as in Table 5. The horizon is the m-month spot rate, n months from now. The units are annualized percentage points. All regressions include exchange fixed effects and compute a within R^2. Standard errors are clustered by exchange and date. The sample is the longest available for each exchange in the main sample.

		PANEL A. RISK PREMIA $\tilde{\mu}_n^m$											
		Mincer-Zarnowitz				Average Error			Error Predictability				
		β_1	$se(\beta_1)$	p-val	R^2	$\tilde{\varepsilon}_t$	$se(\tilde{\varepsilon}_t)$	p-val	β_1	$se(\beta_1)$	p-val	R^2	N
4-Month Horizon: $n+m=4$													
$n=1$	$m=3$	0.90	0.048	*	0.62	0.044	0.073		-0.12	0.040	**	0.03	2269
$n=2$	$m=2$	0.77	0.058	***	0.35	0.055	0.11		-0.20	0.053	***	0.04	2268
$n=3$	$m=1$	0.68	0.065	***	0.20	0.094	0.14		-0.24	0.055	***	0.04	2235
5-Month Horizon: $n+m=5$													
$n=1$	$m=4$	0.94	0.043		0.67	0.067	0.064		-0.075	0.033	*	0.02	2269
$n=2$	$m=3$	0.83	0.053	**	0.42	0.12	0.10		-0.13	0.043	**	0.02	2268
$n=3$	$m=2$	0.74	0.063	***	0.25	0.16	0.13		-0.15	0.049	**	0.02	2241
$n=4$	$m=1$	0.62	0.078	***	0.13	0.27	0.15		-0.19	0.053	***	0.02	2236
6-Month Horizon: $n+m=6$													
$n=1$	$m=5$	0.95	0.037		0.70	0.067	0.058		-0.054	0.029	*	0.01	2269
$n=2$	$m=4$	0.86	0.047	**	0.46	0.14	0.091		-0.094	0.036	**	0.01	2268
$n=3$	$m=3$	0.78	0.058	***	0.30	0.21	0.12		-0.11	0.044	**	0.01	2249
$n=4$	$m=2$	0.67	0.072	***	0.17	0.27	0.14	*	-0.15	0.050	**	0.02	2238
$n=5$	$m=1$	0.57	0.076	***	0.09	0.32	0.17	*	-0.20	0.057	***	0.02	2228
9-Month Horizon: $n+m=9$													
$n=3$	$m=6$	0.81	0.055	***	0.39	0.15	0.087		-0.065	0.036		0.01	2257
$n=4$	$m=5$	0.73	0.066	***	0.26	0.21	0.10	*	-0.10	0.044	**	0.01	2242
$n=5$	$m=4$	0.65	0.071	***	0.17	0.26	0.12	*	-0.14	0.049	**	0.02	2232
$n=6$	$m=3$	0.55	0.071	***	0.10	0.28	0.14	*	-0.17	0.051	***	0.02	2224
12-Month Horizon: $n+m=12$													
$n=3$	$m=9$	0.81	0.049	***	0.44	0.094	0.070		-0.059	0.036		0.01	2258
$n=6$	$m=6$	0.56	0.055	***	0.15	0.17	0.11		-0.16	0.047	***	0.03	2227
$n=9$	$m=3$	0.39	0.097	***	0.06	0.25	0.15		-0.24	0.052	***	0.04	2191

(Continued on the next page)

Table A9
Alternative Horizons (Continued)

		Mincer-Zarnowitz				Average Error			Error Predictability				
		β_1	$se(\beta_1)$	p-val	R^2	ε_t	$se(\varepsilon_t)$	p-val	β_1	$se(\beta_1)$	p-val	R^2	N
4-Month Horizon: $n+m=4$													
$n=1$	$m=3$	0.96	0.042		0.80	-0.074	0.073		-0.17	0.039	***	0.06	2269
$n=2$	$m=2$	0.90	0.057		0.62	-0.21	0.11	*	-0.30	0.056	***	0.09	2268
$n=3$	$m=1$	0.89	0.068		0.50	-0.26	0.14	*	-0.36	0.059	***	0.08	2235
5-Month Horizon: $n+m=5$													
$n=1$	$m=4$	0.99	0.038		0.84	-0.027	0.063		-0.12	0.033	***	0.04	2269
$n=2$	$m=3$	0.95	0.052		0.69	-0.080	0.097		-0.22	0.046	***	0.06	2268
$n=3$	$m=2$	0.94	0.063		0.57	-0.13	0.12		-0.27	0.053	***	0.06	2241
$n=4$	$m=1$	0.92	0.080		0.45	-0.11	0.14		-0.33	0.057	***	0.06	2236
6-Month Horizon: $n+m=6$													
$n=1$	$m=5$	0.99	0.035		0.86	-0.027	0.057		-0.10	0.029	***	0.03	2269
$n=2$	$m=4$	0.97	0.049		0.73	-0.039	0.086		-0.18	0.040	***	0.05	2268
$n=3$	$m=3$	0.97	0.057		0.63	-0.043	0.11		-0.22	0.048	***	0.05	2249
$n=4$	$m=2$	0.95	0.071		0.52	-0.082	0.13		-0.29	0.052	***	0.06	2238
$n=5$	$m=1$	0.93	0.082		0.42	-0.15	0.15		-0.36	0.059	***	0.07	2228
9-Month Horizon: $n+m=9$													
$n=3$	$m=6$	0.97	0.046		0.73	-0.075	0.082		-0.16	0.036	***	0.05	2257
$n=4$	$m=5$	0.97	0.055		0.66	-0.082	0.096		-0.22	0.040	***	0.06	2242
$n=5$	$m=4$	0.97	0.063		0.59	-0.10	0.11		-0.27	0.046	***	0.08	2232
$n=6$	$m=3$	0.94	0.071		0.51	-0.17	0.13		-0.31	0.053	***	0.08	2224
12-Month Horizon: $n+m=12$													
$n=3$	$m=9$	0.96	0.039		0.78	-0.14	0.067	*	-0.15	0.034	***	0.06	2258
$n=6$	$m=6$	0.91	0.061		0.59	-0.28	0.10	**	-0.29	0.047	***	0.10	2227
$n=9$	$m=3$	0.81	0.077	**	0.41	-0.46	0.15	**	-0.41	0.062	***	0.12	2191

PANEL B. EXPECTED RETURNS μ_n^m

(Continued from the previous page)

