NBER WORKING PAPER SERIES

DYNAMICS OF SUBJECTIVE RISK PREMIA

Stefan Nagel
Zhengyang Xu

Working Paper 29803
http://www.nber.org/papers/w29803

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2022

We are grateful for comments and suggestions to Klaus Adam, Nick Barberis, Ching-Tse Chen, Hui Chen, Ken French, Niels Gormsen, Ralph Koijen, Ian Martin, Monika Piazzesi, seminar participants at Boston College, University of Chicago, University of Southern California, University of Utah, and conference participants at the NBER Behavioral Finance Meeting. Nagel gratefully acknowledges financial support from the Fama-Miller Center at the University of Chicago Booth School of Business. Xu gratefully acknowledges financial support from the Research Grants Council of Hong Kong SAR, China (Project No. 21504421). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Dynamics of Subjective Risk Premia
Stefan Nagel and Zhengyang Xu
NBER Working Paper No. 29803
February 2022
JEL No. G12,G41

ABSTRACT

We examine subjective risk premia implied by return expectations of individual investors and professionals for aggregate portfolios of stocks, bonds, currencies, and commodity futures. While in-sample predictive regressions with realized excess returns suggest that objective risk premia vary countercyclically with business cycle variables and aggregate asset valuation measures, subjective risk premia extracted from survey data do not comove much with these variables. This lack of cyclicality of subjective risk premia is a pervasive property that holds in expectations of different groups of market participants and in different asset classes. A similar lack of cyclicality appears in out-of-sample forecasts of excess returns, which suggests that investors' learning of forecasting relationships in real time may explain much of the cyclicality gap. These findings cast doubt on models that explain time-varying objective risk premia inferred from in-sample regressions with countercyclical variation in perceived risk or risk aversion. We further find a link between subjective perceptions of risk and subjective risk premia, which points toward a positive risk-return tradeoff in subjective beliefs.

Stefan Nagel
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
stefan.nagel@chicagobooth.edu

Zhengyang Xu
Department of Economics and Finance
City University of Hong Kong
9-221, 9/F, LAU, Tat Chee Avenue
Kowloon Tong, Kowloon, Hong Kong SAR
zhengyxu@cityu.edu.hk

1 Introduction

Explaining the empirically observed time-variation of risk premia in asset markets has been one of the main challenges in asset pricing research in recent decades. Evidence from in-sample predictive regressions shows that excess returns of aggregate portfolios of equities, bonds, currencies, and commodity futures are predictable with slow-moving predictors such as the dividend-price ratio, yield spreads, or the futures basis. Broadly speaking, the in-sample predicted excess returns in these regressions—which we refer to as *objective* risk premia—are countercyclical, i.e., higher in recessions than in booms.

These findings sparked a large literature of rational expectations (RE) asset pricing models in which time-varying risk premia are generated via time-varying risk or time-varying risk aversion (see Cochrane (2017) for a review). In these models, investors are endowed with perfect knowledge of the underlying stochastic processes that generate payoffs. As a consequence, the objective law of motion of returns that generates the data observable to the econometrician ex post is the same as the law of motion that investors perceive in real time.

Yet, if investors are not endowed with such perfect knowledge, the dynamics of *subjective* risk premia they perceive in real time may differ systematically from the dynamics of the objective risk premia that an econometrician's predictive regressions extract from the data ex post. For example, if investors are learning about stochastic process parameters, there is a wedge between their real-time posterior beliefs about risk premia and the objective risk premia estimated by an econometrician from return data ex post with in-sample regressions. Fading memory, behavioral biases, and other imperfections can further magnify this wedge. In the presence of these wedges, one cannot infer perceived risks and risk aversion from objective risk premia estimates; subjective beliefs data is required.

In this paper, we present a pervasive stylized fact about subjective risk premia dynamics that holds true for different surveys, for expectations of individuals and professionals, for aggregate portfolios in different asset classes (stock market, Treasury bonds, currencies, commodity futures), and for a variety of state variables that have appeared in the literature to capture the general business cycles and asset-class specific valuation cycles: Subjective risk premia are substantially less cyclical than objective risk premia inferred from in-sample predictive regressions. This calls into question the common practice of equating time-variation in objective risk premia with time-varying risk or time-varying risk aversion. Subjectively perceived risk or risk aversion appears to be substantially less sensitive to cyclical state variables than is assumed in RE models calibrated to match time-variation in objective risk premia.

We use survey data from a variety of different sources to construct monthly or quarterly subjective return expectations. For stock market return expectations, we use data on individuals, CFOs, and professional forecasters. For bonds, currencies, and commodity futures, we have professionals' expectations only. To study the dynamics of subjective risk premia, we focus on expected returns in excess of Treasury yields over a one-year forecast horizon in all asset classes. We then project these subjective expected excess returns on state variables. By projecting realized excess returns on the same state variables in predictive regressions, we can then compare the cyclicality of subjective and objective risk premia.

In the first part of our analysis for each asset class, we focus on state variables that the prior literature has found to be good predictors of excess returns, for example, the dividend-price ratio for stock market returns, interest-rate cycle variables for bonds, interest rate spreads for currencies, and the futures basis for commodity futures. In virtually all combinations of state variables and different types of surveys, we find that projections of subjective risk premia on state variables yield coefficients that are substantially smaller in magnitude than projections of realized excess returns on the same state variables. In some cases, the point estimate for subjective risk premia has the opposite sign as for objective risk premia, but the overarching main regularity is that subjective risk premia vary a lot less with the state variables than future realized excess returns do in in-sample predictive regressions.

The state variables in this first part are asset-class specific, chosen based on earlier research that emphasized their role as significant excess return predictor for the specific asset class. On the other hand, many macro-finance models generate time-varying risk premia for different asset classes that are all tied to the state of the business cycle (Cochrane 2017). Could there be better alignment between subjective and objective risk premia if we focus on projections of subjective expected excess returns and realized excess returns for the different asset classes on a common set of business cycle variables? Perhaps the misalignment is due to factors idiosyncratic to the asset classes, not due to the common business-cycle component?

We find that this is not the case. We consider a number of commonly used business cycle state variables, including term and default spreads, industrial production growth, a volatility index, and a real activity factor extracted from a large number of macroeconomic time series. Consistent with prior literature, these cyclical variables forecast excess returns in predictive regressions in several asset classes. But for these business cycle variables, too, we find that their association with subjective expected excess returns is much weaker than in the in-sample predictive regressions using empirically realized returns.

Overall, combining the evidence from asset-specific predictors and the common cyclical predictors and across all asset classes and surveys, we find that the average magnitude by which a one standard deviation change in a predictor moves the subjective expected excess return is only about one fifth of the average magnitude by which it moves the objective expected excess return according to in-sample predictive regressions using realized excess returns. This pervasive lack of cyclical movements in subjective risk premia casts doubt on the time-varying risk or risk aversion mechanisms that have been used in many macro-finance models to explain cyclical fluctuations in asset prices.

Does the lack of cyclicality evident in subjective excess return expectations indicate gross mistakes on the part of investors and forecasters, or is there a plausible belief formation mechanism that does not imply easily detectable prediction errors that could be exploited by sophisticated investors? To shed light on this, we examine the cyclicality of out-of-sample (OOS) excess return forecasts. At every point in time, these OOS forecasts are constructed based on regression coefficients that are estimated using only past data. Moreover, at every point in time, the forecaster evaluates which linear combination of the predictive regression forecast and a simple trailing mean of excess returns would have done best in the past in forecasting OOS. The OOS forecast of excess returns is then based on this optimal linear combination. We find that these OOS forecasts are substantially less cyclical than in-sample forecasts. Moreover, if we construct the OOS forecasts with exponential weighting of past data to reflect fading memory as in Nagel and Xu (2022), or as combination forecasts of single-predictor forecasts as in Rapach et al. (2010), much of the cyclicality gap to the subjective risk premia in survey data disappears. Thus, while it is not our goal in this paper to provide evidence on a specific belief-formation mechanism, the OOS forecast evidence shows that the lack of cyclicality in subjective risk premia does not mean that the subjective forecasts are unreasonable or unsophisticated.

That subjective risk premia are largely unrelated to standard business cycle indicators and standard return predictors does not necessarily imply that subjective risk premia are constant. Indeed, we find that time-varying subjective risk perceptions appear to have some explanatory power for time-variation in subjective risk premia. This is in stark contrast to the literature on time-varying objective risk premia where researchers have not had much success linking them to time-varying risk. Due to data availability constraints, in this analysis we focus only on stock market expectations. For all three groups of market participants (individuals, CFOs, and professional forecasters) we find a positive link between the level of subjectively perceived stock market risk and their subjective risk premia. In other words, in terms of subjective beliefs, there is a positive risk-return tradeoff. The dynamics of perceived risks are heterogeneous between different groups of market participants, though. For example, professional investors' subjectively perceived stock market risk is more closely related to the VIX index than the risk perceived by respondents in the CFO survey. As documented recently by Lochstoer and Muir (2019), the CFO survey respondents seem to put substantially more weight on realized variance during the past year than the VIX does.

Our work connects to a number of areas in the literature. Our findings are most closely related to earlier work by Piazzesi et al. (2015) who show that subjective bond risk premia of professional forecasters are less cyclical than predictive regression forecasts.[1] We show that this relative lack of cyclicality result holds for a much broader range of state variables, not only for bonds but also for return expectations in stock market, for currencies, and for commodity futures, and in surveys of individual investors and CFOs. We further show OOS forecasts of excess return exhibit a similar lack of cyclicality.

Much of the earlier literature on subjective return expectations has focused on whether the sign of subjective return forecast lines up with the sign of predictive regression forecasts. For example, for stock market return expectations, Greenwood and Shleifer (2014) emphasize the procyclicality of individual investor return expectations. In contrast, Dahlquist and Ibert (2021) find that professional forecasters in the Livingston survey as well as asset managers issue countercyclical forecasts (see, also, Wu (2018), Wang (2021)). Our results show that these differences between surveys are overshadowed by an economically much bigger common regularity: the general lack of sensitivity of subjective risk premia to variables that forecast excess returns in predictive regressions (and in rational expectations models built to match the predictive regression evidence).

The lack of cyclical movements in subjective risk premia is consistent with models that generate volatile asset prices through time-varying subjective expectations of fundamentals growth rather than time-varying risk aversion or time-varying perceptions of risk. Examples include models of perpetual learning (Collin-Dufresne et al. (2017), Nagel and Xu (2022)) or diagnostic expectations (Bordalo et al. (2021)). In a similar vein, our evidence supports models that generate bond price and exchange rate variation through subjective beliefs about

[1] Other papers provide evidence that is related, but does not quite pin down the time-variation in the subjective risk premium for aggregate asset classes. Bacchetta et al. (2009) show that at country level, professional forecasters' subjective expectations errors for currency excess returns are predictable with the interest-rate differential, with the same sign and often similar magnitude as realized excess returns. However, from these country-level results it is not clear how objective and subjective risk premia for the aggregate currency portfolio would vary over time because the country-level results could be driven by country-specific components of the interest-rate differential and currency excess returns. From the viewpoint of macro-finance modeling, the properties of risk premia for aggregate strategy are arguably the most relevant. De La O and Myers (2021) find a small positive covariance between the price-dividend ratio and total return expectations from several surveys, but since the risk-free rate is correlated with the price-dividend ratio, their results are not directly informative about comovement of the price-dividend ratio with the subjective risk premium.

future interest rates, as in Piazzesi et al. (2015) and Gourinchas and Tornell (2004), respectively, and commodity futures price variation through subjective beliefs about future spot prices.

Price-growth extrapolation mechanisms as in Adam et al. (2017) and Jin and Sui (2021) are not consistent with a complete lack of comovement between state variables that predict realized excess returns and subjective risk premia. In these models, a positive correlation between past returns and future subjectively expected returns is sustained in equilibrium by a subjective risk premium that is increasing in past price growth. As a consequence, the price level, as captured, say, by the price-dividend ratio, is positively related to the subjective equity premium. However, as Jin and Sui (2021) show, this positive relationship can be quite weak in a calibrated version of their model, which could be broadly consistent with the lack of cyclicality in subjective risk premia that we find in this paper. With such a weak link between past price growth and subjective risk premia, asset price movements largely reflect changing expectations about future fundamentals growth, similar to models that directly target time-varying subjective expectations of fundamentals growth.

Finally, our finding that allowing for real-time learning helps explain the dynamics of subjective excess return expectations relates to recent work by Farmer et al. (2021). They show that a model in which agents are learning in real time about long-run dynamics helps explain macroeconomic forecasting anomalies in professional forecasts data.

The paper is organized as follows. Section 2 introduces the data. Section 3 presents the main results on the cyclicality of subjective and objective risk premia. Section 4 compares the cyclicality of subjective risk premia and OOS forecasts. Section 5 looks at the subjective risk-return tradeoff. Section 6 concludes.

2 Data

In this section we describe the survey data and how we construct subjective return expectations in each asset market. We also describe the return predictors and business cycle variables.

Appendix A provides more detailed information on the data sources and the exact timing assumptions that we use to match data on returns, predictor variables, and survey data.

2.1 Subjective expectations and returns data

Table 1 below reports the sample periods and moments of subjective return expectations in excess of risk-free rates. For comparison, we also report the moments of realized excess returns for the full sample period. Figure 1 plots the subjective excess return expectations in each asset market. We construct these measures as follows.

2.1.1 Stock market

For aggregate stock market return expectations we use a series of individual investor expectations, CFO expectations, and a series of professional forecaster expectations. The individual investor series is from Nagel and Xu (2022) (NX). It covers several surveys in the period 1972 to 1977 and has continuous quarterly coverage from 1987 onward. The series is constructed by combining information from the UBS/Gallup survey, the Conference Board survey, and the Michigan Survey of Consumers, plus several smaller surveys of brokerage and investment firm customers. The second source of stock return expectations is the Graham-Harvey CFO survey (Ben-David et al. 2013) which asks about the expectations of annual S&P 500 returns. The series is available from 2000Q3 at quarterly frequency (with gaps) and we use the mean of respondents' expected return each period.[2] The third source is the Livingston Survey maintained by the Federal Reserve Bank of Philadelphia. The series provides semi-annual forecasts starting in June 1952. From June 1992 onward, we use the ratio of twelve-month to zero-month mean level forecasts of the S&P500 stock index (or its predecessors) to measure price growth expectations. In earlier periods, where the zero-month nowcast is not available, we use the annualized ratio of twelve-month to six-month forecasts. We adjust these price

[2] Survey forecasts in 2001Q3, 2019Q1, and 2020Q1-3 are not available.

Table 1
Moments of Excess Returns and Excess Return Expectations

This table reports the frequencies, sample periods, means (%), and standard deviations (%) of realized one-year excess returns and survey-implied excess return expectations in each asset market. For stock market, we report the excess returns on the aggregate stock market. For Treasury bonds, we report the equal-weighted average excess returns on bonds with maturities of two, five, seven, and ten years. For foreign exchange, we report the average excess returns on a basket of developed market currencies. For commodities, we report the buy-and-hold excess returns on metals and crude oil, respectively.

	Frequency	Sample Period	Mean	S.D.
A. Stock				
Realized	Monthly	1926/12–2020/12	8.15	21.15
NX	Quarterly	1972Q3–2021Q2	5.95	1.98
CFO	Quarterly	2000Q3–2021Q2	3.98	1.75
Livingston	Semi-annually	1952Q2–2020Q4	4.09	5.86
B. Treasury bond				
Realized	Monthly	1952/03–2020/12	1.32	5.08
BCFF	Monthly	1988/01–2020/12	-0.76	1.70
C. Foreign exchange				
Realized	Monthly	1984/10–2021/06	1.78	9.76
CE & FX4casts	Monthly	1986/08–2021/06	0.50	3.16
D. Commodity				
Realized, Metals	Monthly	1978/09–2021/06	1.53	25.54
CE, Metals	Monthly	1995/08–2021/06	0.70	5.67
Realized, Oil	Monthly	1984/12–2021/06	8.49	33.68
CE, Oil	Monthly	1995/08–2021/06	2.46	9.70

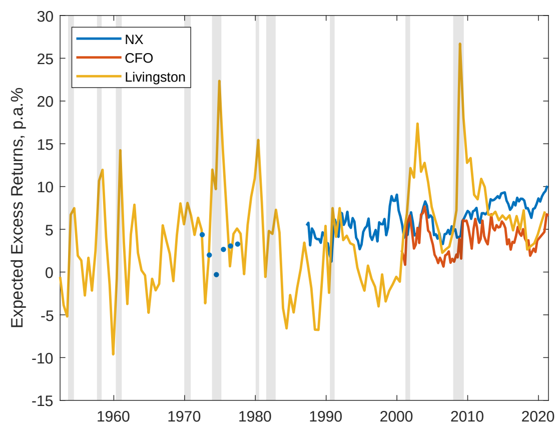

(A) Stock market

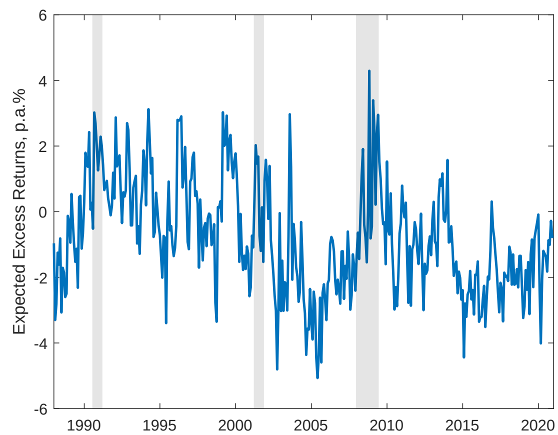

(B) Treasury bond market

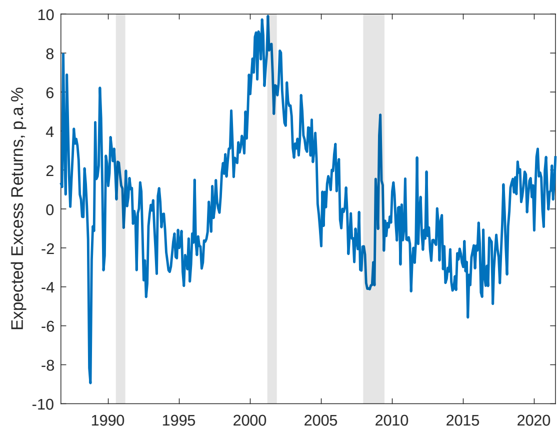

(C) FX market

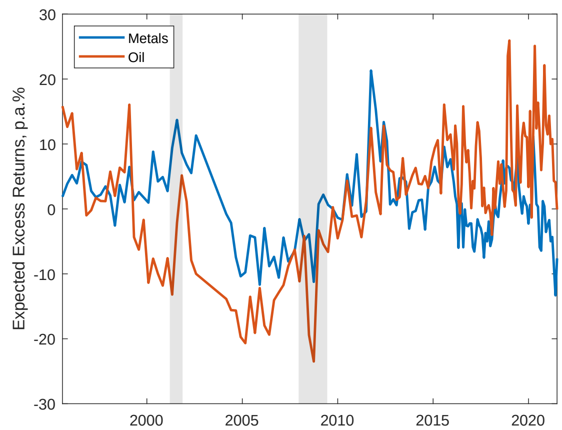

(D) Commodity market

FIGURE 1
Subjective Excess Return Expectations

Each panel plots the subjective excess return expectations implied by surveys. The grey-shaded areas indicate NBER recessions.

growth expectations with the dividend yield to arrive at return expectations.[3]

For each subjective one-year expected stock return series, $\tilde{\mathbb{E}}_t R_{t+1}$, we then construct the subjective risk premium by subtracting the one-year Treasury yield $R_{f,t}$ that prevailed at the time of the survey,

$$\tilde{\mathbb{E}}_t R^e_{t+1} = \tilde{\mathbb{E}}_t R_{t+1} - R_{f,t}. \tag{1}$$

For comparison, we also examine realized returns on the CRSP value-weighted index over the one-year forecast period. Realized returns are also expressed as excess returns over and above the one-year Treasury yield.

2.1.2 Treasury bond market

For bond return expectations of professional forecasters, we use the Blue Chip Financial Forecasts (BCFF) survey. Par yield forecasts are available for 6-month, 1-year, 2-year, 5-year, and 10-year bond maturities starting from January 1988. We take the average across forecasters in each survey month. Survey participants forecast the quarterly average of yields of a particular maturity for the current quarter, next quarter, and up to several quarters ahead. We follow Kim and Orphanides (2012) and treat these forecasts as approximately equal to a mid-quarter forecast. We interpolate linearly between these mid-quarter forecasts to obtain forecasts at a one-year horizon. To estimate a full expected par yield curve, we fit a Nelson-Siegel model to these forecasts period-by-period, as in Diebold and Li (2006). We then bootstrap zero-coupon yield forecasts from these par yield forecasts. The final step for each maturity is then to construct implied return expectations by relying on the approximation

$$\tilde{\mathbb{E}}_t R^n_{t+1} \equiv \tilde{\mathbb{E}}_t \left[\frac{(1+Y^n_t)^n}{\left(1+Y^{n-1}_{t+1}\right)^{n-1}} - 1 \right] \approx \frac{(1+Y^n_t)^n}{\left(1+\tilde{\mathbb{E}}_t Y^{n-1}_{t+1}\right)^{n-1}} - 1. \tag{2}$$

[3] We use the decomposition $\tilde{\mathbb{E}}_t R_{t+1} = \tilde{\mathbb{E}}_t(P_{t+1}/P_t) + (D_t/P_t)\tilde{\mathbb{E}}_t\,(D_{t+1}/D_t) - 1$. Similar to Adam et al. (2017), we set the expected dividend growth $\tilde{\mathbb{E}}_t(D_{t+1}/D_t)$ equal to the sample average of S&P annual dividend growth over the post-WWII sample period 1946–2020, which is 1.064. The dividend-price ratio is calculated using the S&P500 Index from CRSP.

Given the time-t one-year ahead forecast of the $(n-1)$-year zero-coupon yield $\tilde{\mathbb{E}}_t Y_{t+1}^{n-1}$ and the current n-year zero-coupon yield, we can then calculate $\tilde{\mathbb{E}}_t R_{t+1}^n$ for $n \in \{2, 5, 7, 10\}$. We aggregate to a bond portfolio return forecast $\tilde{\mathbb{E}}_t R_{t+1}$ by taking an equal-weighted average across maturities. Finally, we construct subjective expected excess returns, $\tilde{\mathbb{E}}_t R_{t+1}^e$, by subtracting the one-year Treasury yield from $\tilde{\mathbb{E}}_t R_{t+1}$.

To compare the subjective risk premium with predictive regression forecasts of realized excess returns, we use CRSP Treasury Index returns for maturities $n \in \{2, 5, 7, 10\}$. As for the subjective expectations, we form an equal-weighted average of these returns and subtract the one-year Treasury yield to obtain excess returns.

2.1.3 Foreign exchange market

In our analysis of currencies, we focus on the expected returns on a portfolio of developed market currencies from the perspective of a US investor. We use professional forecasts of exchange rates from FX4casts and Consensus Economics (CE) to construct subjective expected returns on this currency portfolio. We primarily use CE forecasts and supplement with FX4casts when CE forecasts are unavailable. Consensus forecasts in CE are arithmetic averages of individual forecasts; FX4casts uses a geometric average. With the time-t mean forecast $\tilde{\mathbb{E}}_t S_{t+1}^i$ of the one-year ahead spot exchange rate S_{t+1}^i of country i (in units of foreign currency per dollar), we approximate the subjective expected excess return on a US investor's position in an individual country's currency i as

$$\tilde{\mathbb{E}}_t R_{t+1}^{i,e} \equiv \tilde{\mathbb{E}}_t \left[\frac{F_t^i}{S_{t+1}^i} - 1 \right] \approx \frac{F_t^i}{\tilde{\mathbb{E}}_t S_{t+1}^i} - 1, \tag{3}$$

where F_t^i is the one-year forward rate at time t. Developed market currencies in our sample include Australia, Canada, Denmark, Germany (replaced by the Euro from 1999 onward), Japan, New Zealand, Norway, Sweden, Switzerland, and United Kingdom. To avoid a shift in sample composition at the Euro introduction in 1999, we use the German Mark as the only representative from the Eurozone currencies prior to Euro introduction. We calculate the expected excess return on the developed market currency portfolio as an equal-weighted average of the individual currencies' $\tilde{\mathbb{E}}_t R^{i,e}_{t+1}$.

The realized excess returns that we compare with subjective risk premia are constructed in similar fashion from forward and spot rates as $R^{i,e}_{t+1} = F^i_t / S^i_{t+1} - 1$.

2.1.4 Commodity futures market

We use commodity spot price forecasts to compute subjective expected returns on buy-and-hold futures positions. We obtain consensus spot price forecasts from CE that cover major commodities in metals and energy. The publication was quarterly before 2012, bi-monthly until 2015, and switched to a monthly release cycle since February 2016.[4] We use the set of commodities with the longest history and highest coverage in CE: WTI crude oil and several metals (aluminum, copper, gold, and silver). These commodities cover the major components of the S&P GSCI index.[5] Since CE commodity forecasts are also for quarterly average, we use the same linear interpolation as earlier for bonds to obtain one-year forecasts.

Using these spot price forecasts, we then calculate the expected excess return from entering a one-year futures position at time t at the one-year futures price F_t and holding it until it gets settled at maturity at $t+1$ at the spot price S_{t+1} (or a futures price very close to the spot price shortly before maturity):

$$\tilde{\mathbb{E}}_t R^e_{t,t+1} \equiv \frac{\tilde{\mathbb{E}}_t S_{t+1}}{F_t} - 1. \tag{4}$$

We obtain spot price and futures data from Datastream and Bloomberg. When spot price data is not available for a commodity, we follow Koijen et al. (2018) and extrapolate spot prices from Bloomberg generic futures data. We construct a sector-level subjective risk premium series for metals by taking an equal-weighted average of the subjective expected

[4] During the period of August 2002 to March 2004, and the third quarter of 2007, there are no forecasts available as no surveys were undertaken by the previous owners of the publication.

[5] Based on the 2021 reference percentage dollar weights, WTI crude oil accounts for 40.4% of the energy sector and the four metals we include account for 84.4% of the metals sector.

excess returns across metals.

We construct realized excess returns along similar lines as $R^e_{t+1} = S_{t+1}/F_t - 1$, with equal-weighted average within the metals sector.

2.2 Return Predictors and Cyclical Indicators

In this section, we introduce the main return predictors and cyclical indicators that we use in our analysis. For each asset class, we examine several asset class-specific return predictors that have appeared in the prior literature. We also examine a common set of business cycle indicators. Table 2 reports the correlations of return predictors and cyclical indicators.

2.2.1 Asset-class specific predictors

The stock market return predictors are the log consumption-wealth ratio (CAY, quarterly, starting in 1951Q1) from Lettau and Ludvigson (2001), the repurchase-adjusted log dividend-price ratio of the CRSP value-weighted index (D/P, monthly, starting end of 1926) and a long-run exponential average of past per-capita real aggregate dividend growth (EXPD, quarterly, starting end of 1926) from Nagel and Xu (2022), as well as net equity expansion (NTIS, monthly, starting from end of 1926) from Welch and Goyal (2008), calculated as the ratio of twelve-month moving sums of net issues by NYSE-listed stocks divided by the total market capitalization of NYSE stocks at the end of the twelve-month window.

For Treasury bonds, we use two asset-class specific predictors. The first predictor is the macro factor from Ludvigson and Ng (2009) (LN). It is constructed as the fitted value from regressing equal-weighted bond portfolio excess returns on lagged principal components of a broad set of macro variables. We use the six-factor version of their model. The second predictor is the cycle factor from Cieslak and Povala (2015) (CYCLE). To construct this factor, we first regress zero-coupon yields of maturities between one to fifteen years from Liu and Wu (2021) on trend inflation to obtain the short-maturity cycle and the average longer-maturity cycles. The fitted value from regressing equal-weighted bond portfolio excess

Table 2

Correlations of Return Predictors and Cyclical Indicators

This table reports the unconditional correlations of asset-class specific predictors and business-cycle indicators. For countercyclical business-cycle indicators, N-IP is the negative of the 12-month log change in the U.S. industrial production index; TERM is the spread between long interest rates from Robert Shiller's website and 3-month Treasury yields; DEFAULT is the spread between Moody's Seasoned Baa and Aaa Corporate Bond yields; F1 is the real factor from Ludvigson and Ng (2009); VIX^2 is the square of the CBOE volatility index; CAY is the aggregate log consumption-wealth ratio from Lettau and Ludvigson (2001); D/P and EXPD are the repurchase-adjusted log dividend-price ratio and the experienced payout growth from Nagel and Xu (2022), respectively; NTIS is the net equity expansion from Welch and Goyal (2008); LN is the macro factor from Ludvigson and Ng (2009); CYCLE is the cycle factor from Cieslak and Povala (2015); FD is the average forward discount on the basket of developed market currencies; BASIS is the percentage difference between futures and spot prices and OI is futures open interest growth for metals (M) and oil (E), respectively.

	N-IP	TERM	DEFAULT	F1	VIX^2	CAY	D/P	EXPD	NTIS	LN	CYCLE	FD	BASIS (M)	OI (M)	BASIS (E)	OI (E)
N-IP	1.00															
TERM	0.05	1.00														
DEFAULT	0.26	0.20	1.00													
F1	0.63	0.15	0.51	1.00												
VIX^2	0.08	-0.00	0.31	0.31	1.00											
CAY	-0.17	0.22	-0.05	0.08	-0.04	1.00										
D/P	0.17	0.10	0.50	0.31	0.15	-0.04	1.00									
EXPD	0.15	-0.39	-0.13	-0.05	0.06	-0.11	-0.04	1.00								
NTIS	0.01	-0.16	-0.18	-0.29	-0.02	0.02	-0.14	0.04	1.00							
LN	0.48	0.54	0.21	0.50	0.18					1.00						
CYCLE	-0.19	0.50	-0.13	-0.02	-0.01					0.18	1.00					
FD	0.22	0.43	0.12	0.29	0.02							1.00				
BASIS (M)	0.02	-0.31	0.47	-0.00	-0.10								1.00			
OI (M)	-0.03	0.17	-0.14	-0.16	-0.17								-0.14	1.00		
BASIS (E)	0.25	-0.04	0.28	0.20	0.28								-0.03	0.00	1.00	
OI (E)	-0.30	0.04	-0.19	-0.31	-0.20								-0.00	0.34	-0.43	1.00

returns on these two cycle components then represents the cycle factor.

In our analysis of foreign exchange, we consider a US investor's position in a portfolio of developed market currencies as in Lustig et al. (2014). Following their analysis of excess return predictability, we use the average one-year forward discount as predictor. For each currency i we construct the forward discount,

$$FD_t^i \equiv \frac{F_t^i}{S_t^i} - 1, \tag{5}$$

from one-year forward rates F_t^i and spot rates S_t^i. We then average these currency-level forward discounts across all currencies in the portfolio (FD).

Finally, for commodities we use two asset-class specific predictors. The first predictor, as in Hong and Yogo (2012), is the one-year futures basis defined as

$$Basis_t \equiv \frac{F_t}{S_t} - 1, \tag{6}$$

where F_t is the one-year futures price and S_t is the spot price (approximated by extrapolating from generic futures if not available, as discussed above). We average the commodity-specific bases within metals futures (BASIS). The second predictor is open interest growth following Hong and Yogo (2012). Using CFTC Commitments of Traders data, we multiply open interest in terms of number of outstanding contracts by the number of units per contract and the spot price per unit to get dollar open interest.[6] We then sum up dollar open interests within each sector. As a final step, we calculate percentage growth rates over one-year periods (OI).

2.2.2 Business-cycle indicators

The first business-cycle indicator is industrial production growth constructed as the year-over-year log growth using data from FRED database at the Federal Reserve Bank of St.

[6] For Aluminum, we use open interest data from London Metal Exchange instead given the limited coverage of Aluminum in CFTC Commitments of Traders data.

Louis. For ease of interpretation, we take the negative of industrial production growth (N-IP) so that the series is countercyclical. The second indicator is the term spread (TERM) between long- and short-maturity yields. For long-maturity yields, we use the long interest rates from Robert Shiller's website. For short-maturity yields, we use the annualized three-month Treasury bill yields. The third indicator is the default spread (DEFAULT) between the Moody's Seasoned Baa and Aaa yields, obtained from FRED. The fourth indicator is the real factor (F1) from Ludvigson and Ng (2009) which is the first principal component of a broad set of macroeconomic variables. The last indicator we use is the square of CBOE Volatility Index (VIX^2).[7] Given its relatively short history, we supplement the series with the news-implied volatility index (NVIX) from Manela and Moreira (2017). We estimate the relation between VIX and NVIX in the overlapping samples and impute VIX before 1990 with the fitted coefficients from this regression.[8] All cyclical indicators are available at a monthly frequency. We lag N-IP and F1 by one-month to account for publication lags.

3 Dynamics of Subjective and Objective Risk Premia

We now examine how much subjective and objective risk premia vary with asset-class specific predictors and the business-cycle indicators. We regress subjective expected excess returns from surveys on the predictor variables and compare the results to regressions in which future realized excess returns are the dependent variable. For ease of interpretation, we standardize all predictor variables to have unit standard deviations in the full sample that we use in the realized excess return regressions.[9] With standardized predictors we can easily compare the amount of variation in objective and subjective risk premia associated with

[7] We use VIX^2 instead of VIX because asset pricing theories typically relate risk premium to variance, not standard deviation.

[8] The regression uses monthly data from January 1990 to March 2016 and yields:

$$VIX = -2.940 + 0.919 \times NVIX,$$

with an R^2 of 0.59.

[9] If predictor data is available, the full sample start date is December 1925, and otherwise at the earliest date of predictor availability.

different predictors by comparing the magnitudes of slope coefficients. For parsimony, while all regressions include an intercept, we omit the intercept estimates from all tables.

The realized returns regressions are run with monthly or quarterly data depending on data availability for the predictor variable. To match with the one-year forecast horizon in our subjective expected return data, we use realized excess returns compounded over one-year periods as the dependent variable. We use a stationary block bootstrap to account for the autocorrelations in residuals induced by the overlapping return windows in the dependent variable. We also use the bootstrap to adjust coefficients for the predictive regression bias discussed in Stambaugh (1999). Appendix B provides more details of the bootstrap approach. In regressions where subjective expected returns are the dependent variable, we adjust p-values for heteroskedasticity and autocorrelation in regression residuals using the equal-weighted cosine (EWC) approach from Lazarus et al. (2018).

3.1 Stock market

To provide a basis for comparison, Table 3 first reports the results for predicting future one-year realized excess returns. All asset-class specific predictors except NTIS are at least marginally significant ($p < 0.10$) when used as sole predictor variable. The magnitude of slope coefficient point estimates are similar and quite large for all of them. Based on bias-adjusted slope coefficient estimates, a one standard deviation move in the predictor variable is associated with a substantial change in objective expected returns of 3.24 to 6.44 percentage points over a one-year forecast horizon.

For cyclical indicators, TERM and F1 are the only ones that are individually at least marginally significant. The magnitudes of point estimates are also much more mixed than for the asset-class specific predictors, ranging from 1.55 for N-IP to 4.54 for F1. However, without exception, for all predictors in both groups, the point estimates indicate a countercyclical pattern of objective risk premia: low following booms, high following downturns in asset prices or the business cycle.