Table A10
Additional Robustness Checks

This table reports additional robustness checks for option-based risk premia (Panel A) and expected returns (Panel B). See Appendix C.5 for more details. Mincer-Zarnowitz regressions test $H_0 : \beta_1 = 1$, as in Table 2. The average forecast error tests $H_0 : \bar{\varepsilon}_t = 0$, as in Table 4. Error-predictability regressions test $H_0 : \beta_1 = 0$, as in Table 5. The horizon is the 6-month spot rate, 6 months from now. The units are annualized percentage points. All regressions include exchange fixed effects and compute a within R^2. Standard errors are clustered by exchange and date. The sample is the longest available for each exchange in the main sample.

| | PANEL A. RISK PREMIA $\tilde{\mu}_{t+6}^{(6)}$ | | | | | | | | | | | |
| | Mincer-Zarnowitz | | | | Average Error | | | Error Predictability | | | | |
	β_1	$se(\beta_1)$	p-val	R^2	$\tilde{\varepsilon}_t$	$se(\tilde{\varepsilon}_t)$	p-val	β_1	$se(\beta_1)$	p-val	R^2	N
Baseline: 199001 to 202106	0.56	0.055	***	0.15	0.17	0.11		-0.16	0.047	***	0.03	2227
Winsorization: 199001 to 202106	0.60	0.059	***	0.19	0.15	0.093		-0.15	0.048	**	0.03	2227
Trimming: 199001 to 202106	0.57	0.056	***	0.18	0.091	0.078		-0.10	0.049	*	0.01	2029
Balanced Panel: 200610 to 202010	0.49	0.058	***	0.12	0.22	0.13		-0.20	0.050	***	0.05	1674
First Half: 199001 to 201110	0.49	0.078	***	0.11	0.37	0.19	*	-0.25	0.064	***	0.07	1113
Second Half: 201111 to 202106	0.30	0.12	***	0.06	-0.024	0.099		-0.11	0.081		0.01	1114

| | PANEL B. EXPECTED RETURNS $\mu_{t+6}^{(6)}$ | | | | | | | | | | | |
| | Mincer-Zarnowitz | | | | Average Error | | | Error Predictability | | | | |
	β_1	$se(\beta_1)$	p-val	R^2	ε_t	$se(\varepsilon_t)$	p-val	β_1	$se(\beta_1)$	p-val	R^2	N
Baseline: 199001 to 202106	0.91	0.061		0.59	-0.28	0.10	**	-0.29	0.047	***	0.10	2227
Winsorization: 199001 to 202106	0.91	0.048	*	0.65	-0.31	0.092	***	-0.28	0.051	***	0.11	2227
Trimming: 199001 to 202106	0.90	0.045	*	0.65	-0.33	0.079	***	-0.20	0.051	***	0.06	2041
Balanced Panel: 200610 to 202010	0.89	0.080		0.55	-0.28	0.13	*	-0.35	0.048	***	0.14	1674
First Half: 199001 to 201110	0.78	0.11	*	0.30	-0.14	0.18		-0.40	0.061	***	0.18	1113
Second Half: 201111 to 202106	0.43	0.085	***	0.21	-0.41	0.11	***	-0.19	0.088	*	0.04	1114

Table A11
Model Calibration: Objective Parameters

This table reports AR(3) time-series regressions for 3-month spot rates:

$$\mu_t^{(3)} = \left(1 - \sum_{j=1}^{3} \phi_j\right) \bar{\mu} + \phi_1 \mu_{t-1}^{(3)} + \phi_2 \mu_{t-2}^{(3)} + \phi_3 \mu_{t-3}^{(3)} + e_t$$

The units are annualized percentage points. Standard errors are from 10,000 Monte Carlo simulations of length T months. The sample is the longest available for each exchange in the main sample.

	AUS	CHE	DEU	ESP	FRA	GBR	HKG	ITA	SX5E	USA
ϕ_1	0.59	0.78	0.79	0.76	0.87	0.90	0.74	0.66	0.77	0.94
	(0.07)	(0.07)	(0.07)	(0.08)	(0.07)	(0.07)	(0.07)	(0.08)	(0.07)	(0.05)
ϕ_2	0.02	-0.01	-0.06	-0.07	-0.14	-0.19	0.16	-0.02	-0.03	-0.26
	(0.08)	(0.08)	(0.08)	(0.09)	(0.09)	(0.08)	(0.09)	(0.09)	(0.08)	(0.07)
ϕ_3	0.22	0.05	0.11	0.09	0.12	0.14	-0.05	0.11	0.11	0.18
	(0.07)	(0.06)	(0.06)	(0.08)	(0.07)	(0.06)	(0.07)	(0.07)	(0.06)	(0.05)
$\bar{\mu}$	6.19	6.13	8.89	10.49	7.69	6.73	9.30	11.63	8.99	6.40
	(1.36)	(1.00)	(1.62)	(1.48)	(1.23)	(1.21)	(2.44)	(1.27)	(1.59)	(0.89)
σ_e	3.23	2.86	3.82	4.41	2.89	2.89	5.04	4.31	3.69	2.56
	(0.16)	(0.13)	(0.18)	(0.23)	(0.14)	(0.13)	(0.27)	(0.23)	(0.17)	(0.09)
T	208	237	237	180	222	237	184	180	237	384

Figure A1
Filtered Option Prices

Panel A plots the number of options after filters. Panel B plots the share of filtered options with positive open interest. Each bar is the annual median from daily data. The black line is the full sample median from daily data. The sample is the longest available for each exchange. Option prices have maturity $30 \leq \tau \leq 365$ days. See Appendix B.1 for more details.