Table 3
Regressing Realized Stock Market Excess Returns on Predictors

Dependent variable is the one-year cumulative excess returns on the CRSP index. The full sample period is from December 1926 to December 2020. Column variables are described in Section 2.2. R^e_{past} denotes the past one-year excess return. All independent variables are standardized to have unit standard deviations in the full sample period. In each block, the first row reports the OLS estimates multiplied by 100; bootstrap bias-adjusted coefficients are reported in braces; bootstrapped p-values are reported in parentheses. We use a stationary bootstrap with an optimal block length determined as in Politis and White (2004).

	Asset-Class Specific								Business-Cycle									
	CAY		D/P		EXPD		NTIS		N-IP		TERM		DEFAULT		F1		VIX^2	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
Coeff	4.78	5.34	6.40	6.68	-6.48	-7.29	-5.23	-5.23	1.74	1.53	4.13	4.30	3.32	3.24	4.24	3.92	1.49	1.22
{bias-adj.}	{3.24}	{3.71}	{6.01}	{5.83}	{-6.44}	{-7.19}	{-4.87}	{-4.93}	{1.55}	{1.26}	{4.40}	{4.62}	{2.84}	{2.51}	{4.54}	{4.16}	{1.82}	{1.48}
(p-value)	(0.08)	(0.10)	(0.00)	(0.00)	(0.03)	(0.04)	(0.16)	(0.16)	(0.51)	(0.49)	(0.00)	(0.01)	(0.12)	(0.13)	(0.00)	(0.04)	(0.31)	(0.39)
R^e_{past}		-3.52		0.87		-3.51		-1.25		-0.70		-1.28		-0.24		-1.07		-1.01
{bias-adj.}		{-3.32}		{0.36}		{-3.70}		{-1.47}		{-1.23}		{-1.39}		{-0.55}		{-0.65}		{-1.24}
(p-value)		(0.12)		(0.71)		(0.09)		(0.51)		(0.74)		(0.55)		(0.91)		(0.69)		(0.62)
Adj. R^2	0.03	0.06	0.09	0.09	0.08	0.11	0.06	0.06	0.01	0.01	0.04	0.04	0.02	0.02	0.06	0.06	0.00	0.00
N	272	272	1117	1117	377	373	1117	1117	1129	1117	1129	1117	1129	1117	717	717	1129	1117

When past one-year realized excess returns are included as an additional predictor in the regressions, past returns are never statistically significant except marginally so when used along with EXPD ($p = 0.09$).

Table 4 shows the results of regressing subjective excess return expectations from different surveys on the same predictor variables as in Table 3. Panel A presents the results for individual investor expectations. Results for CFO expectations are shown in Panel B. Looking across all regressions, a striking commonality is that the magnitude of slope coefficients is generally much smaller than in the realized excess return regressions in Table 3. For example, in the single-predictor regressions, the biggest coefficient in Table 4 is 1.12 for TERM in Panel B. In contrast, the largest coefficient estimate in absolute magnitude is 6.44 for EXPD in in Table 3 (for TERM it is 4.40). Therefore, subjective expected excess returns of individual investors and CFOs vary much less with these standard predictor variables than objective risk premia do.

With subjective expected excess returns from the Livingston survey in Panel C, the magnitudes of estimated slope coefficients for the asset-class specific predictors are also fairly small. Only for the cyclical indicators, we obtain coefficients (5.62 for N-IP, 4.18 for DEFAULT) that are in the ballpark of the magnitudes that we obtained in the realized return regressions. However, for these two predictors, the coefficients in the realized returns regression in Table 3 are actually quite small (1.55 and 2.84, respectively), so there is not a great match between subjective and objective risk premia in these cases either.

Since evidence for extrapolation from recent past returns play an important role in earlier studies of return expectations data, e.g., Greenwood and Shleifer (2014), we also run versions of these regressions with past one-year excess returns as a regressor. Past returns often have substantial explanatory power, with positive effects on subjective risk premia of individuals and CFOs, indicating extrapolation, and negative effects for professionals in the Livingston survey, indicating a contrarian influence of past returns. Therefore, recent past returns are a potentially important influence on subjective return expectations that one should seek to

Table 4
Regressing Survey Stock Market Excess Return Expectations on Predictors

Dependent variables are the one-year-ahead excess return expectations from Nagel and Xu (2022), the CFO survey, and the Livingston survey, respectively. Column variables are described in Section 2.2. R^e_{past} denotes the past one-year excess return. All predictors are standardized to have unit standard deviations in the full sample period. The first row in each block reports the OLS estimates multiplied by 100. EWC p-values following Lazarus et al. (2018) are reported in parentheses.

	Asset-Class Specific								Business-Cycle									
	CAY		D/P		EXPD		NTIS		N-IP		TERM		DEFAULT		F1		VIX^2	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
A. Individual (NX)																		
Coeff	-0.75	-0.75	-0.24	-0.00	0.17	0.20	-0.36	-0.55	0.98	1.57	0.44	0.46	-0.30	0.57	0.15	0.63	-0.01	0.27
(p-value)	(0.04)	(0.03)	(0.52)	(1.00)	(0.63)	(0.54)	(0.48)	(0.23)	(0.17)	(0.01)	(0.20)	(0.18)	(0.35)	(0.13)	(0.68)	(0.08)	(0.95)	(0.14)
R^e_{past}		0.87		0.87		0.88		0.98		1.16		0.88		1.07		1.13		1.13
(p-value)		(0.01)		(0.02)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.01)
Adj. R^2	0.21	0.33	0.00	0.11	0.00	0.12	0.01	0.16	0.04	0.23	0.04	0.16	-0.00	0.13	-0.00	0.16	-0.01	0.15
N	142	142	142	142	142	142	142	142	143	142	143	142	143	142	142	142	143	142
B. CFO																		
Coeff	-0.23	-0.15	-0.47	-0.41	-0.94	-0.92	0.74	0.67	0.32	1.64	1.12	1.15	-0.15	0.47	0.08	0.59	0.01	0.23
(p-value)	(0.64)	(0.76)	(0.32)	(0.31)	(0.01)	(0.02)	(0.01)	(0.03)	(0.55)	(0.02)	(0.00)	(0.00)	(0.60)	(0.18)	(0.68)	(0.07)	(0.92)	(0.16)
R^e_{past}		0.35		0.30		0.16		0.20		0.89		0.39		0.63		0.85		0.73
(p-value)		(0.18)		(0.29)		(0.60)		(0.48)		(0.01)		(0.08)		(0.08)		(0.03)		(0.05)
Adj. R^2	-0.00	0.01	0.04	0.05	0.35	0.34	0.08	0.07	-0.01	0.12	0.37	0.44	-0.01	0.03	-0.01	0.06	-0.01	0.04
N	77	77	77	77	77	77	77	77	79	77	79	77	79	77	77	77	79	77
C. Livingston																		
Coeff	-1.70	-1.42	1.01	0.42	-0.37	-0.59	-1.05	-0.88	5.62	5.08	1.06	1.03	4.18	3.41	2.67	2.23	1.65	1.23
(p-value)	(0.08)	(0.09)	(0.48)	(0.73)	(0.40)	(0.17)	(0.43)	(0.34)	(0.00)	(0.00)	(0.38)	(0.37)	(0.03)	(0.04)	(0.00)	(0.00)	(0.00)	(0.02)
R^e_{past}		-2.70		-2.78		-2.94		-2.83		-2.28		-2.87		-1.79		-1.67		-2.23
(p-value)		(0.02)		(0.01)		(0.01)		(0.01)		(0.01)		(0.01)		(0.05)		(0.08)		(0.06)
Adj. R^2	0.06	0.17	0.01	0.13	-0.00	0.14	0.01	0.14	0.27	0.35	0.03	0.16	0.22	0.26	0.23	0.26	0.13	0.20
N	138	138	138	138	138	138	138	138	138	138	138	138	138	138	122	122	138	138

understand. The heterogeneity in past returns' influence on expectations between individuals/CFOs on one hand and professional forecasters on the other hints that a model with heterogeneous beliefs may be needed to understand the asset pricing consequences of belief formation, perhaps as in Barberis et al. (2015) where extrapolators coexist with traders that are rationally contrarian (although predictions from this model would be difficult to square with the lack of predictability associated with past returns in the realized return regressions in Table 3). However, for our purposes, the important take-away from the table is that the inclusion of past returns does not have a substantial impact on the slope coefficient estimates for the slower-moving asset-class specific predictors and cyclical indicators that are the focus of our analysis. Whatever the mechanism that let past returns affect return expectations, this mechanism seems to be largely distinct from the predictor variables and the low-frequency valuation cycles that have been emphasized in macro-finance modeling of asset markets.

Figure 2 provides a visual summary of our results on objective and subjective risk premia in the stock market. While our discussion above focused on estimates in single-predictor regressions, this figure shows the coefficient estimates when past returns are included in the regression along with the predictor. The blue bars show the coefficients from the realized return regressions in Table 3. The other bars show the corresponding coefficients for the same predictor with the survey expected excess returns as dependent variable from Table 4. The figure clearly illustrates that subjective risk premia have much lower sensitivity to the predictor variables—with the exception of the Livingston survey for N-IP and DEFAULT—than the objective risk premia captured by the predictive regressions with realized returns.

In stark contrast to the results in Figure 2, in a RE model of time-varying risk premia, the sensitivities of objective and subjective risk premia would be the same. Specific details of the results for particular predictors aside, the fact that subjective risk premia do not seem to move much with standard cyclical variables casts doubt on the empirical importance of time-varying risk or time-varying risk aversion as a major source of aggregate stock market price movements.

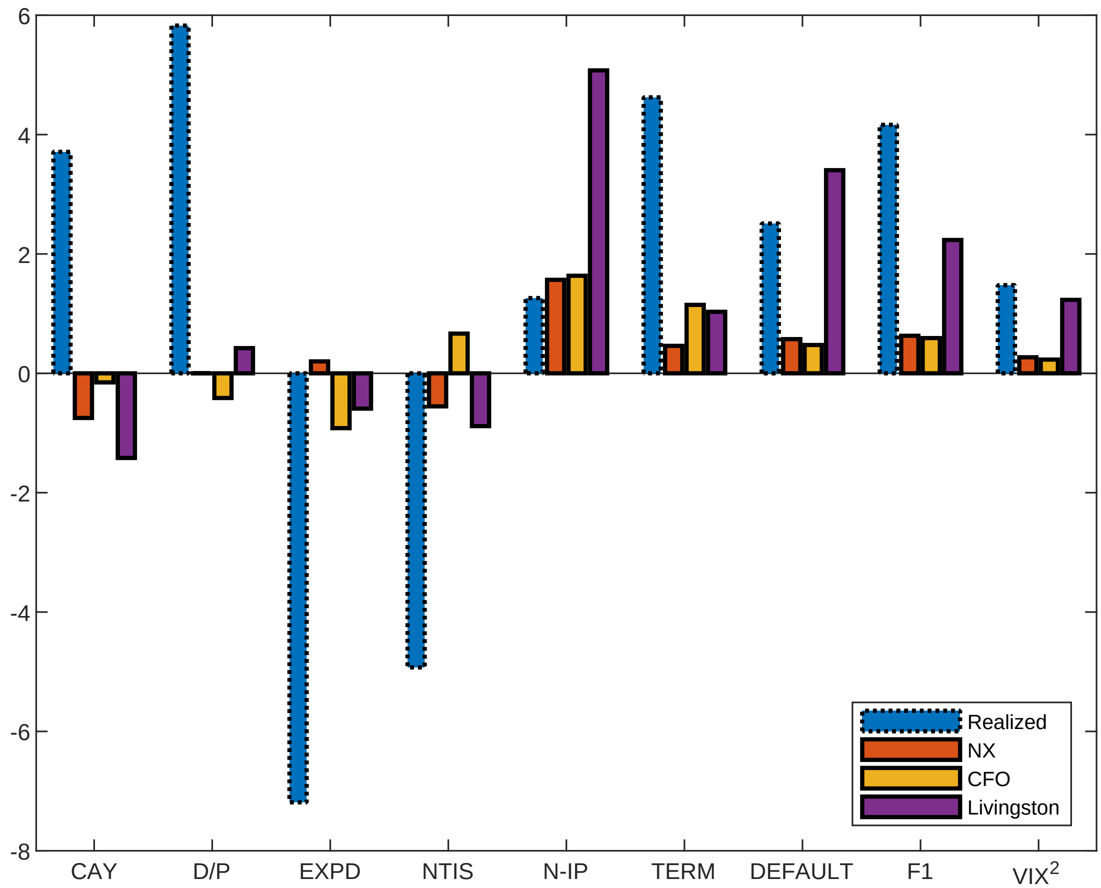

Figure 2
Coefficients from Regressing Stock Excess Returns on Predictors

The blue bars show the slope coefficients from regressing realized stock market excess returns on asset-class specific predictors and business-cycle indicators. The estimates are bootstrap bias-adjusted. The red, yellow, and purple bars use survey excess return expectations from Nagel and Xu (2022), the CFO survey, and the Livingston survey as dependent variable, respectively. All regressions control for past one-year realized excess returns.

Price-growth extrapolation mechanisms as in Adam et al. (2017) and Jin and Sui (2021) produce pro-cyclical movements in subjective risk premia and hence, for example, a negative relationship between D/P and subjective risk premia. This relationship could be sufficiently weak to still be consistent with a largely acyclical subjective stock market risk premium. Jin and Sui (2021) report that a one standard deviation increase in their sentiment state variable pushes up annualized subjective expected stock market returns by about 0.60 percentage points. This would correspond roughly to a coefficient of -0.60 for CAY or D/P in our regressions in Table 4—which is in the ballpark of what we find for individual investor and CFO expectations in Panels A and B. Of course, most of these point estimates are statistically not significantly different from zero so we cannot say with much confidence whether these coefficients are truly different from zero.

3.2 Treasury bond market

Table 5 turns to the Treasury bond market. The dependent variable in Panel A is an equal-weighted average of one-year excess returns on CRSP Treasury Indexes with maturities of two, five, seven, and ten years. We observe realized returns at monthly frequency, hence as in the earlier stock market return regressions, the dependent variable is measured over overlapping annual windows and we use block-bootstrapped p-values to account for this overlap. In Panel B, the dependent variable is the subjective expected excess returns from the Blue Chip survey of professional forecasters. As in the case of realized returns, we use an equal-weighted average of subjective expected excess returns over two-, five-, seven-, and ten-year maturities.

In Panel A, for realized excess returns, the two asset-class specific predictors LN and CYCLE are statistically significant at conventional levels ($p < 0.01$ for LN; $p < 0.05$ for CYCLE). If used as single predictors, they get point estimates of 1.94 and 2.29, respectively. This means that a one standard deviation move in the predictor is associated with a roughly 2pp movement in objective expected excess returns. Among the cyclical indicators, TERM

Table 5
Regressing Treasury Bond Average Excess Returns on Predictors

In Panel A, dependent variable is the equal-weighted average of one-year excess returns on CRSP Treasury Indexes with maturities of two, five, seven, and ten years. The sample period is from March 1952 to December 2020. In Panel B, dependent variable is the equal-weighted average of one-year-ahead excess return expectations from BCFF with the same set of maturities as in Panel A. The sample period is from January 1988 to December 2020. Column variables are described in Section 2.2. R^e_{past} denotes the past one-year average excess returns. All independent variables are standardized to have unit standard deviations in the full sample period. The first row in each panel reports the OLS estimates multiplied by 100. In Panel A, bootstrap bias-adjusted coefficients are reported in braces; bootstrapped p-values are reported in parentheses. We use a stationary bootstrap with an optimal block length determined as in Politis and White (2004). In Panel B, EWC p-values following Lazarus et al. (2018) are reported in parentheses.

	Asset-Class Specific				Business-Cycle									
	LN		CYCLE		N-IP		TERM		DEFAULT		F1		VIX^2	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
A. Realized														
Coeff	1.94	2.33	2.85	2.85	0.57	0.66	1.50	1.88	1.20	1.26	1.00	1.07	0.31	0.32
{bias-adj.}	{1.94}	{2.32}	{2.29}	{2.38}	{0.75}	{0.85}	{1.48}	{1.92}	{1.36}	{1.35}	{1.08}	{1.13}	{0.30}	{0.33}
(p-value)	(0.00)	(0.00)	(0.01)	(0.02)	(0.61)	(0.60)	(0.00)	(0.00)	(0.12)	(0.13)	(0.04)	(0.07)	(0.22)	(0.20)
R^e_{past}		-0.90		-0.16		-0.13		-0.91		-0.22		-0.22		-0.09
{bias-adj.}		{-0.85}		{-0.37}		{-0.12}		{-0.94}		{-0.13}		{-0.10}		{0.02}
(p-value)		(0.20)		(0.79)		(0.84)		(0.08)		(0.73)		(0.77)		(0.88)
Adj. R^2	0.14	0.16	0.24	0.24	0.00	0.00	0.09	0.12	0.02	0.02	0.03	0.03	0.00	0.00
N	717	717	578	578	826	814	826	814	826	814	717	717	826	814
B. BCFF														
Coeff	0.37	0.50	0.83	0.83	0.54	0.67	0.30	0.33	0.30	0.41	0.57	0.75	0.17	0.22
(p-value)	(0.16)	(0.05)	(0.00)	(0.00)	(0.31)	(0.22)	(0.15)	(0.11)	(0.47)	(0.35)	(0.02)	(0.00)	(0.13)	(0.07)
R^e_{past}		-0.35		0.02		-0.25		-0.22		-0.24		-0.44		-0.28
(p-value)		(0.08)		(0.90)		(0.27)		(0.26)		(0.32)		(0.06)		(0.24)
Adj. R^2	0.03	0.05	0.23	0.22	0.02	0.03	0.03	0.04	0.01	0.02	0.07	0.11	0.02	0.04
N	397	397	397	397	397	397	397	397	397	397	397	397	397	397

and F1 are also significant predictors at conventional levels, and DEFAULT marginally so. The magnitude of expected return variation associated with the cyclical indicators is generally smaller than for LN and CYCLE.

Interestingly, for all of these predictor variables, we again see much smaller coefficient estimates when subjective expected excess returns are the dependent variable in Panel B. Here we obtain the maximum coefficient with CYCLE, but at 0.83 the magnitude is still much smaller than in Panel A (2.29). Including excess returns over the past year as an additional predictor variable generally has very little effect, both in Panel A and B.

Figure 3 summarizes the results, again focusing on coefficient estimates in regressions with past returns included as controls. The figure shows very clearly the different magnitudes of the slope coefficients in the regressions with realized excess returns, shown as blue bars, and regressions with subjective expected excess returns, shown as red bars. Thus, in the Treasury bond market, too, subjective risk premia are much less sensitive to cyclical predictor variables than objective risk premia are.

3.3 Foreign exchange market

Table 6 presents results for foreign exchange. The dependent variable in Panel A is the realized excess return on a portfolio of developed market currencies from the viewpoint of a US-based investor. The asset-class specific predictor is the average forward discount on the developed market currencies as in Lustig et al. (2014).

In Panel A, all predictors except TERM produce statistically significant or at least marginally significant coefficient estimates. The magnitudes are substantial. The one standard deviation change in the average forward discount is associated with a change of 3.75pp in predicted realized excess returns. Among the cyclical indicators, the strongest predictor is N-IP. A one standard deviation move in N-P is associated with a change of 6.66pp in the objective expected excess returns on a US investors' currency portfolio.

The estimates in Panel B for subjective expected excess returns show a sharp contrast to

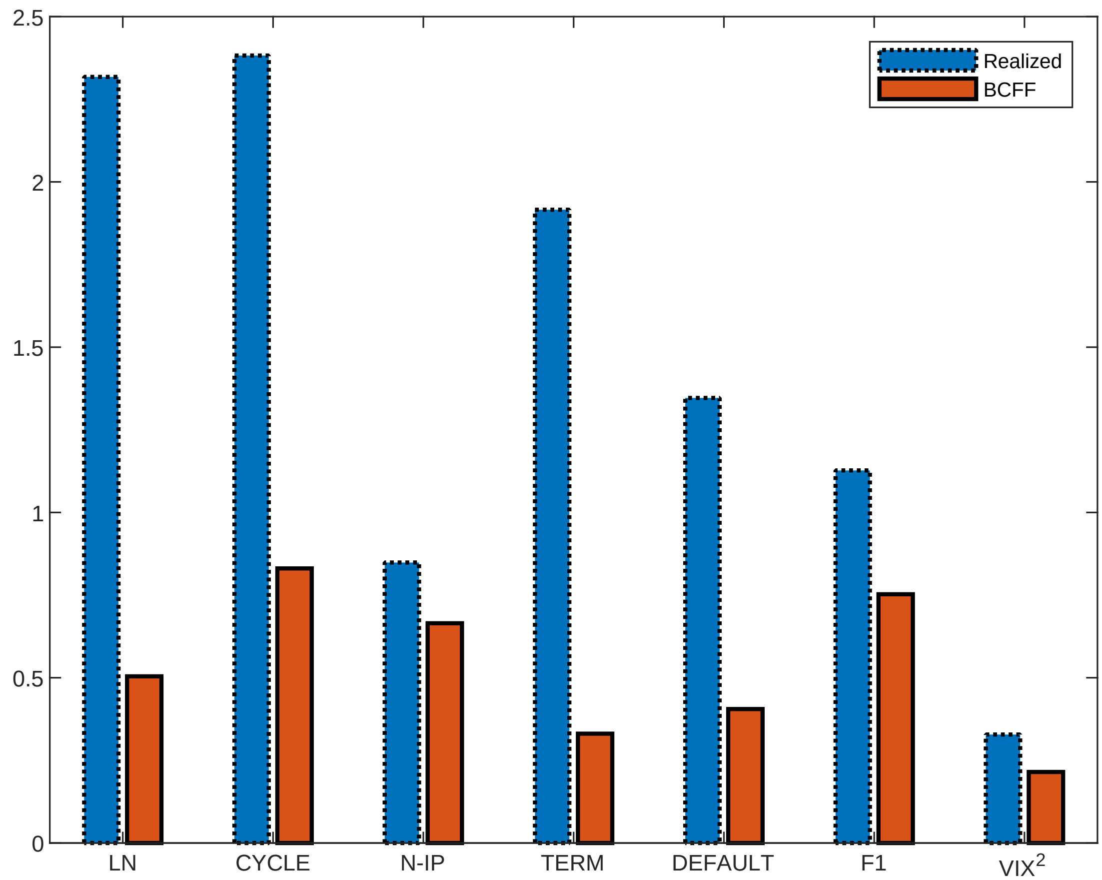

FIGURE 3
Coefficients from Regressing Treasury Bond Average Excess Returns on Predictors

The blue bars plot the slope coefficients from regressing the equal-weighted average of realized excess returns on CRSP Treasury Indexes on asset-class specific predictors and business-cycle indicators. The maturities include two, five, seven, and ten years. The coefficients are bias-adjusted using bootstrap. The red bars use the equal-weighted average of survey excess return expectations with the same set of maturities from BCFF as dependent variable. All regressions control for past one-year realized average excess returns.

Panel A. Point estimates are generally much smaller and none of the coefficient estimates is significantly different from zero according to conventional levels.

As earlier in the case of bond excess returns, controlling for the past one-year return on the currency portfolio does not have much effect on the point estimates of the other predictors, neither for realized excess returns, nor for subjective expected excess returns.

Figure 4 visualizes the magnitude of the coefficients, focusing on the regressions that include past returns along with the predictor variable. The figure clearly shows the sharp contrast between the large coefficients obtained in realized excess return prediction regressions and the extremely small coefficients in the subjective expected excess return regressions. There is therefore little support for the notion that the objective return premia obtained in ex-post regressions with realized excess returns reflect ex-ante subjective risk premia from the viewpoint of the professional forecasters.

3.4 Commodity futures market

In our analysis of commodity futures risk premia, we look at the energy and metals sectors separately. The economic risks that the underlying commodities are exposed to could be quite different and the dynamics of risk premia could therefore be potentially different as well. Following Hong and Yogo (2012), the asset-class specific predictors for commodity futures are the futures basis and past one-year open interest growth.

3.4.1 Metals

The dependent variable in Panel A of Table 7 is the average realized one-year buy-and-hold excess return on a portfolio of metals futures. As shown in columns (1) to (4), the futures basis and open interest growth are economically strong predictors, with one standard deviation change in the predictor associated with around 7pp change in the predicted excess returns. However, the commodity futures excess returns are so noisy and the sample sufficiently short that these estimates are not statistically significantly different from zero at conventional levels

Table 6
Regressing Foreign Exchange Average Excess Returns on Predictors

In Panel A, dependent variable is the average realized one-year excess returns on the basket of developed currencies. The sample period is from October 1984 to June 2021. In Panel B, dependent variable is the one-year-ahead average excess return expectations on the same basket of developed currencies from CE and FX4casts. The sample period is from August 1986 to June 2021. Column variables are described in Section 2.2. R^e_{past} denotes the past one-year excess returns. All independent variables are standardized to have unit standard deviations in the full sample period. The first row in each panel reports the OLS estimates multiplied by 100. In Panel A, bootstrap bias-adjusted coefficients are reported in braces; bootstrapped p-values are reported in parentheses. We use a stationary bootstrap with an optimal block length determined as in Politis and White (2004). In Panel B, EWC p-values following Lazarus et al. (2018) are reported in parentheses.

	Asset-Class Specific		Business-Cycle									
	FD		N-IP		TERM		DEFAULT		F1		VIX^2	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A. Realized												
Coeff	3.83	4.25	7.43	6.74	1.84	2.06	5.75	6.79	2.51	1.94	1.24	1.29
{bias-adj.}	{3.75}	{3.79}	{6.66}	{5.55}	{1.99}	{2.24}	{5.99}	{7.05}	{2.65}	{2.50}	{1.44}	{1.24}
(p-value)	(0.03)	(0.08)	(0.04)	(0.10)	(0.27)	(0.29)	(0.07)	(0.07)	(0.06)	(0.14)	(0.03)	(0.04)
R^e_{past}		-1.27		1.00		0.07		0.65		0.82		0.89
{bias-adj.}		{-1.23}		{1.03}		{-0.12}		{0.50}		{0.76}		{0.75}
(p-value)		(0.42)		(0.62)		(0.97)		(0.67)		(0.69)		(0.66)
Adj. R^2	0.15	0.13	0.10	0.07	0.03	0.04	0.11	0.15	0.04	0.02	0.03	0.04
N	442	430	442	430	442	430	442	430	442	430	442	430
B. CE & FX4casts												
Coeff	-0.19	0.01	0.27	0.10	-0.68	-0.62	0.27	0.23	0.06	-0.03	0.56	0.53
(p-value)	(0.75)	(0.99)	(0.74)	(0.91)	(0.24)	(0.25)	(0.69)	(0.77)	(0.89)	(0.94)	(0.10)	(0.13)
R^e_{past}		-0.41		-0.40		-0.20		-0.40		-0.43		-0.25
(p-value)		(0.54)		(0.57)		(0.76)		(0.57)		(0.54)		(0.71)
Adj. R^2	0.00	0.01	-0.00	0.01	0.04	0.04	-0.00	0.01	-0.00	0.01	0.07	0.07
N	419	419	419	419	419	419	419	419	414	414	419	419

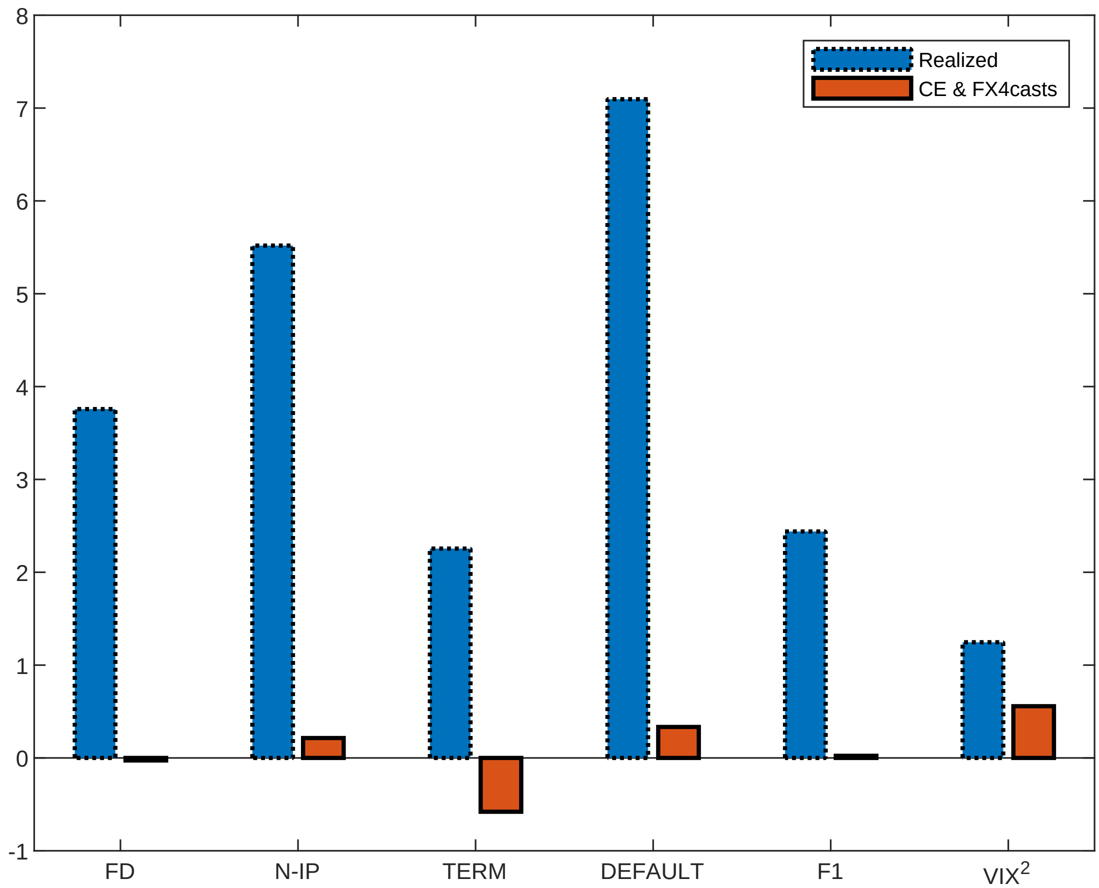

FIGURE 4

Coefficients from Regressing Foreign Exchange Average Excess Returns on Predictors

The blue bars the slope coefficients from regressing average realized one-year excess returns of developed currencies on asset-class specific predictors and business-cycle indicators. The coefficients are bootstrap bias-adjusted. The red bars use average survey excess return expectations on the same basket of currencies from CE and FX4casts as dependent variable. All regressions control for past one-year average realized excess returns.

despite the large magnitude of the point estimates. With a p-value of 0.11, open interest growth, when used as a single predictor, is at least close to marginally significant. Among the cyclical indicators, only VIX^2 is marginally significant on its own ($p < 0.10$).

Panel B shows the corresponding regressions for subjective expected excess returns. Looking across all specifications, we again see the general pattern that the magnitudes of the slope coefficient estimates for the predictors are almost all lower than in Panel A (and often of different sign). Figure 5 illustrates this graphically.

Interestingly, Panel B also shows that subjective risk premia of professional forecasters in the metals futures market seem to be strongly contrarian with respect to recent past excess returns on metals futures. A fairly large share of time-variation in subjective risk premia can be traced to this contrarian effect, even though the relationship of past returns to objective risk premia is, if anything, positive in Panel A.

3.4.2 The crude oil market

Table 8 presents the results for futures excess returns in the WTI crude oil market. The results for realized excess return prediction in Panel A show that there are no statistically significant predictors of excess returns. Some of the point estimates are large, though. For example, the point estimate for N-IP implies that a one standard deviation fall in industrial production growth is associated with a rise of predicted excess returns of 10.26pp. But the p-value of 0.21 indicates that there is a very large degree of uncertainty about the true magnitude of the effect.

Panel B presents the result for subjective expected excess returns. As in the case of metals futures, subjective risk premia are strongly contrarian with respect to recent past excess returns. Other than this past returns effect, there is little else in terms of a clear relationship to predictors. The basis in column (1) shows up with a statistically significant relationship on its own, but this goes away when past returns are controlled for. The DEFAULT variable in column (10) is a statistically significant predictor only when past returns are controlled

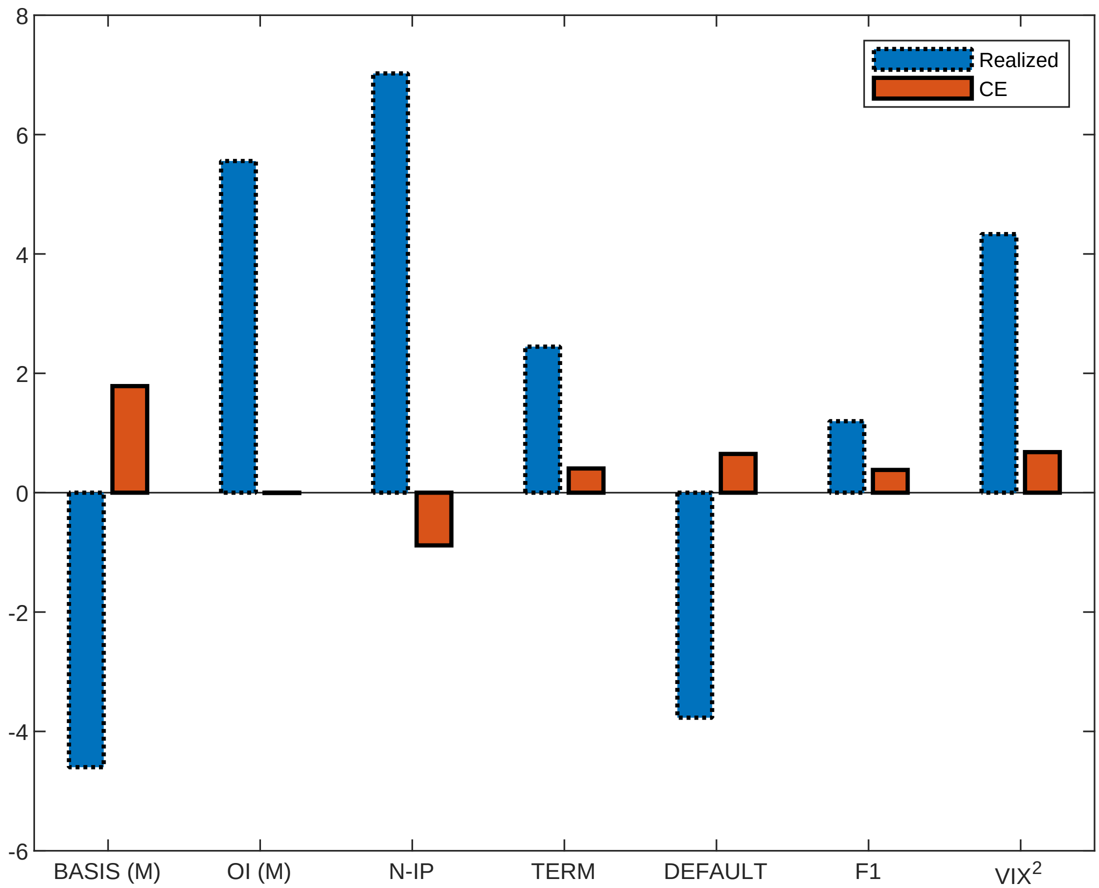

FIGURE 5

Coefficients from Regressing Metals Futures Excess Returns on Predictors

The blue bars plot the slope coefficients from regressing average realized one-year buy-and-hold excess returns of metals futures on asset-class specific predictors and business-cycle indicators. The coefficients are bootstrap bias-adjusted. The red bars use average survey excess return expectations from CE as dependent variable. All regressions control for past one-year average realized excess returns.