PANEL A. NUMBER OF OBSERVATIONS
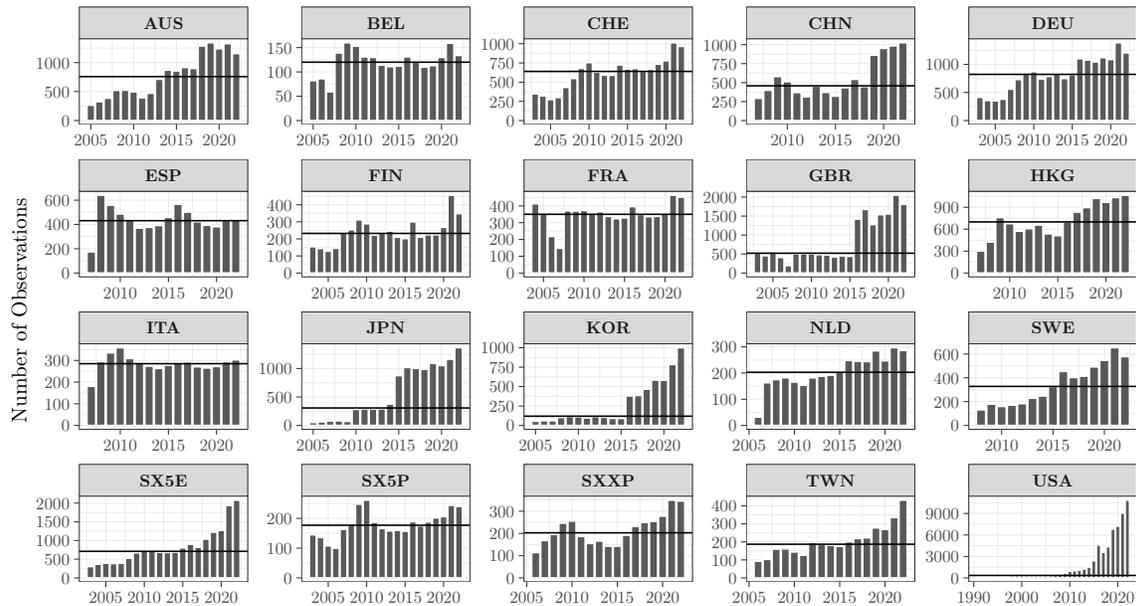
PANEL B. POSITIVE OPEN INTEREST
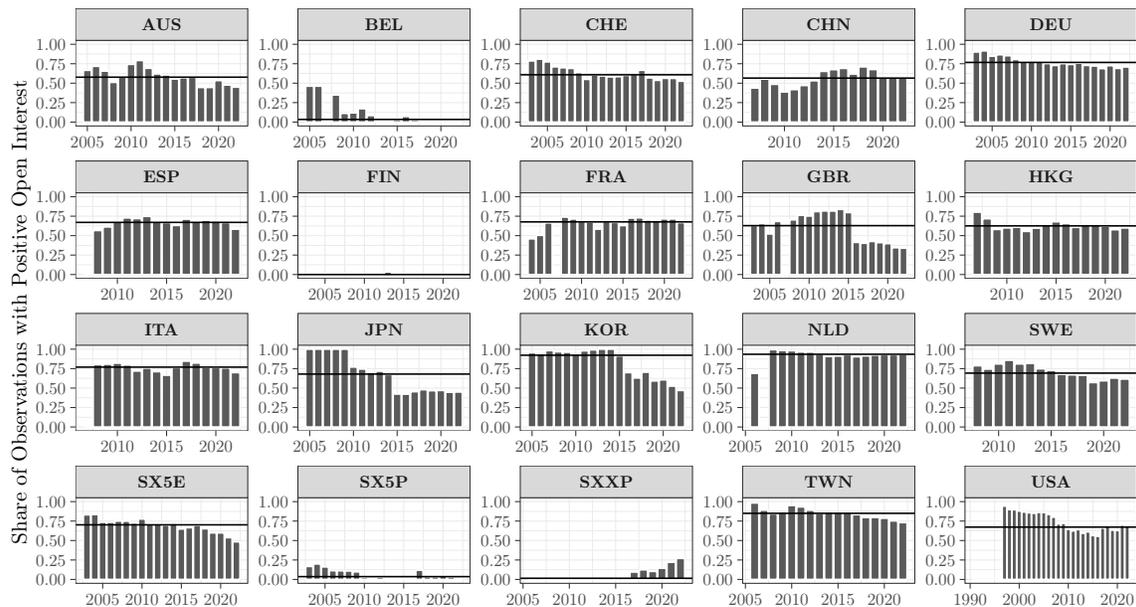
Figure A2
Minimum/Maximum Strike Price by Maturity

This figure plots the minimum/maximum strike price by maturity bin. The minimum/maximum is the annual median from daily data. The black line is the full sample minimum/maximum from daily data. The units are risk-neutral standard deviations from the index price. The sample is the longest available for each exchange. See Appendix B.1 for more details.

PANEL A. 92 TO 182 DAYS-TO-MATURITY

PANEL B. 274 TO 365 DAYS-TO-MATURITY

Figure A3
Timeline: Current/Realized Spot Rates and Forward Rates

See Section 3 for more details.

$$\mu_{t+n}^{(m)}$$

$$\mu_t^{(n+m)}$$

$$\mu_t^{(n)}$$

$$f_t^{(n,m)}$$

t $\qquad\qquad\qquad\qquad$ $t+n$ $\qquad\qquad\qquad\qquad$ $t+n+m$

- $\mu_t^{(n)} = \mathcal{L}_t^{(n)}$ is the n-month spot rate at time $= t$

- $\mu_t^{(n+m)} = \mathcal{L}_t^{(n+m)}$ is the $n+m$-month spot rate at time $= t$

- $f_t^{(n,m)} = \mu_t^{(n+m)} - \mu_t^{(n)}$ is the m-month forward rate n months from time $= t$

- $\mu_{t+n}^{(m)} = \mathcal{L}_{t+n}^{(m)}$ is the m-month spot rate at time $= t+n$
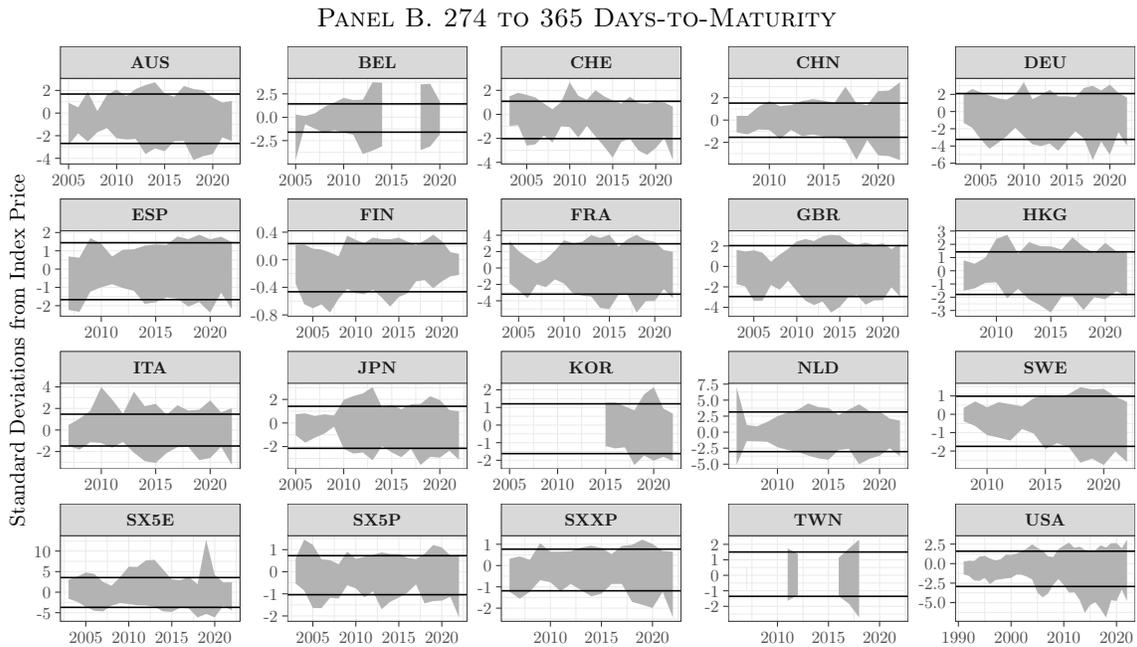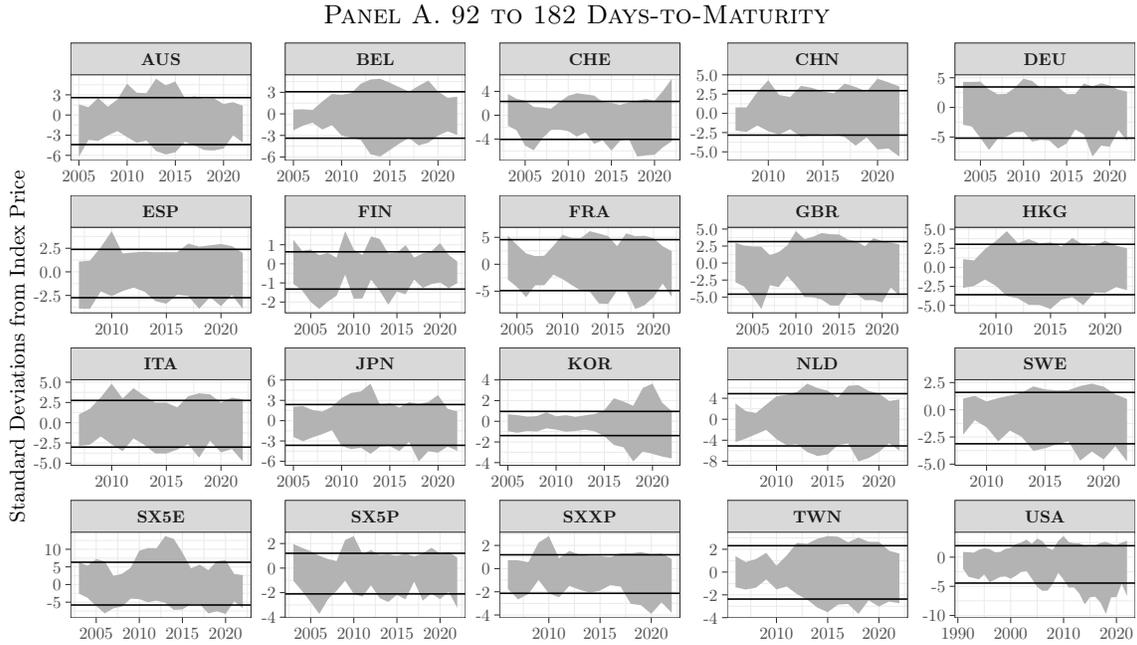
Figure A4
Current Spot and Forward Rates in the Full Sample

This figure 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) in the full sample. Spot and forward rates are for equity risk premia and are from option-based expectations. The sample is the longest available for each exchange.
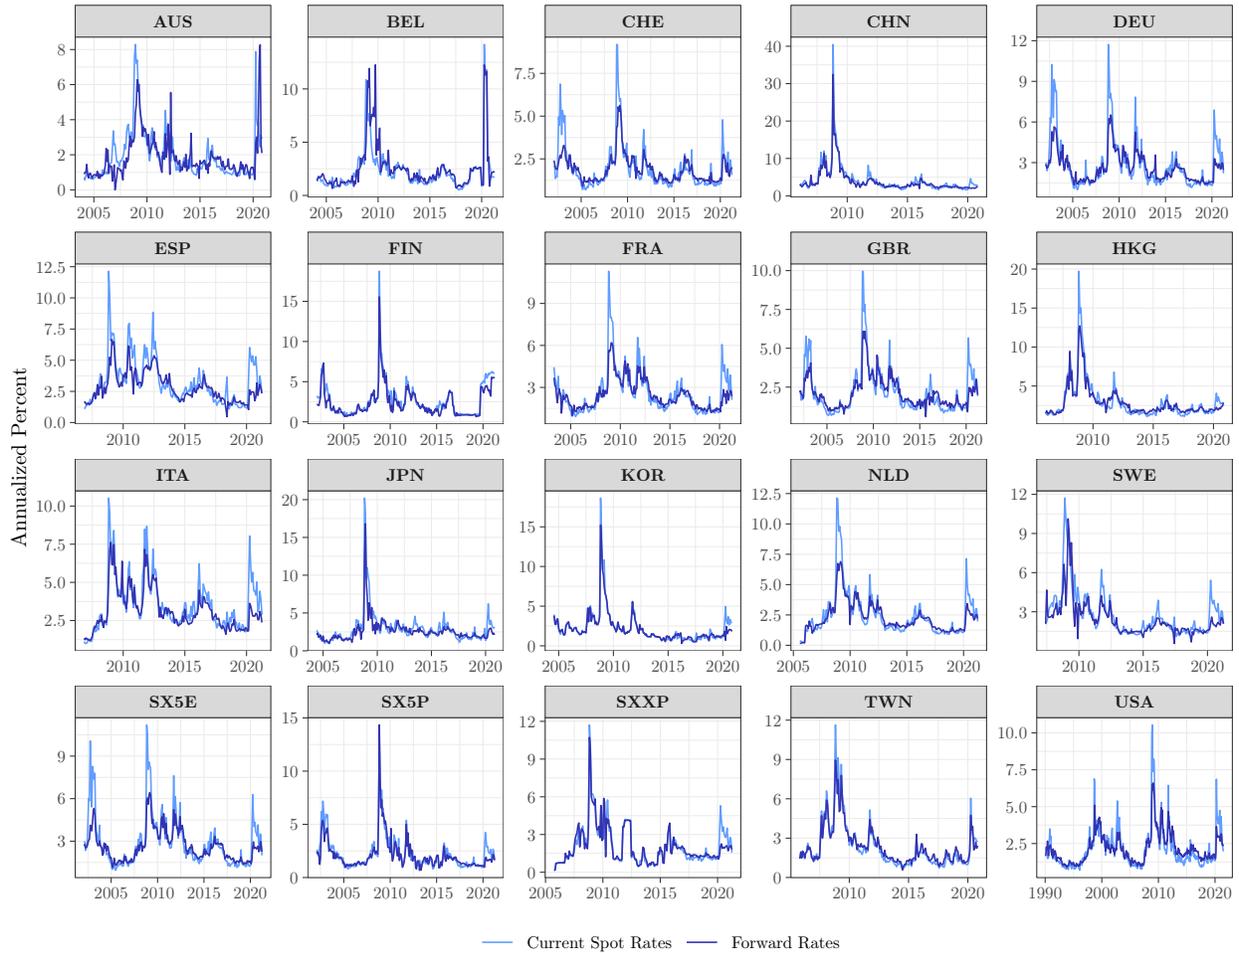
Figure A5
Measurement Error: Spot and Forward Rates