TABLE 7

Regressing Metals Futures Excess Returns on Predictors

In Panel A, dependent variable is the average realized one-year buy-and-hold excess returns on metals. The sample period is from September 1978 to June 2021. In Panel B, dependent variable is the one-year-ahead average excess return expectations from CE. The sample period is from August 1995 to June 2021. Column variables are described in Section 2.2. R^e_{past} denotes the past one-year excess returns. All independent variables are standardized to have unit standard deviations in the full sample period. In Panel A, bootstrap bias-adjusted coefficients are reported in braces; bootstrapped p-values are reported in parentheses. We use a stationary bootstrap with an optimal block length determined as in Politis and White (2004). In Panel B, EWC p-values following Lazarus et al. (2018) are reported in parentheses.

	Asset-Class Specific				Business-Cycle									
	BASIS		OI		N-IP		TERM		DEFAULT		F1		VIX^2	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
A. Realized														
Coeff	-6.57	-6.69	6.38	5.72	5.84	7.80	1.51	1.98	-3.51	-1.61	-1.18	0.25	3.11	3.83
{bias-adj.}	{-5.86}	{-4.57}	{7.45}	{5.51}	{5.91}	{7.03}	{2.33}	{2.48}	{-5.19}	{-3.78}	{-0.88}	{1.24}	{3.80}	{4.31}
(p-value)	(0.20)	(0.31)	(0.11)	(0.19)	(0.55)	(0.43)	(0.75)	(0.64)	(0.69)	(0.86)	(0.80)	(0.96)	(0.06)	(0.03)
R^e_{past}		3.45		1.27		4.68		4.70		4.20		4.50		5.31
{bias-adj.}		{4.50}		{3.61}		{5.68}		{5.98}		{4.58}		{5.51}		{6.10}
(p-value)		(0.34)		(0.81)		(0.27)		(0.20)		(0.33)		(0.27)		(0.13)
Adj. R^2	0.06	0.09	0.09	0.09	0.01	0.04	0.00	0.03	0.00	0.02	-0.00	0.02	0.03	0.07
N	516	504	441	441	516	504	516	504	516	504	516	504	516	504
B. CE														
Coeff	7.27	1.79	-2.22	-0.01	-0.05	-0.89	0.47	0.41	2.28	0.65	1.24	0.38	1.09	0.68
(p-value)	(0.02)	(0.43)	(0.01)	(0.99)	(0.97)	(0.52)	(0.52)	(0.62)	(0.28)	(0.71)	(0.32)	(0.72)	(0.05)	(0.20)
R^e_{past}		-3.83		-4.12		-4.19		-4.12		-4.01		-3.83		-3.86
(p-value)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)
Adj. R^2	0.12	0.35	0.13	0.34	-0.01	0.35	-0.00	0.35	0.04	0.35	0.02	0.30	0.08	0.38
N	150	150	150	150	150	150	150	150	150	150	145	145	150	150

for. Figure 6 shows that the overall picture of absolute coefficient magnitudes in the two panels is mixed, with no clear difference in magnitudes in the objective and subjective risk premia specifications. This is to be expected when none of the variables really has a strong relationship to objective or subjective risk premia.

3.5 Summary measures, variations, and statistical inference

We now construct summary measures that combine our findings across predictors and asset classes. This also allows us to compactly report the sensitivity of our results to a number of robustness checks. For each survey, we first calculate the following two ratios:

$$M_1 \equiv \frac{\sum_{k\in\mathbb{K}} \operatorname{sgn}\left(\beta_k^{Adj}\right)\cdot\beta_k^{Sub}}{\sum_{k\in\mathbb{K}} \operatorname{sgn}\left(\beta_k^{Adj}\right)\cdot\beta_k^{Adj}}, \tag{7}$$

and

$$M_2 \equiv \frac{\sum_{k\in\mathbb{K}} |\beta_k^{Sub}|}{\sum_{k\in\mathbb{K}} |\beta_k^{Adj}|}, \tag{8}$$

where β_k^{Sub} is the slope coefficient on asset-class specific predictor or business-cycle indicator k in the regression with survey excess return expectations, β_k^{Adj} is the bias-adjusted coefficient in the regression with realized excess returns, and $\mathbb{K}$ is the group of all asset-class specific predictors and cyclical indicators combined.

Because the sign of predictors can be arbitrary, we need to first align the signs before we can average coefficients across predictors. Multiplying with $\operatorname{sgn}\left(\beta_k^{Adj}\right)$ in the numerator and denominator in (7) is like switching the signs of the predictor variables such that the coefficients in the realized return regressions are all positive, and hence the denominator of M_1 is always positive. Then, if subjective risk premia have the same variation with respect to predictors as objective risk premia do, M_1 will be about 1; if subjective risk premia have the exact opposite variation, M_1 will be about -1.

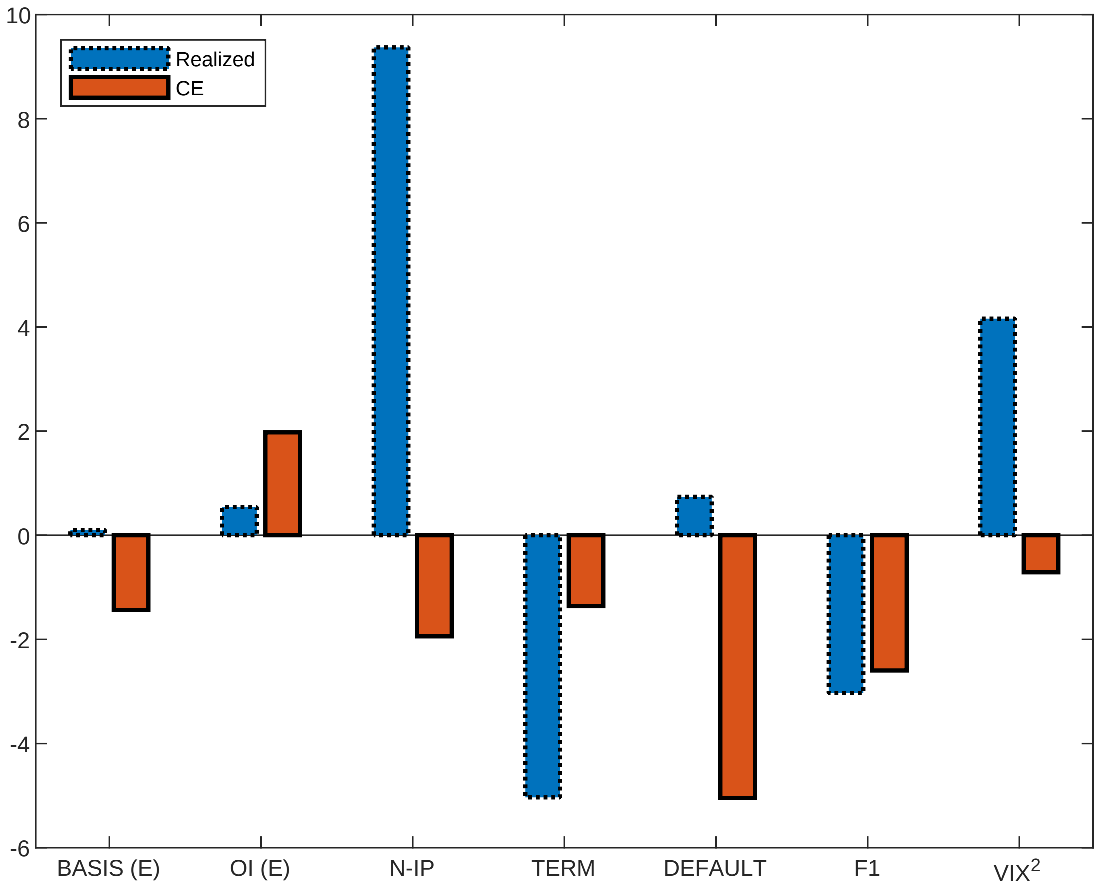

FIGURE 6
Coefficients from Regressing Crude Oil Futures Excess Returns on Predictors

The blue bars plot the slope coefficients from regressing realized one-year buy-and-hold excess returns of WTI crude oil futures on asset-class specific predictors and business-cycle indicators. The coefficients are bootstrap bias-adjusted. The red bars use survey excess return expectations from CE as dependent variable. All regressions control for past one-year realized excess returns.

TABLE 8

Regressing Crude Oil Futures Excess Returns on Predictors

In Panel A, dependent variable is the realized one-year buy-and-hold excess returns on the WTI crude oil futures. The sample period is from December 1984 to June 2021. In Panel B, dependent variable is the one-year-ahead buy-and-hold excess return expectations from CE. The sample period is from August 1995 to June 2021. Column variables are described in Section 2.2. R^e_{past} denotes the past one-year excess returns. All independent variables are standardized to have unit standard deviations in the full sample period. In Panel A, bootstrap bias-adjusted coefficients are reported in braces; bootstrapped p-values are reported in parentheses. We use a stationary bootstrap with an optimal block length determined as in Politis and White (2004). In Panel B, EWC p-values following Lazarus et al. (2018) are reported in parentheses.

	Asset-Class Specific				Business-Cycle									
	BASIS		OI		N-IP		TERM		DEFAULT		F1		VIX^2	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
A. Realized														
Coeff	1.20	-0.71	0.36	5.14	9.89	7.82	-4.25	-4.33	2.43	2.03	-0.63	-2.35	3.88	3.64
{bias-adj.}	{1.74}	{0.24}	{-0.71}	{0.45}	{10.26}	{9.32}	{-4.46}	{-5.16}	{0.58}	{0.66}	{-0.47}	{-2.98}	{4.05}	{4.18}
(p-value)	(0.81)	(0.90)	(0.95)	(0.58)	(0.21)	(0.48)	(0.53)	(0.57)	(0.71)	(0.81)	(0.84)	(0.51)	(0.25)	(0.33)
R^e_{past}		-3.30		-6.39		-1.78		-3.10		-2.53		-3.34		-2.03
{bias-adj.}		{-1.38}		{-2.59}		{-0.66}		{-2.44}		{-2.19}		{-2.52}		{-0.51}
(p-value)		(0.67)		(0.55)		(0.77)		(0.57)		(0.67)		(0.55)		(0.75)
Adj. R^2	-0.00	-0.00	-0.00	0.00	0.01	0.01	0.01	0.02	-0.00	0.00	-0.00	0.00	0.03	0.02
N	439	427	426	426	439	427	439	427	439	427	439	427	439	427
B. CE														
Coeff	2.46	-1.43	-3.61	1.98	3.82	-1.94	-1.73	-1.36	0.19	-5.05	0.22	-2.60	0.52	-0.71
(p-value)	(0.04)	(0.36)	(0.13)	(0.38)	(0.23)	(0.55)	(0.35)	(0.41)	(0.95)	(0.00)	(0.92)	(0.19)	(0.52)	(0.29)
R^e_{past}		-6.47		-7.14		-5.93		-5.41		-6.60		-7.24		-5.83
(p-value)		(0.01)		(0.00)		(0.01)		(0.00)		(0.00)		(0.00)		(0.01)
Adj. R^2	0.07	0.33	0.14	0.33	0.03	0.32	0.01	0.32	-0.01	0.38	-0.01	0.43	0.00	0.32
N	150	150	150	150	150	150	150	150	150	150	145	145	150	150

It may also be of interest to simply compare the volatility of the predictor-related components in subjective and objective risk premia, without paying attention to the direction in which a predictor moves these premia. This is captured by ratio M_2. This ratio uses the absolute values of coefficients.

The final summary measures are calculated by first taking the average of M_1 and M_2 within asset classes (if there are multiple surveys) then across asset classes.[10] Figure 7 presents the results.

To explore the robustness of our conclusions, we report results from several different specifications. First, we use coefficient point estimates obtained with (*Control*) or without (*Single*) controlling for past returns. Second, we re-run all of the analysis with a different definition of the realized excess return sample. In our baseline analysis above, we used the maximum sample length that was available given data availability constraints on the predictors and returns data. With this baseline approach, one might worry that the law of motion of objective risk premia has undergone structural change and may be different in the samples for which we have survey data coverage, which are often much shorter than the samples for which we have realized returns and predictor data. For this reason, we re-run all realized excess regressions with the realized returns sample restricted to the time periods for which we have survey data to construct subjective expected excess returns. In Figure 7, the results from our baseline analysis with the full sample are denoted with *(F)*, those with the return sample matched to the survey sample are denoted with *(S)*.

The four pairs of bars on the left-hand side of the figure show the results. The blue bars show the summary measure of the M_1 ratios, the red bars show the summary of the M_2 ratios. The baseline case in the left-most blue bar shows that a one percentage point move in the objective risk premium is associated, on average, only with 0.12 percentage point move in the same direction of the subjective risk premium. Based on the comparison of the first two blue bars, including or not including past returns has very little effect on the ratio of

[10] We treat metals and crude oil as in the same asset class. Treating them as separate asset classes does not quantitatively change the results.

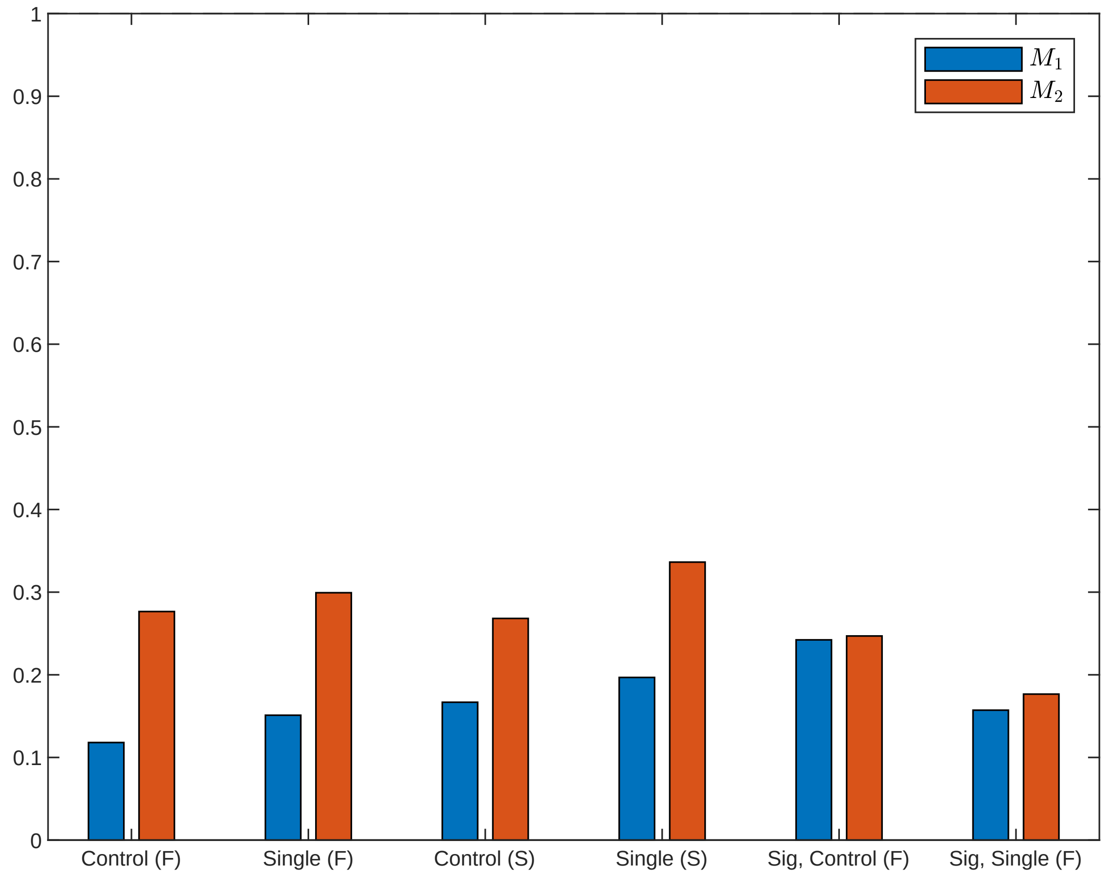

FIGURE 7

Comparing Magnitudes of Coefficients from Regressing Survey Expectations and Realized Returns

For each blue and red bar, we first calculate the following ratios for each survey:

$$M_1 \equiv \frac{\sum_{k \in \mathbb{K}} \operatorname{sgn}\left(\beta_k^{Adj}\right) \cdot \beta_k^{Sub}}{\sum_{k \in \mathbb{K}} \operatorname{sgn}\left(\beta_k^{Adj}\right) \cdot \beta_k^{Adj}}, \quad M_2 \equiv \frac{\sum_{k \in \mathbb{K}} |\beta_k^{Sub}|}{\sum_{k \in \mathbb{K}} |\beta_k^{Adj}|},$$

respectively, where β_k^{Sub} is the slope coefficient on asset-class specific predictor or business-cycle indicator k in the regression with survey excess return expectations, β_k^{Adj} is the bias-adjusted coefficient in the regression with realized excess returns, and $\mathbb{K}$ is the set of all predictors. We then take the average of M_1 and M_2 within each asset class if there are multiple surveys and finally across asset classes. *(F)* denotes estimates using the full return sample; *(S)* denotes estimates using the matched-to-survey return sample; *Control* and *Single* denote estimates with and without controlling for past excess returns, respectively; *Sig* denotes estimates only using predictors whose full-sample p-values are smaller than 0.05.

coefficients for objective and subjective premia. In both cases, we get a ratio of close to 0.15. Using the full sample for realized returns (first and second blue bar) or a shorter sample matched to the survey data set (third and fourth blue bar) does not have much effect on the results either, with a slight increase of about 0.05 in both *Control* and *Single* cases. Overall, subjective risk premia vary substantially less with the predictor variables than objective risk premia.