This figure describes truncation bias in the Black-Scholes model (left panel) and the stochastic volatility jump (SVJ) model (right panel). Integration bounds are in moneyness K/P units from the index price. Bar labels are in volatility standard deviations from the index price. Black-Scholes parameters: $P_t = 100$, $r = 0.05$, $q = 0.02$. SVJ parameters under $\mathbb{Q}$ are from Bakshi, Cao, and Chen (1997):

θ_v	κ_v	σ_v	ρ	μ_J	σ_J	λ
0.040	2.030	0.380	-0.570	-0.050	0.070	0.590

Under Black-Scholes (SVJ), good times correspond to low volatility $IV = 10\%$ ($\sqrt{v_t} = 10\%$), bad times to high volatility $IV = 60\%$ ($\sqrt{v_t} = 60\%$). The units are non-annualized basis points. See Appendix B.2 for more details.

PANEL A. BLACK-SCHOLES: WITHOUT EXTRAPOLATION

PANEL B. STOCHASTIC VOLATILITY JUMP: WITHOUT EXTRAPOLATION

PANEL C. STOCHASTIC VOLATILITY JUMP: WITH EXTRAPOLATION
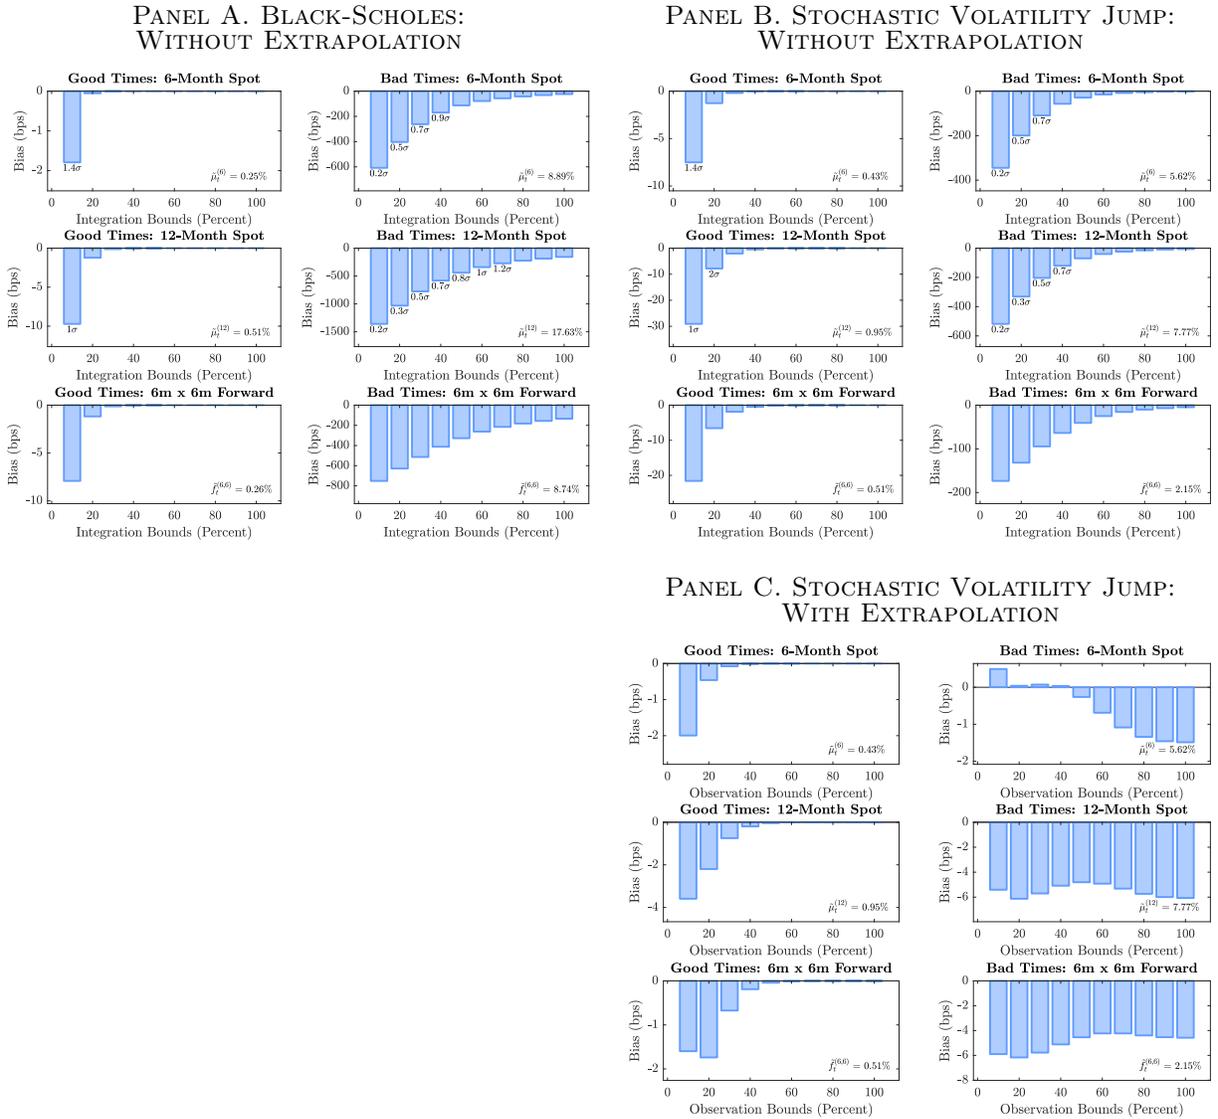
Figure A6
Measurement Error: Coibion-Gorodnichenko Regressions

This figure quantifies how much correlated measurement error is necessary to produce the Coibion-Gorodnichenko regression slopes in the data with monthly forecast revisions (left panel) and quarterly forecast revisions (right panel). The solid lines are the slopes in simulations. The shaded regions are 95% confidence bands in 50,000 samples. The blue circles are slopes in the data. The sample is the longest available for each exchange in the main sample. See Appendix C.1 for more details.