The red bars shows that the picture is quite similar if we ignore potential directional mismatch between the movements of subjective and objective risk premia and just look at the relative volatility of predictor-related components using the ratio M_2. Even then, the ratio is far below one, indicating a much smaller volatility of the cyclical-variable component in subjective risk premia. Neither controlling for past returns nor using matched-to-survey samples has much effect on the summary measure of the M_2 ratios.

One potential concern about the analysis so far is that the set of predictors might include ones that are simply irrelevant for risk premia. If the predictor variable represents just irrelevant noise that is not truly related to risk premia, neither objective nor subjective, the true slope coefficients are zero. Due to estimation error, the slope coefficient estimates will be non-zero and they could have similar absolute magnitudes in the regressions with realized returns and subjective expectations. This would bias the M_2 ratio toward one and the M_1 ratio toward zero.

For this reason, the last two pairs of bars on the right-hand side include, within each asset class, only on predictor variables that are significant at conventional levels ($p \leq 0.05$) in predicting realized excess returns using the full available sample of realized return and predictor data. As the figure shows, for these variables, the relative magnitude of the subjective expected excess return regression coefficients is still small. Depending on whether we include past returns in the regression or not, we obtain ratios of slightly above or below 0.20 for both M_1 and M_2 ratios.

So far we have focused on point estimates in our summary measures. However, it would also be useful to assess the joint statistical significance of the cyclicality wedges between subjective and objective risk premium dynamics. We develop an asymptotic inference approach in Appendix C. This is not straightforward as the regressions involve different predictors, sample periods, data sets, and measurement frequencies. Our approach treats all these regressions jointly as one big system of regression equations. We then construct a composite estimator that aggregates the differences between realized-return-based and survey-based coefficient estimates as

$$d = \boldsymbol{e}_2' \boldsymbol{B} \boldsymbol{w}, \tag{9}$$

where the $K \times J$ matrix $\boldsymbol{B}$ horizontally concatenates the regression coefficient row vectors from J realized-return-based and survey-based regressions. The weight column vector $\boldsymbol{w}$ is such that the elements that multiply coefficients from the predictive regressions with excess returns as dependent variable are positive, of equal magnitude, and sum to 1 across all regressions within an asset class, while those that multiply coefficients from the survey data regressions are negative, of equal magnitude, and sum to -1 within an asset class; $\boldsymbol{e}_2$ is a vector with one as the second element and zeros otherwise to pick out the aggregated slope coefficient differences corresponding to asset-class specific predictors and business-cycle indicators. Thus, d represents the average wedge between realized excess returns and survey expectations associated with a one-standard-deviation change in the predictor variable. For the purposes of statistical inference, formulating this wedge as a difference is more tractable than the ratio that we used in our earlier summary measures in Figure 7. Since signs of regression coefficients matter for d, we flip the signs of predictors if necessary such that they forecast future realized excess returns positively, as we did for the ratio M_1 earlier.

The top panel in Table 9 reports the results for individual asset classes. Focusing on the full-sample estimates in the first two columns, the point estimate for the aggregate wedge d is between two to four times as big as the standard error, and hence statistically significant at conventional levels, for all asset classes except metals with past return control and oil.

To further aggregate across asset classes, we construct two composite measures. The first one equally weights all asset classes. The first two columns of the bottom panel in Table 9 show that this composite cyclicality wedge is highly statistically significant with the point estimate more than five times as big as the standard error. To prevent asset classes with high return volatility from dominating, we construct a second composite measure that weights asset classes by the inverse of their realized excess return volatility. This point estimate for this second composite measure is smaller as it puts a higher weight on Treasury bonds, but the standard errors shrink to a similar degree, so that it remains highly statistically significant.

As a robustness check, we also repeat the analysis using the matched-to-survey sample. All results hold except that the cyclicality wedge narrows for Treasury bonds but widens for metals futures.

4 Comparison with Out-of-Sample Forecasts

Our analysis so far compared the cyclicality of subjective expected excess returns with the cyclicality of forecasts implied by estimates of in-sample predictive regressions. For evaluating RE models, this is the appropriate comparison. In RE models, agents already know the underlying data-generating process, including the values of its parameters. In-sample predictive regressions are then the most efficient way for an econometrician to estimate what economic agents already know. Yet, the empirical evidence on the large gap in cyclicality between subjective expectations and in-sample predictive regression forecasts is hard to square with RE models. This brings up the question whether non-RE belief-formation mechanisms that do not endow agents with so much knowledge of the data-generating process could explain the lack of cyclicality in subjective expected excess returns.

One natural possibility is that, unlike in RE models, investors do not know the parameters of the data-generating process. Instead, they learn about these parameters from observed data in real time. If so, an OOS forecast, not the fitted value from an in-sample regression run ex post over the whole sample, should be close to their observed subjective expectations. For this reason, we now examine whether OOS forecasts of excess returns are closer to the

Table 9
Joint Statistical Significance of the Cyclicality Wedges

This table reports the point estimates and standard errors of the following measure:

$$d = \boldsymbol{e}_2'\boldsymbol{B}\boldsymbol{w},$$

where the $K \times J$ matrix $\boldsymbol{B}$ horizontally concatenates the regression coefficient row vectors from J realized-return-based and survey-based regressions. The weight column vector $\boldsymbol{w}$ is such that the elements that multiply coefficients from the predictive regressions with excess returns as dependent variable are positive, of equal magnitude, and sum to 1 across all regressions within an asset class, while those that multiply coefficients from the survey data regressions are negative and sum to -1 within an asset class; $\boldsymbol{e}_2$ is a vector with one as the second element and zeros otherwise to pick out the aggregated slope coefficient differences corresponding to asset-class specific predictors and business cycle indicators. The composite measure either equally weights the four asset classes or weights them in inverse proportion to the volatility of their realized excess returns σ. We multiply the point estimates d by 100. Standard errors calculated using the EWC estimator from Lazarus et al. (2018) are reported in parentheses (also multiplied by 100).

	Full-Sample		Survey-Sample	
	w/o R^e_{past}	with R^e_{past}	w/o R^e_{past}	with R^e_{past}
Stock Market	3.56	3.55	3.37	3.09
	(0.96)	(1.01)	(1.16)	(1.36)
Treasury Bonds	0.90	0.95	0.40	0.62
	(0.35)	(0.40)	(0.42)	(0.45)
Foreign Exchange	3.67	3.76	3.29	3.55
	(0.92)	(1.00)	(0.93)	(0.97)
Commodities				
Metals	4.97	4.25	10.43	10.70
	(2.27)	(2.40)	(2.22)	(2.46)
Oil	2.54	3.97	3.68	5.95
	(1.86)	(2.64)	(2.49)	(3.45)
Composite				
Equal-weighted	2.97	3.09	3.53	3.90
	(0.53)	(0.61)	(0.66)	(0.67)
σ^{-1}-weighted	2.23	2.31	2.04	2.29
	(0.40)	(0.43)	(0.44)	(0.44)

subjective excess return expectations in terms of cyclicality than the in-sample excess return predictions.

We first construct OOS forecasts at time t based on each individual predictor x_i by performing the following predictive regressions over expanding windows:

$$R^e_k = \alpha_{i,t} + \beta_{i,t} x_{i,k-1} + \varepsilon_{i,k}, \quad k = 2, \dots, t. \tag{10}$$

Given the negligible role of past excess returns in predicting future excess returns in our earlier analyses, we focus here on specifications with just a single predictor without including past one-year excess returns. We assume that investors estimate $\alpha_{i,t}$ and $\beta_{i,t}$ as

$$(\hat{\alpha}_{i,t},\ \hat{\beta}_{i,t}) \equiv \underset{\alpha,\ \beta}{\operatorname{argmin}} \sum_{k=2}^{t} \omega_{t,k} (R^e_k - \alpha - \beta x_{i,k-1})^2. \tag{11}$$

Our first set of analyses sets $\omega_{t,k} = 1$ and hence $\hat{\alpha}_{i,t}$ and $\hat{\beta}_{i,t}$ are simply OLS estimates from recursively expanding windows.

An OOS forecast of $t+1$ excess returns based on $\hat{\alpha}_{i,t}$ and $\hat{\beta}_{i,t}$ would, however, impose extreme confidence that the variable x_i is truly a predictor of excess returns. In practice, investors cannot be sure. If it's not, then a forecast that constrains $\beta_{i,t}$ to zero could perform better. For this reason, we assume that investors evaluate, based on the historical data that has accumulated until time t, whether the predictive regression forecast actually adds value, in terms of OOS forecasting, relative to simply setting the forecast equal to the historical mean of R^e_k

$$\bar{R}^e_t = \frac{\sum_{k=2}^{t} \omega_{t,k} R^e_k}{\sum_{k=2}^{t} \omega_{t,k}}. \tag{12}$$

More precisely, we assume that investors check in past data until period t which combination of $\bar{R}^e_t$ and the predictive regression forecast $\hat{\alpha}_{i,t} + \hat{\beta}_{i,t} x_{i,t}$ produced the best OOS forecasts.

Investors' forecast is then a weighted average:

$$\hat{\mathbb{E}}_{i,t} R^e_{t+1} \equiv s_{i,t}(\hat{\alpha}_{i,t} + \hat{\beta}_{i,t} x_{i,t}) + (1 - s_{i,t})\bar{R}^e_t, \tag{13}$$

where the weight $s_{i,t}$ is chosen to maximize the historical OOS performance starting at t_{min} after a burn-in period

$$s_{i,t} = \underset{s}{\operatorname{argmin}} \sum_{k=t_{min}}^{t} \omega_{t,k} u_k(s)^2, \tag{14}$$

$$u_k(s) \equiv R^e_k - s(\hat{\alpha}_{i,k-1} + \hat{\beta}_{i,k-1} x_{i,k-1}) - (1 - s)\bar{R}^e_{k-1}. \tag{15}$$

In other words, we let investors look for the optimal degree of shrinkage of the predictive regression forecast toward the historical mean.

In our empirical implementation, we choose t_{min} of 10 years as burn-in period. For foreign exchange and commodity futures, we do not have enough data after the burn-in periods to have much statistical power. For this reason, we focus on stock and Treasury bond markets only in this analysis. For our cyclicality tests, we then use the OOS forecast series starting from the dates when survey data becomes available (June 1952 for the stock market and January 1988 for the Treasury bond market). We regress the OOS forecasts $\hat{\mathbb{E}}_{i,t} R^e_{t+1}$ on $x_{i,t}$, just as we regress the subjective expected excess returns $\tilde{\mathbb{E}}_t R^e_{t+1}$ from the survey data on $x_{i,t}$.

Both of the asset-specific predictors for Treasury bond excess returns use look-ahead information in their construction that is not available to investors in real time: the bond excess return predictor of Ludvigson and Ng (2009) is constructed as a linear combination of macro factors with weights chosen to maximize in-sample bond excess return predictability subject to a model complexity penalty; the cycle factor of Cieslak and Povala (2015) is constructed by projecting in-sample bond yields on trend inflation and then bond excess returns on the residuals to construct the cycle factor. For the purpose of this OOS exercise, we therefore create real-time versions of these factors.[11] Appendix D.1 describes the details of their construction.

To parsimoniously present the results from this large number of regressions, we again use the ratios introduced in (7) and (8), but now with numerators replaced with coefficients from using OOS forecasts:

$$M_1^* \equiv \frac{\sum_{k\in\mathbb{K}} \operatorname{sgn}\left(\beta_k^{Adj}\right)\cdot\beta_k^{OOS}}{\sum_{k\in\mathbb{K}} \operatorname{sgn}\left(\beta_k^{Adj}\right)\cdot\beta_k^{Adj}}, \quad M_2^* \equiv \frac{\sum_{k\in\mathbb{K}} |\beta_k^{OOS}|}{\sum_{k\in\mathbb{K}} |\beta_k^{Adj}|}. \tag{16}$$

We report the detailed individual coefficients from these regressions in Appendix D.2. If the OOS forecasts constructed in our exercise get close, in terms of cyclicality, to the subjective expected excess returns observed in the survey data, then we should see these ratios to be similar to the ones we reported in Figure 7 with coefficients from the subjective expected excess return regressions in the numerators.

As a baseline for comparison, the left-most blue bars in the two panels of Figure 8 therefore repeat the earlier estimates from Figure 7 in the *Single (F)* specification, i.e., without controlling for past excess returns, where the ratios were 0.16 for the stock market and 0.33 for the Treasury bond market. The neighboring blue bars to the right, labeled *Shrink (E)*, show the ratio (16) with coefficients from using OOS forecasts in the numerators. As the figure shows, that OOS forecasts differ from in-sample forecasts goes some way of explaining the cyclicality gap. If OOS forecasts were as cyclical as in-sample forecasts, the ratio would be about 1.0; if the OOS forecasts were as acyclical as subjective expected excess returns in the survey data, the ratio would be about equal to the ratio represented by the baseline bar. The actual estimate is somewhere inbetween, with 0.46 for the stock market and 0.51 for the Treasury bond market. Using absolute values of coefficients in calculating the measure has

[11] The equity return predictor CAY and the business cycle variable F1 also use look-ahead information in their construction. For CAY, this involves estimation of a cointegration relationship over the full sample ex post; for F1 the estimation of principal components using full sample information. Therefore, these variables may also be contaminated with information that was not available to investors in real time, but the concern is perhaps less severe for CAY and F1 as their construction does not directly use realized excess returns.

little effect, as can be seen from the red bars in the first two pairs.

That investors make use of the full historical record of asset return data in constructing their forecasts may overstate their reliance on historical data from the distant past. In Nagel and Xu (2022) we argued that a reasonable and empirically plausible way of representing investors' learning is to let memory of historical data slowly fade over time as observations recede into the distant past. This could reflect actually fading memory or investors' belief that parameter drift renders data from the distant past irrelevant for forecasting. In line with the learning model in Nagel and Xu (2022), we therefore construct an alternative series of OOS forecasts where investors downweight past data with exponential weights $\omega_{t,k} = \lambda^{t-k}$. We use the same value of $\lambda = 0.982$ as in Nagel and Xu (2022) for quarterly data (and its third root for monthly data). We then repeat our earlier exercise of regressing the OOS forecasts on each individual predictor and we compute the ratio in (16). The result is shown in Figure 8 in the third pair of bars from the left, labeled *Shrink (W)*. Especially for the stock market, where the M_1^* ratio drops by about half to 0.25, the exponential weighting approach brings the ratio substantially closer to the baseline column. For the Treasury bond market it also drops, albeit to a lesser extent, from 0.43 to 0.33. The red bars show a similar pattern.

Another potentially relevant consideration is that investors ex-ante may not know which one of the candidate predictor variables really is a useful OOS predictor. Rapach et al. (2010) show that combination forecasts constructed as average individual-predictor forecasts provide superior OOS performance. Such combination forecasts may therefore be an alternative way of representing how investors construct excess return predictions. We therefore use the individual-predictor OOS forecasts $\hat{\alpha}_{i,t} + \hat{\beta}_{i,t}x_{i,t}$ to construct combination OOS forecasts as a simple equal-weighted average across the N predictors:

$$\hat{E}_t^C R_{t+1}^e = \frac{1}{N}\sum_{i=1}^{N}(\hat{\alpha}_{i,t} + \hat{\beta}_{i,t}x_{i,t}). \tag{17}$$

We then regress these combination forecasts on the individual predictors, one at a time, to

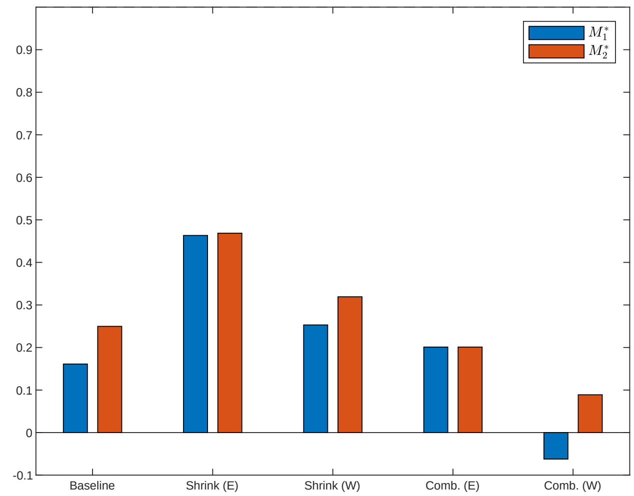

(A) Stock market

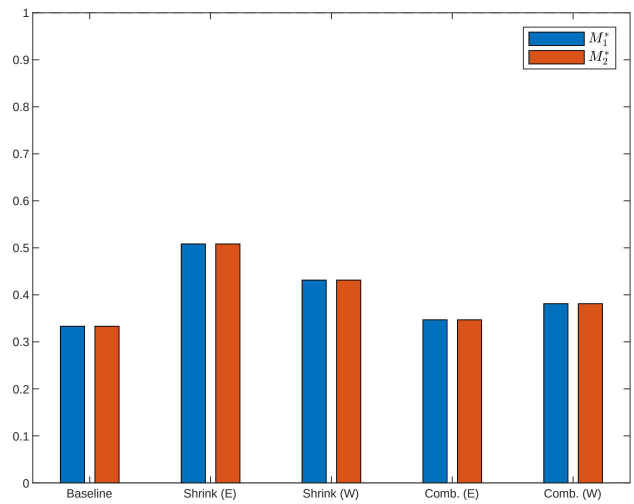

(B) Treasury bond market

FIGURE 8

Comparing Magnitudes of Coefficients from Regressing OOS Forecasts and Realized Returns

The left-most pair of bars labeled *Baseline* repeats the estimates from Figure 7 in the *Single (F)* specification. The next four pairs of bars plot the following ratios, respectively,

$$M_1^* \equiv \frac{\sum_{k\in\mathbb{K}} \operatorname{sgn}\left(\beta_k^{Adj}\right)\cdot \beta_k^{OOS}}{\sum_{k\in\mathbb{K}} \operatorname{sgn}\left(\beta_k^{Adj}\right)\cdot \beta_k^{Adj}}, \quad M_2^* \equiv \frac{\sum_{k\in\mathbb{K}} |\beta_k^{OOS}|}{\sum_{k\in\mathbb{K}} |\beta_k^{Adj}|},$$

where β_k^{OOS} are the slope coefficients from regressing OOS excess return forecasts on predictor k; $\mathbb{K}$ is the set of all predictors. *Shrink* applies shrinkage to the predictive regression forecast toward the historical mean. *Comb.* uses combination forecasts as an equal-weighted average of individual-predictor forecasts. *(E)* denotes equal weighting historical data and *(W)* denotes exponential weighting of past data with a quarterly discount factor of 0.982 (and its third root for monthly data).

construct the ratio in (16). The two right-most pairs of bars in the two panels of Figure 8 present the result, one with forecasts based on equally weighting historical data, labeled *Comb. (E)*, the other one based on exponential weighting, labeled *Comb. (W)*. These combination forecasts are very close, in terms of their weak cyclicality, to the subjective expected excess returns from survey data.

Overall, the results in this section show that simply moving away from the extreme assumption of RE models that investors know the data-generating process can go a long way of explaining why there is a gap between subjective risk premium dynamics in survey data and the dynamics of objective risk premia implied by in-sample predictive regressions.

5 Subjective Risk-Return Tradeoff

The evidence so far suggests that to the extent there is any substantial variation in subjective risk premia, standard predictor variables do not capture much of it. However, this does not imply that subjective risk premia are necessarily constant. They do seem to be largely acyclical with regards to typical proxies for business and asset-price cycles, but they could potentially vary systematically in a different way, unrelated to these standard proxies. As a final step in our analysis, we explore whether time-variation in subjective perceptions of risk could generate time-varying subjective risk premia. Due to data constraints, the analysis in this section focuses on the stock market only.

Standard asset-pricing models predict an approximately linear and positive relation between the conditional equity premium and the conditional equity return variance (Campbell and Cochrane 1999; Bansal and Yaron 2004). Empirically, however, evidence on the risk-return trade-off is mixed and inconclusive (see, e.g., Lettau and Ludvigson (2010) for a recent review). This existing evidence is based on measures of objective risk premia and objective measures of time-varying risk. Our analysis focuses on subjective risk premia and subjective perceptions of risk.

Following Lochstoer and Muir (2019), we construct three proxies of perceived stock market risk. The first measure is from the Graham-Harvey CFO survey. In this survey, respondents provide their assessment of the 10th and 90th percentile of the stock market return distribution at a one-year horizon. We convert the range between these percentiles into an estimate of respondents' subjective variance by taking the square of the range and dividing by the square of 2.56 (which would be accurate if the distribution was normal). The remaining two measures are constructed from the United States Crash Confidence Index from the International Center for Finance at Yale (Shiller (2000), Goetzmann et al. (2016)). This index has two series, one for individuals and one for institutions. They are calculated as the percent of respondents who think there is less than 10% chance of a stock market crash in the next six months.[12] We take the negative of this index to proxy for the level of individuals' and institutions' perceived crash risk. Surveys were initially conducted at six-month intervals.[13] Starting in July 2001, the index reports a six-month moving average of monthly surveys. We match this moving average to surveys in the following months. Thus, for example, the index for November 2018 is an average of results from surveys between June 2018 and November 2018, and we match this index to Livingston survey subjective return expectations in December 2018 (obtained from surveys conducted in November).

Our interest now centers on the relationship between perceived risk and subjective excess return expectations. Figure 9 plots the subjective variance data for the CFO survey and the subjective expected excess return series. The plot shows a clear positive relationship between the two series. In particular, the series both increase substantially in the early 2000s, following the technology crash and recession, in the wake of the financial crisis around 2008, and most recently during the COVID crisis in 2020.

Table 10 presents the results in regression form. In addition to the CFO series, we also

[12] The exact survey question is: "What do you think is the probability of a catastrophic stock market crash in the U. S., like that of October 28, 1929 or October 19, 1987, in the next six months, including the case that a crash occurred in the other countries and spreads to the U. S.?"

[13] The survey dates (first mailing dates) were: July 5, 1989; January 17, 1990; July 27, 1990; January 31, 1991; August 20, 1991; January 31, 1992; August 20, 1992; February 12, 1993; August 6, 1993; February 28, 1994; September 8, 1994; March 4, 1995; September 1, 1995; March 1, 1996; July 30, 1996; March 17, 1997; September 5, 1997; March 2, 1998; and September 9, 1998.

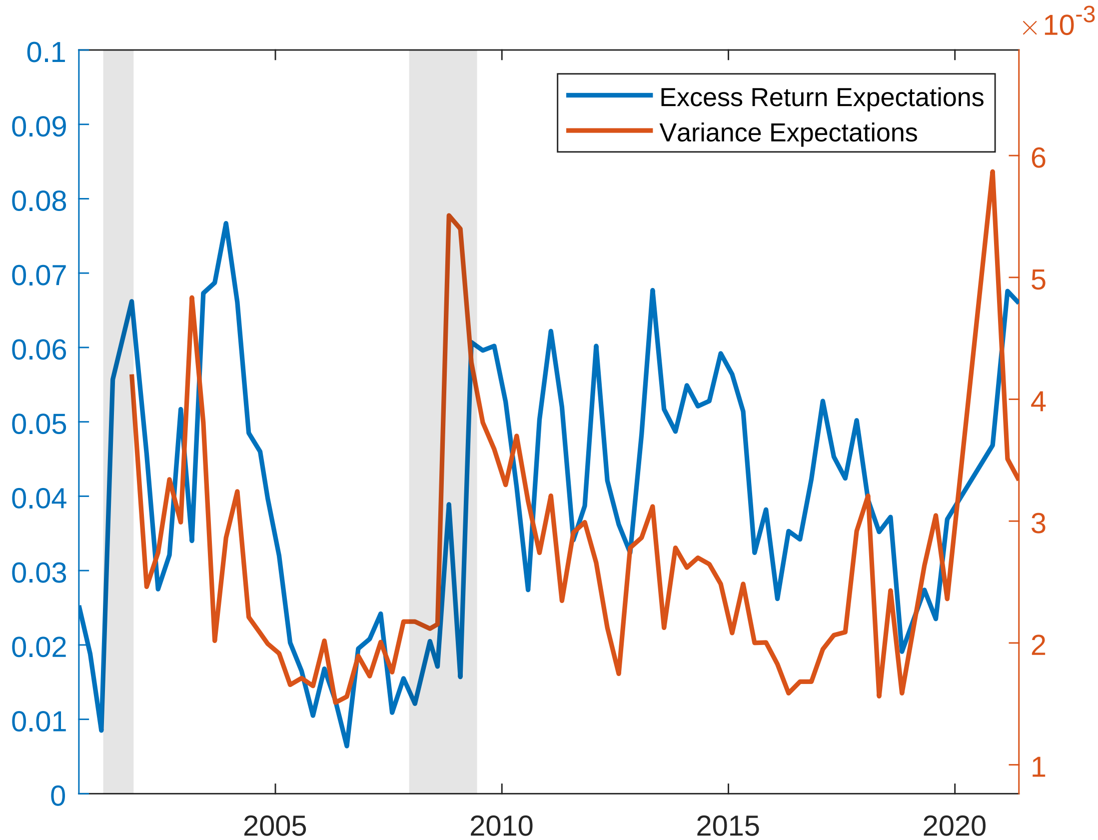

FIGURE 9
CFO Excess Return Expectations and Variance Expectations

The blue line plots the one-year excess return expectations from the CFO survey. The red line plots an approximate measure of the perceived variance: the square of the difference between the mean expectations of 90th and 10th percentiles of returns from the CFO survey, divided by the square of 2.56. The grey-shaded areas indicate NBER recessions.

examine the relationship of individual investor and Livingston survey subjective risk premia to the (negative of) the individual and institutional investor crash confidence indices. To facilitate interpretation, we standardize the crash confidence series to unit standard deviations. We do not standardize the subjective variance series from the CFO survey, as the slope coefficient in a regression of subjective expected excess returns on subjective variance then has a natural economic interpretation as a relative risk aversion coefficient (in a model where investors' wealth is fully invested in the stock market).

Columns (1) and (2) show a weak positive relationship between individuals investors' perceived risk of a stock market crash and their subjective expected excess returns. The relationship is not statistically significant, however. Controlling for past returns in column (2) raises the magnitude of the coefficient on perceived risk, but it is still statistically insignificant at conventional levels ($p = 0.11$). The coefficient of 0.60 implies that a one standard deviation rise in the perceived risk measure is associated with an increase of 0.60pp in the subjective risk premium. For comparison, this is somewhat weaker than the effect of past returns, where a one standard deviation change is associated with a change of 1.01pp in the subjective risk premium.

Columns (3) and (4) show the results for the CFO survey data. As anticipated from Figure 9, subjective variance is positively related to the subjective risk premium.[14] The magnitudes of the slope coefficient estimates are also economically plausible as they would imply a relative risk aversion somewhere between 6 and 9. The strong link between subjective variance and subjective risk premium may seem surprising given that columns (17) and (18) in Table 4, Panel B show the loading of the CFO subjective risk premium on VIX^2 is neither statistically nor economically significant. This indicates that the CFO variance expectations contain information orthogonal to VIX^2 that is important for determining excess return expectations. One potential source of such an orthogonal component is highlighted in Lochstoer and Muir

[14] Lochstoer and Muir (2019) document that survey *return* expectations and variance expectations are uncorrelated. This is also true in our sample period: Using return expectations instead of excess return expectations, the adjusted R^2 drops to zero in both columns (3) and (4).

(2019). They find that CFO variance expectations load much more strongly on realized variance several months ago than VIX^2 does.

Columns (5) and (6) present the results for professional forecasters from the Livingston survey. The relationship between subjective risk premium and subjective risk, as measured by (the negative of) the institutional investor crash confidence index, is much stronger than for individuals. The point estimate in column (5) implies that a one standard deviation rise in perceived risk is associated with an increase in the subjective risk premium of 2.53pp. Statistically, however, the magnitude of the effect is rather uncertain, as the estimate is only marginally significant ($p = 0.10$). To some extent, this is due to the small number of observations. The Livingston survey is available only semi-annually and hence we are left with a small number of observations in the part of the sample that overlaps with the investor crash confidence index series.

For the Livingston survey subjective expectations, the relationship with the risk perceptions captured by the institutional crash confidence index series is quantitatively quite similar to the relationship to VIX^2 in Table 4. For comparison, a one standard deviation move in VIX^2 is associated with a 1.65pp move in the subjective risk premium. In contrast to the CFO expectations, a component of perceived risk orthogonal to the VIX^2 seems to be less important for professional forecasters than for CFOs.[15]

Overall, there is evidence for a positive risk-return tradeoff in subjective beliefs. At the same time, different groups of market participants appear to form subjective beliefs about risks in different ways. For example, using the institutional crash confidence index to explain individuals' subjective expected excess returns would result in a substantial deterioration in explanatory power. Further exploration of the formation of risk perceptions for different groups of market participants seems like a fruitful area for future research.

From the results in this section it is also clear that while subjective risk premia appear to vary with subjective risk perception, this variation does not seem to help explain asset

[15] Relatedly, Buraschi et al. (2021) find that the subjective risk premium of professional forecasters for Treasury bonds comoves with objective measures of the quantity of risk.

Table 10
Subjective Excess Return Expectations and Variance Expectations

The dependent variable is the one-year-ahead subjective excess return expectation from surveys. NX, CFO, and Livingston denote survey excess return expectations from Nagel and Xu (2022), the CFO survey, and the Livingston survey, respectively. Var. Exp. is a perceived return variance from the CFO survey and calculated as the square of the range between reported 10th and 90th percentiles of the stock market return distribution divided by the square of 2.56. Crash (Indiv.) and Crash (Inst.) denote the negative of standardized Crash Confidence Index of individuals and institutions, respectively. R^e_{past} are one-year excess returns standardized to unit standard deviations in the full sample period. We report the OLS estimates with the coefficients for the Crash Index variables multiplied by 100. EWC p-values following Lazarus et al. (2018) are reported in parentheses.

	NX		CFO		Livingston	
	(1)	(2)	(3)	(4)	(5)	(6)
Const	0.08	0.09	0.02	0.01	0.19	0.17
(p-value)	(0.00)	(0.00)	(0.06)	(0.25)	(0.03)	(0.02)
Crash (Indiv.)	0.38	0.60				
(p-value)	(0.36)	(0.11)				
Var. Exp.			6.72	8.65		
(p-value)			(0.06)	(0.04)		
Crash (Inst.)					2.53	1.76
(p-value)					(0.10)	(0.14)
R^e_{past}		1.01		0.74		-2.84
(p-value)		(0.00)		(0.02)		(0.00)
Adj. R^2	0.04	0.28	0.12	0.20	0.24	0.40
N	87	86	75	73	39	39

price volatility because it is disconnected from the predictor variables that capture asset price cycles in in Table 3.

6 Conclusion

Objective risk premia of major asset classes implied by in-sample predictive regressions of excess returns vary over time with cyclical state variables. In stark contrast, we find that subjective risk premia extracted from forecasts in surveys show very little movement with these cyclical variables. This lack of cyclicality of subjective risk premia is pervasive: it holds for stocks, bonds, currencies, and commodity futures and for subjective beliefs of individuals, CFOs, and professional forecasters. While the properties of subjective expected excess return forecasts differ in some respects between groups of market participants and between asset classes—for example, whether they are extrapolative or contrarian with regards to recent realized returns—the pervasive lack of cyclical movement with standard return predictors and business cycle variables is shared among all of them.

Much of the cyclicality gap between subjective risk premia from survey data and statistical forecasts of excess returns disappears when the latter are constructed as out-of-sample rather than in-sample predictions. Thus, moving away from rational expectations models in which investors know the data-generating process and its parameters to models in which they learn about them in real time can help reconcile the evidence from statistical forecasts and subjective excess return expectations data.

While subjective risk premia do not vary much with standard return predictors and business cycle variables, they are not constant, though. In particular, they comove to some extent with time-varying subjective perceptions of risk, consistent with a positive risk-return tradeoff in subjective beliefs. However, the variation in subjective risk premia associated with movements in perceived risk not sufficiently aligned with asset price cycles to contribute much to an explanation of asset price fluctuations.

Therefore, to match the joint set of stylized facts of volatile asset prices, cyclical objective risk premia, and acyclical subjective risk premia, asset-pricing models need to have a time-varying beliefs wedge between subjective and objective forecasts of fundamentals, such as future dividend growth in the case of stocks or future short-term interest rates and inflation in the case of bonds. This wedge must do most of the work in generating large persistent movements in asset prices. Approaches that generate volatile asset prices instead through time-variation in risk or risk aversion—or other mechanisms that induce variation in the excess return that investors require and subjectively expect to earn—are not consistent with the evidence that subjective risk premia are largely unrelated to the state variables that capture asset price cycles.

Appendix

A Data Sources and Timing Assumptions

A.1 Data sources

Interest Rates: The main interest rate we use to calculate one-year excess returns is the one-year Treasury Constant Maturity rate. To extend the time series, we use annualized three-month Treasury bill yields: starting from 1934, we use *3-Month Treasury Bill: Secondary Market Rate*; before 1934, we use *Yields on Short-Term United States Securities, Three-Six Month Treasury Notes and Certificates, Three Month Treasury Bills for United States*. All data are available from FRED.

Stock Returns: For monthly stock returns, we use the value-weighted returns on the CRSP index since January 1926 from CRSP. To obtain longer-horizon returns we compound one-month returns.

Bond Returns: Monthly returns on two-, five-, seven-, and ten-year Treasury indexes come from CRSP. We only use realized returns starting from April 1951 after the Treasury-Fed Accord ended the yield control. To obtain longer-horizon returns we compound one-month returns. We subtract the one-year interest rates described in the first paragraph to obtain excess returns.

Bond Yields: Monthly zero-coupon yields come from Liu and Wu (2021) which are available starting from June 1961.[16] We use these yields to construct subjective bond return expectations and the cycle factor from Cieslak and Povala (2015).

FX Rates and Interbank Interest Rates: Daily spot and forward exchange rates used to construct returns are from Datastream. The main provider on Datastream is World Markets PLC/Reuters (WMR) and we supplement data from other providers including Barclays Bank PLC (BB), Thomson/Reuters (TR), and HSBC. The mnemonics are *idiso*SP for spot rates and *idiso*1Y or *idiso*YF for 1-year forward rates.[17] *id* corresponds to the data provider: *US* for WMR, *BB* for BB, *TD* for TR, and *MB* for HSBC. *iso* is the currency ISO code of each country. The spot rate series are also supplemented with Sterling-based quotes from WMR (converted to USD-based).

We obtain daily interbank interest rates and Eurocurrency deposit rates from Datastream and Global Financial Data. The providers of Eurocurrency deposit rates on Datastream include Refinitiv and Garban Information Services. Their mnemonics are EC*iso*1Y and GS*iso*1Y, respectively. The providers of interbank rates on Datastream are national sources including central banks. Interbank rates from Global Financial Data have mnemonics in the format of IB*gfd*12D. *gfd* is the ISO 3166-1 code.