PANEL A. SHORT-HORIZON FORECAST REVISIONS
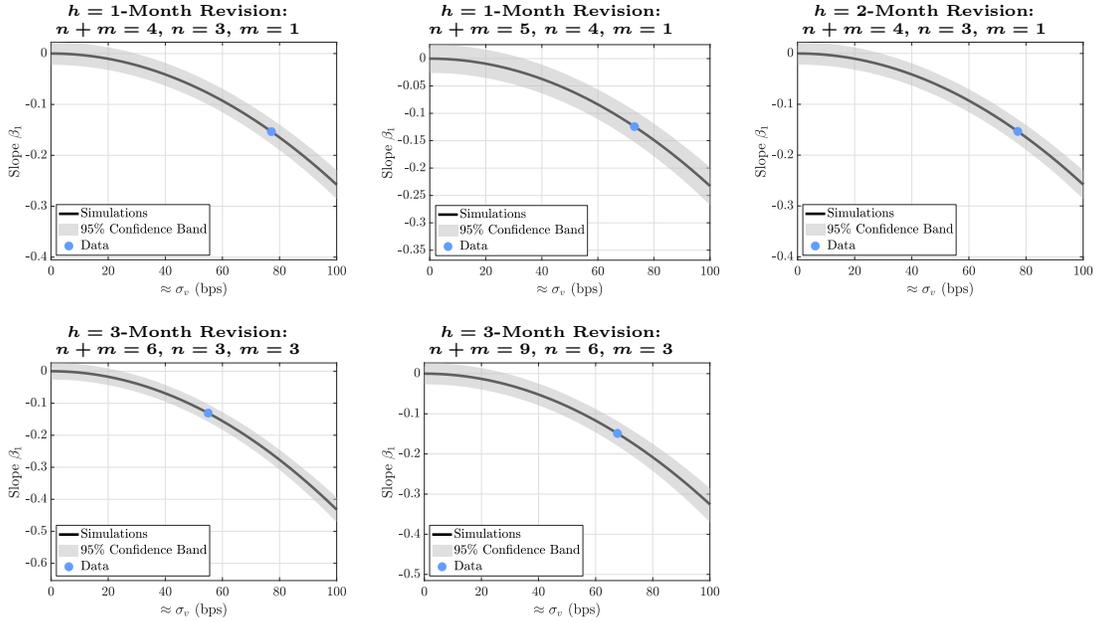
PANEL B. LONG-HORIZON FORECAST REVISIONS

OA-30
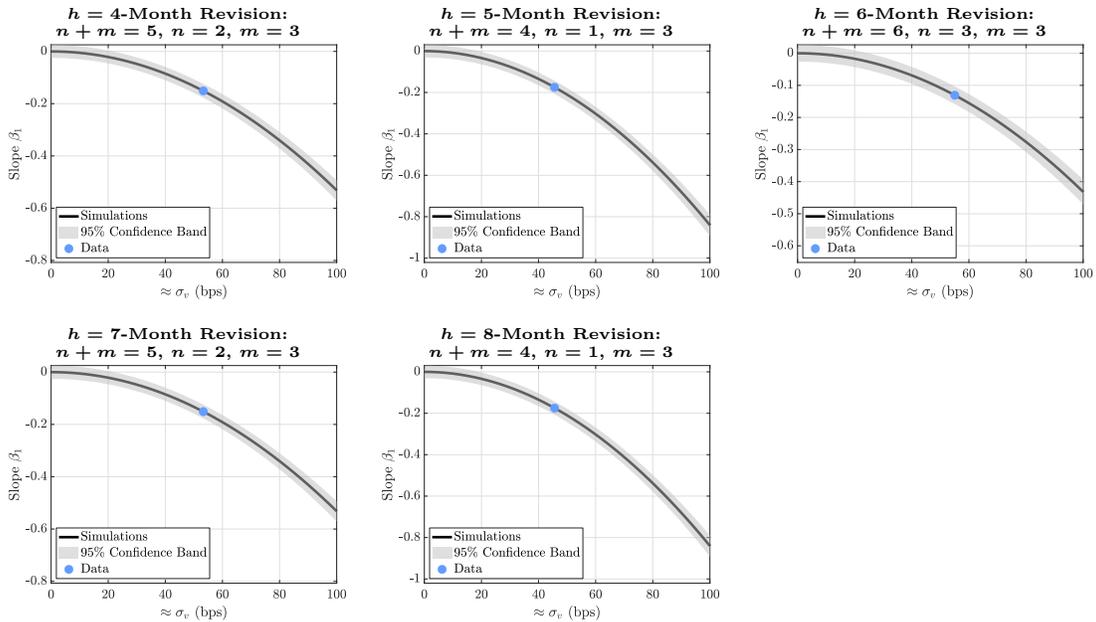
Figure A7
Power Utility: Regression Slopes and Average Forecast Errors

This figure reports regression slopes and average forecast errors from the standpoint of an unconstrained power utility investor fully invested in the market for option-based risk premia. See Panel A of Table A5 for more details. The sample is the longest available for each exchange in the main sample.
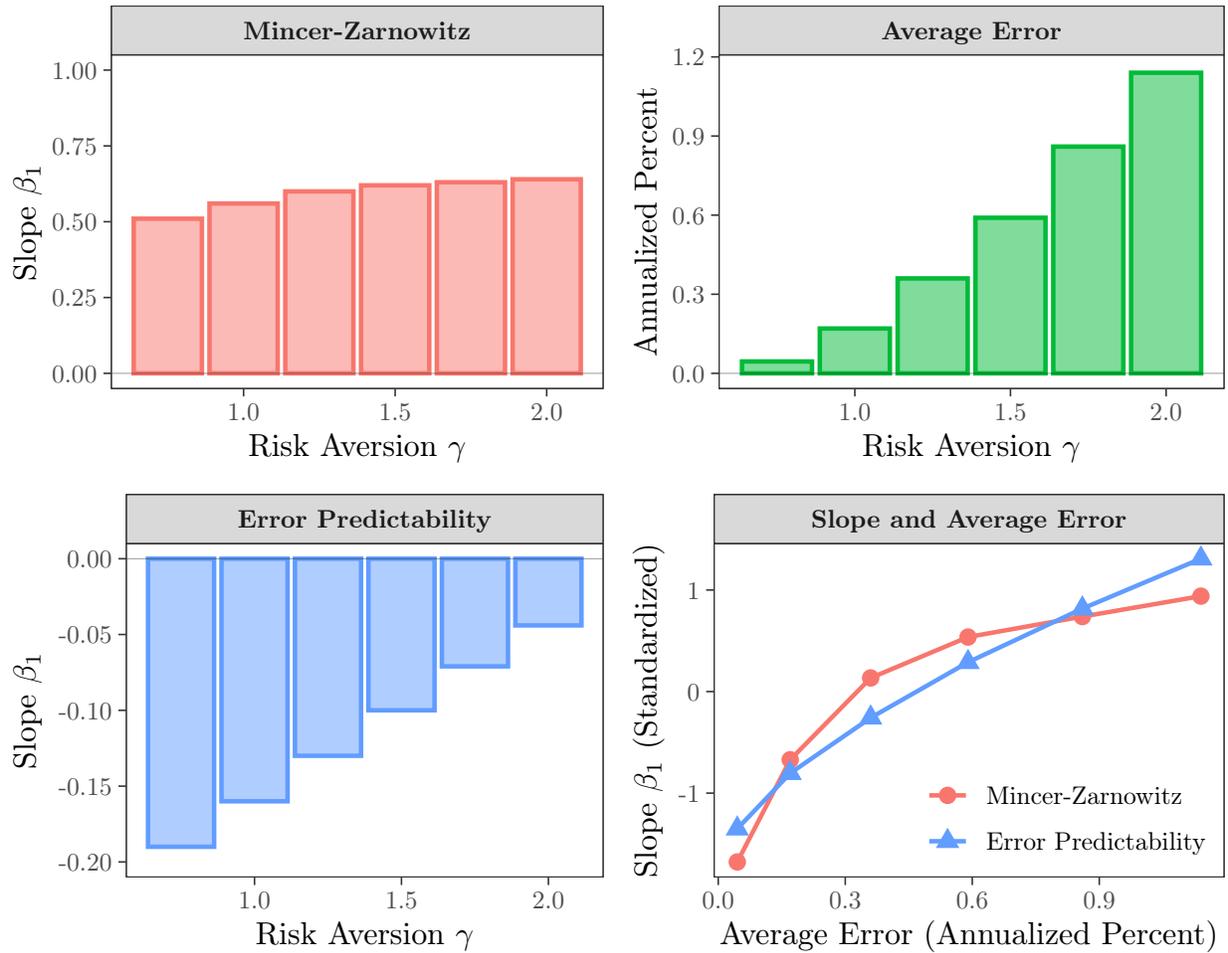
Figure A8
Model Calibration: Regression Fit

This figure reports regression R^2s in the calibrated model of expectation errors. The model is calibrated from the standpoint of an unconstrained log utility investor fully invested in the market. Table A11 reports the objective parameters. The solid lines are model-implied population R^2s in a single long sample. The shaded regions are model-implied 95% confidence bands in 10,000 short samples. The blue circles are model-implied R^2s under rational expectations with $\theta = 0$. The red squares are model-implied R^2s under diagnostic expectations with $\theta = 0.91$ from Bordalo, Gennaioli, and Shleifer (2018). The green triangles are R^2s in the data. The sample is the longest available for each exchange in the main sample.
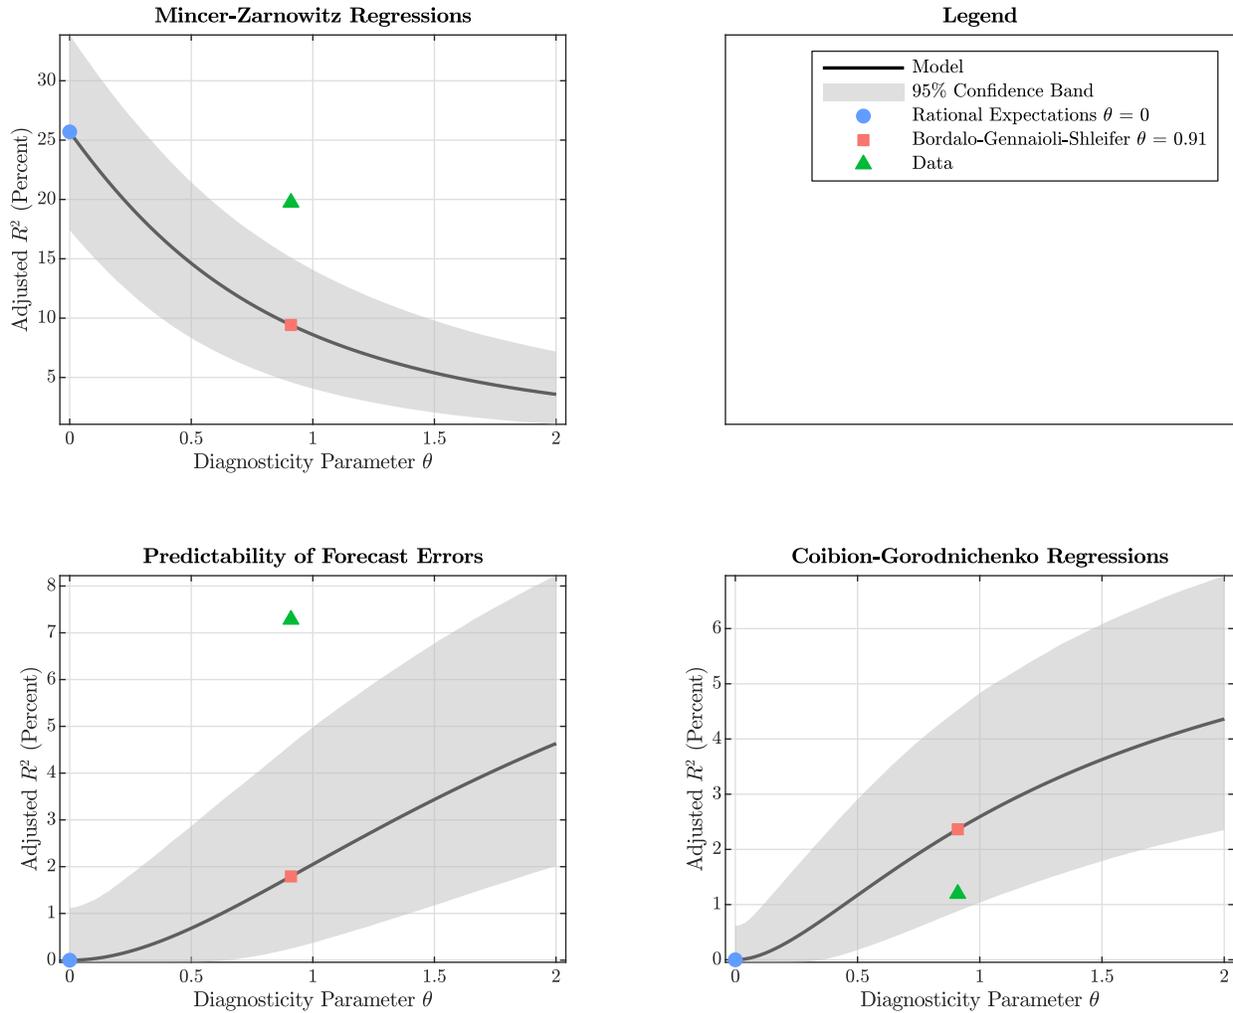
Appendix References

Acosta, Miguel. 2023. "The Perceived Causes of Monetary Policy Surprises." Working paper, Federal Reserve Board.

Aït-Sahalia, Yacine and Jefferson Duarte. 2003. "Nonparametric option pricing under shape restrictions." *Journal of Econometrics* 116 (1-2):9–47.

Andersen, Torben G., Oleg Bondarenko, and Maria T. Gonzalez-Perez. 2015. "Exploring Return Dynamics via Corridor Implied Volatility." *Review of Financial Studies* 28 (10):2902–2945.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen. 1997. "Empirical Performance of Alternative Option Pricing Models." *Journal of Finance* 52 (5):2003–2049.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan. 2003. "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options." *Review of Financial Studies* 16 (1):101–143.

Bakshi, Gurdip and Dilip Madan. 2000. "Spanning and derivative-security valuation." *Journal of Financial Economics* 55 (2):205–238.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. 2015. "X-CAPM: An extrapolative capital asset pricing model." *Journal of Financial Economics* 115 (1):1–24.

Bates, David S. 1996. "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options." *Review of Financial Studies* 9 (1):69–107.

Beason, Tyler and David Schreindorfer. 2022. "Dissecting the Equity Premium." *Journal of Political Economy* 130 (8):2203–2222.

Berger, David, Ian Dew-Becker, and Stefano Giglio. 2020. "Uncertainty Shocks as Second-Moment News Shocks." *Review of Economic Studies* 87 (1):40–76.

Bernanke, Ben S. and Kenneth N. Kuttner. 2005. "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *Journal of Finance* 60 (3):1221–1257.

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, Wouter Hueskes, Ralph Koijen, and Evert Vrugt. 2013. "Equity yields." *Journal of Financial Economics* 110 (3):503 – 519.

Black, Fischer and Myron Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81 (3):637–654.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2018. "Diagnostic Expectations and Credit Cycles." *Journal of Finance* 73 (1):199–227.

Campbell, John Y. 2017. "Discussion of Lawrence Jin and Pengfei Sui, 'Asset Pricing with Return Extrapolation'." NBER Behavioral Finance Meeting.

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. 2001. "Optimal positioning in derivative securities." *Quantitative Finance* 1 (1):19.

Carr, Peter and Liuren Wu. 2009. "Variance Risk Premiums." *Review of Financial Studies* 22 (3):1311–1341.

Coibion, Olivier and Yuriy Gorodnichenko. 2015. "Information rigidity and the expectations formation process: A simple framework and new facts." *American Economic Review* 105 (8):2644–78.

Constantinides, George M., Jens Carsten Jackwerth, and Alexi Savov. 2013. "The Puzzle of Index Option Returns." *Review of Asset Pricing Studies* 3 (2):229–257.

Culp, Christopher L., Yoshio Nozawa, and Pietro Veronesi. 2018. "Option-Based Credit Spreads." *American Economic Review* 108 (2):454–88.

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

Dew-Becker, Ian and Stefano Giglio. 2021. "Cross-sectional Uncertainty and the Business Cycle: Evidence from 40 Years of Options Data." *American Economic Journal: Macroeconomics* forthcoming.

Duarte, Jefferson, Christopher S. Jones, and Junbo L. Wang. 2022. "Very Noisy Option Prices and Inference Regarding the Volatility Risk Premium." Working paper, USC.

Figlewski, Stephen. 2010. "Estimating the Implied Risk-Neutral Density for the US Market Portfolio." In *Volatility and Time Series Econometrics: Essays in Honor of Robert Engle*, edited by Tim Bollerslev, Jeffrey Russell, and Mark Watson. Oxford University Press, 323–353.

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.

Gatheral, Jim. 2011. *The Volatility Surface: A Practitioner's Guide*. John Wiley & Sons.

Gatheral, Jim and Antoine Jacquier. 2011. "Convergence of Heston to SVI." *Quantitative Finance* 11 (8):1129–1132.

———. 2014. "Arbitrage-Free SVI Volatility Surfaces." *Quantitative Finance* 14 (1):59–71.

Golez, Benjamin and Ben Matthies. 2023. "Equity Term Structure Response to FOMC Announcements." Working paper, University of Notre Dame.

Gormsen, Niels Joachim and Christian Skov Jensen. 2022. "Higher-Moment Risk." Working paper, University of Chicago.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking* 1 (1):55–93.

Jiang, George J. and Yisong S. Tian. 2007. "Extracting Model-Free Volatility from Option Prices: An Examination of the VIX Index." *Journal of Derivatives* 14 (3):35–60.

Kelly, Bryan, Hanno Lustig, and Stijn Van Nieuwerburgh. 2016. "Too-Systemic-to-Fail: What Option Markets Imply about Sector-Wide Government Guarantees." *American Economic Review* 106 (6):1278–1319.

Kelly, Bryan, Ľuboš Pástor, and Pietro Veronesi. 2016. "The Price of Political Uncertainty: Theory and Evidence from the Option Market." *Journal of Finance* 71 (5):2417–2480.

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.

Leombroni, Matteo, Andrea Vedolin, Gyuri Venter, and Paul Whelan. 2021. "Central bank communication and the yield curve." *Journal of Financial Economics* 141 (3):860–880.

Malz, Allan M. 2014. "A Simple and Reliable Way to Compute Option-Based Risk-Neutral Distributions." Staff Report 677, Federal Reserve Bank of New York.

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

Martin, Ian W.R. and Christian Wagner. 2019. "What is the Expected Return on a Stock?" *Journal of Finance* 74 (4):1887–1929.

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.

Nakamura, Emi and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics* 133 (3):1283–1330.

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