Some of the series clearly contain outliers due to data errors. Following Hassan and Mano (2019), we calculate the log forward premium across data providers, exclude observations where deviations of the log interest rate differential exceed 50 bps, and splice data together. Finally, we build monthly series as the end-of-calendar-month values.

[16] We thank the authors for providing the data on their website.

[17] Mnemonics of spot rates from WMR do not have a clear pattern. For example, *AUSTDO$* corresponds to quotes of the Australian Dollar.

Table A.1
Mnemonics of Commodity Futures from Bloomberg and Datastream

This table reports the mnemonics of spot and first-, second-, and third-genereic futures prices from Bloomberg. Datastream mnemonics are reported in the second row for each metal.

Commodity	Spot	First	Second	Third
Panel A: Energy				
Crude Oil		CL1 Comdty	CL2 Comdty	CL3 Comdty
Panel B: Metals				
Aluminium		LMAHDY LME Comdty	LMAHDS03 LME Comdty	LMAHDS15 LME Comdty
		LAHCASH (P)	LAH3MTH (P)	LAH15MT (P)
Copper		LMCADY LME Comdty	LMCADS03 LME Comdty	LMCADS15 LME Comdty
		LCPCASH (P)	LCP3MTH (P)	LCP15MT (P)
Gold	GOLDLNPM Index	GC1 Comdty	GC2 Comdty	GC3 Comdty
	GOLDBLN (P)			
Silver	SLVRLND Index	SI1 Comdty	SI2 Comdty	SI3 Comdty
	SILVUSL (P)			

Commodity Prices and Open Interest: Spot rates and futures prices are obtained from Bloomberg. For metals traded on London Metal Exchange (LME), we supplement data from Datastream. Table A.1 reports the mnemonics from these databases. For commodities without reliable spot prices, we interpolate the futures curve to compute the synthetic spot prices (e.g., crude oil). For metals traded on LME, only spot prices and futures prices at three- and fifteen-month horizons are available. We interpolate the twelve-month futures prices linearly using adjacent futures as knots. The open interest data comes from *Commitments of Traders* available at the Commodity Futures Trading Commission. Open interest data before 2010 comes from Hong and Yogo (2012).[18]

A.2 Survey details and timing assumptions

For each survey, its survey month is the reported calendar month in which survey results are released. We describe in detail how surveys are conducted below:

Nagel and Xu (2022): The major surveys used are the UBS/Gallup survey, the Conference Board survey, and the Michigan Survey of Consumers. The UBS/Gallup survey typically interviews households during the first two weeks in the survey month. Questionnaires from the Conference Board survey reach households around the first of each survey month. The responses are collected before the eighteenth of the month. Interviews for the Michigan Survey of Consumers typically begin at the beginning of the survey month or the end of previous month.

Duke CFO: The CFO Survey is conducted quarterly within a two-week fielding period

[18] We thank the authors for providing the data on their website.

(running from Monday of the first week through Friday of the second week). (A special case is the 2020Q1 survey, which was conducted over six weeks between February 25 and April 3.) The survey typically starts late in the middle or early in the last calendar month of each quarter and the starting date is provided. We define the survey month as the month that the starting date belongs to.

Livingston: Survey months for the Livingston survey are June and December. The questionnaires are mailed to participants in May and November right after the release of CPI (typically in the last week).

BCFF: The survey is typically released on the first day of each calendar month. The responses are collected during the last week of the month before the survey month.

Consensus Economics: For currency spot rate forecasts, the responses are recorded on the first or second Monday in each survey month. For commodity spot price forecasts, the survey responses are normally collected on the second Monday or third Monday in each survey month.

FX4casts: The survey questions are usually sent out on the last Friday of the survey month with responses collected during Friday and the following Monday and Tuesday.

This leads to the following timing assumptions: In a given survey month, we assume participants in all surveys except the FX4casts survey have access to monthly macro information with a two-month lag and price information with a one-month lag, e.g., June survey participants have information of industrial production up to April and stock prices up to May. This assumption is consistent with the Livingston Survey Documentation and the approach in Nagel and Xu (2022). We match the one-month lags of FX4casts survey to the Consensus Economics survey when constructing the currency forecasts, e.g., May survey from FX4casts to June survey from the Consensus Economics.

B Bootstrap Methods for Predictive Regressions

To correct for the potential small-sample bias in the predictive regressions, we assume the following system:

$$z_{t+1} = \alpha + \boldsymbol{\beta}'\boldsymbol{x}_t + \eta_{t+1}, \tag{B.1}$$

$$\boldsymbol{x}_{t+1} = \boldsymbol{\kappa} + \boldsymbol{\Phi}\boldsymbol{x}_t + \boldsymbol{\iota}_{t+1}, \tag{B.2}$$

where $\boldsymbol{x}_t$ is a $K \times 1$ vector of predictors, $\boldsymbol{\beta}$ is a $K \times 1$ vector of loadings, $\boldsymbol{\Phi}$ is a $K \times K$ coefficient matrix, and $\eta_t = \boldsymbol{\rho}'\boldsymbol{\iota}_t + \xi_t$ where ξ_t is independent from $\boldsymbol{\iota}_t$. We implement a bootstrap that involves the following steps:

1. We start with estimating a VAR(1) system for the predictors, $\boldsymbol{x}_t$. To correct for the bias in the VAR estimates, we use the approach in Amihud et al. (2009) that relies on the analytical expressions from Nicholls and Pope (1988). We denote the bias-adjusted coefficients and shocks as $\tilde{\kappa}$, $\tilde{\boldsymbol{\Phi}}$, and $\{\tilde{\boldsymbol{\iota}}_t\}$.

2. We then estimate the predictive regression in (B.1) with OLS. We obtain coefficient estimates $\hat{\alpha}$, $\hat{\boldsymbol{\beta}}$ and residuals $\hat{\eta}_{t+1}$. We then construct pseudo samples by bootstrapping

the time-series of residual vectors $(\tilde{\boldsymbol{\iota}}_t, \hat{\eta}_t)$ to preserve the cross-sectional correlations. To account for the potentially autocorrelated $\{\eta_t\}$, we use a circular block bootstrap. For the sample i, with bootstrapped residuals $(\tilde{\boldsymbol{\iota}}_t^i, \hat{\eta}_t^i)$, we impose the null of no predictability by generating data as

$$z_{t+1}^{null,i} = \bar{z} + \hat{\eta}_{t+1}^i, \tag{B.3}$$

$$\boldsymbol{x}_{t+1} = \tilde{\boldsymbol{\kappa}} + \tilde{\boldsymbol{\Phi}}\boldsymbol{x}_t + \tilde{\boldsymbol{\iota}}_{t+1}^i. \tag{B.4}$$

We then re-run the predictive regression in the sample i and record the t-statistic τ^i. The $\{\tau^i\}$ are used to obtain the small-sample p-value by comparing with the sample t-statistic $\hat{\tau}$. For the sample j, with bootstrapped residuals $(\tilde{\boldsymbol{\iota}}_t^j, \hat{\eta}_t^j)$, we generate data under the alternative that $\boldsymbol{\beta} = \hat{\boldsymbol{\beta}}$ as

$$z_{t+1}^{alter,j} = \hat{\alpha} + \hat{\boldsymbol{\beta}}'\boldsymbol{x}_t + \hat{\eta}_{t+1}^j, \tag{B.5}$$

$$\boldsymbol{x}_{t+1} = \tilde{\boldsymbol{\kappa}} + \tilde{\boldsymbol{\Phi}}\boldsymbol{x}_t + \tilde{\iota}_{t+1}^j. \tag{B.6}$$

We re-run the predictive regression in the sample j and record the coefficients $\hat{\boldsymbol{\beta}}^j$. By comparing the average of the $\{\hat{\boldsymbol{\beta}}^j\}$ with $\hat{\boldsymbol{\beta}}$, we obtain the finite-sample bias in $\hat{\boldsymbol{\beta}}$.

C Asymptotic Inference Approach for Composite Measures of Regression Coefficients

We develop an approach to perform statistical inference on composite measures of the differences in regression coefficients between the realized excess return regressions and the survey data regressions. Looking across all J specifications of dependent variables (realized excess returns and survey expectations) and predictors (asset-class specific and business cycle variables), we can stack all observations in specification j as

$$\underbrace{\begin{pmatrix} y_{j,1} \\ \cdots \\ y_{j,N_j} \end{pmatrix}}_{\boldsymbol{y}_j} = \underbrace{\begin{pmatrix} \boldsymbol{x}_{j,1} \\ \cdots \\ \boldsymbol{x}_{j,N_j} \end{pmatrix}}_{\boldsymbol{X}_j} \boldsymbol{b}_j + \underbrace{\begin{pmatrix} \varepsilon_{j,1} \\ \cdots \\ \varepsilon_{j,N_j} \end{pmatrix}}_{\boldsymbol{\epsilon}_j}, \quad j = 1, 2, ..., J, \tag{C.1}$$

where $\boldsymbol{x}_{j,k}$ represents k-th observation of the $1 \times K$ dimensional predictor variable vector used in specification j (with first element equal to unity).[19] In regressions with realized excess returns $y_{j,k}$ is an excess return, while in regressions with survey data, $y_{j,k}$ is a subjective risk premium measured in survey data. Returns are typically from overlapping windows, but we did not write this out explicitly here. Our statistical inference procedure takes the overlap into account, though. Across different specifications j, the predictor variables can be the same, but even then the sample period may be different due to different availability

[19] In our specification, $K = 2$ or 3 depending on whether past excess return is included as an additional control variable.

of excess return data or survey data in different asset classes. Measurement frequency can differ between specifications, too. We denote the corresponding calendar time of available data points in each specification as $\{T_{j,k}\}_{k=1}^{N_j}$ and measure calendar time in months. Note that

$$\hat{\boldsymbol{b}}_j = \boldsymbol{b}_j + (\boldsymbol{X}_j'\boldsymbol{X}_j)^{-1}\boldsymbol{X}_j'\boldsymbol{\epsilon}_j. \tag{C.2}$$

Under regularity conditions, we have

$$\sqrt{N_j}(\hat{\boldsymbol{b}}_j - \boldsymbol{b}_j) \xrightarrow{p} \mathcal{N}(\boldsymbol{0},\ \boldsymbol{Q}_j), \tag{C.3}$$

where

$$\hat{\boldsymbol{Q}}_j = \left(\frac{1}{N_j}\sum_{k=1}^{N_j}\boldsymbol{x}_{j,k}'\boldsymbol{x}_{j,k}\right)^{-1}\hat{\boldsymbol{\Omega}}_j\left(\frac{1}{N_j}\sum_{k=1}^{N_j}\boldsymbol{x}_{j,k}'\boldsymbol{x}_{j,k}\right)^{-1} \tag{C.4}$$

and $\hat{\boldsymbol{\Omega}}_j$ is an estimator of long-run variance of $\boldsymbol{x}_{j,t}'\epsilon_t$. We may write

$$\sqrt{T}(\hat{\boldsymbol{b}}_j - \boldsymbol{b}_j) \xrightarrow{p} \mathcal{N}\left(\boldsymbol{0},\ \frac{T}{N_j}\boldsymbol{Q}_j\right). \tag{C.5}$$

Here T represents the maximum sample period length expressed in terms of monthly observations for any of our regressions. In our settings, N_j/T is either a constant (quarterly v.s. monthly data) or converges to 1 (missing early observations) as T grows.

For two specifications i and j, we can estimate the covariance of $\hat{\boldsymbol{b}}_i - \boldsymbol{b}_i$ and $\hat{\boldsymbol{b}}_j - \boldsymbol{b}_j$ as

$$\hat{\boldsymbol{V}}_{i,j} \equiv \hat{\text{cov}}\big(\sqrt{N_i}(\hat{\boldsymbol{b}}_i - \boldsymbol{b}_i),\ \sqrt{N_j}(\hat{\boldsymbol{b}}_j - \boldsymbol{b}_j)\big) \tag{C.6}$$

$$= \left(\frac{1}{N_i}\sum_{k=1}^{N_i}\boldsymbol{x}_{i,k}'\boldsymbol{x}_{i,k}\right)^{-1}\hat{\boldsymbol{S}}_{ij}\left(\frac{1}{N_j}\sum_{k=1}^{N_j}\boldsymbol{x}_{j,k}'\boldsymbol{x}_{j,k}\right)^{-1}, \tag{C.7}$$

where $\hat{\boldsymbol{S}}_{ij}$ is an estimator of long-run covariance between $\boldsymbol{x}_{i,k}'\boldsymbol{\epsilon}_i$ and $\boldsymbol{x}_{j,k}'\boldsymbol{\epsilon}_j$.

We stack all $\boldsymbol{b}_j$ in the $K \times J$ matrix $\boldsymbol{B}$:

$$\boldsymbol{B} \equiv (\boldsymbol{b}_1,\ \boldsymbol{b}_2,\ \ldots, \boldsymbol{b}_J). \tag{C.8}$$

We treat all these regressions jointly as one big system of regression equations. We can then construct an estimator of the difference between realized return-based and survey-based coefficient estimates as a linear combination of coefficients as

$$d = \boldsymbol{e}_2'\boldsymbol{B}\boldsymbol{w}, \tag{C.9}$$

where the $J \times 1$ weight vector $\boldsymbol{w} \equiv (w_1,\ w_2, \ldots, w_J)'$ is such that the elements that multiply coefficients from the predictive regressions with excess returns as dependent variable are positive, of equal magnitude, and sum to 1 across all regressions within an asset class, while those that multiply coefficients from the survey data regressions are negative, of equal magnitude, and sum to -1 within an asset class; $\boldsymbol{e}_2$ is a $K \times 1$ vector with one as the second element and zeros otherwise.[20]

Let $\hat{\boldsymbol{B}}$ and $\hat{d}$ denote the estimators of $\boldsymbol{B}$ and d, respectively. We have

$$(\hat{\boldsymbol{B}} - \boldsymbol{B})\boldsymbol{w} = \sum_{k=1}^{J} w_k (\boldsymbol{X}_k' \boldsymbol{X}_k)^{-1} (\boldsymbol{X}_k' \boldsymbol{\epsilon}_k). \tag{C.10}$$

Combining (C.5) and (C.7), we have

$$\sqrt{T}(\hat{\boldsymbol{B}} - \boldsymbol{B})\boldsymbol{w} \xrightarrow{p} \mathcal{N}(\boldsymbol{0},\ \boldsymbol{W}) \tag{C.11}$$

where

$$\boldsymbol{W} = \sum_{k=1}^{J} w_k^2 \frac{T}{N_k} \boldsymbol{Q}_k + \sum_{i \neq j} w_i w_j \frac{T}{\sqrt{N_i N_j}} \boldsymbol{V}_{i,j}. \tag{C.12}$$

It follows that

$$\sqrt{T}(\hat{d} - d) \xrightarrow{p} \mathcal{N}(0,\ \boldsymbol{e}_2' \boldsymbol{W} \boldsymbol{e}_2). \tag{C.13}$$

To estimate $\boldsymbol{\Omega}_j$ and $\boldsymbol{S}_{i,j}$, we use HAR estimators (Sun 2013) to account for autocorrelations in residuals. If the data were complete across all specifications, i.e., $\{T_{i,k}\}_{k=1}^{N_i} = \{T_{j,k}\}_{k=1}^{N_j}$, $\forall i, j$, we would have

$$\hat{\boldsymbol{\Omega}}_j = \frac{1}{N_j} \sum_{k=1}^{N_j} \sum_{s=1}^{N_j} \boldsymbol{x}_{j,k}' \epsilon_{j,k} K_G \left(\frac{k}{N_j}, \frac{s}{N_j} \right) (\boldsymbol{x}_{j,s}' \epsilon_{j,s})', \tag{C.14}$$

$$\hat{\boldsymbol{S}}_{i,j} = \frac{1}{\sqrt{N_i N_j}} \sum_{k=1}^{N_j} \sum_{s=1}^{N_j} \boldsymbol{x}_{i,k}' \epsilon_{i,k} K_G \left(\frac{k}{N_i}, \frac{s}{N_j} \right) (\boldsymbol{x}_{j,s}' \epsilon_{j,s})', \tag{C.15}$$

where $K_G(r, s) \equiv \frac{1}{B} \sum_{b=1}^{B} \phi_b(r) \phi_b(s)$ and the $\{\phi_k\}$ are a sequence of orthonormal basis functions. We choose the EWC estimator from Lazarus et al. (2018) and use

$$\phi_b \left(\frac{t}{T} \right) = \sqrt{2} \cos \left[\pi b \left(\frac{t - 1/2}{T} \right) \right]. \tag{C.16}$$

To deal with missing observations, for each $j \in \{1, \ldots, J\}$, we first project $\boldsymbol{x}_{j,k}' \epsilon_{j,k}$ onto the basis functions

$$\sqrt{2} \cos \left[\pi b \left(\frac{t - 1/2}{T_{j,N_j} - T_{j,1} + 1} \right) \right], \quad b = 1, \ldots, B, \tag{C.17}$$

[20] As a concrete example, consider the set of regressions in the equity market which uses the dividend-price ratio and industrial production growth as sole predictors. We have $J = 8$ (two predictive regressions and six survey expectations regressions) and $K = 2$. We stack the results from realized excess returns in the first two columns in $\boldsymbol{B}$ and results from survey expectations regressions in the remaining columns. The weight vector in this case is a 8×1 vector:

$$\boldsymbol{w} = (\frac{1}{2},\ \frac{1}{2},\ -\frac{1}{6},\ -\frac{1}{6},\ -\frac{1}{6},\ -\frac{1}{6},\ -\frac{1}{6},\ -\frac{1}{6})'.$$

Then d is the linear combination of the slope coefficients from the eight regressions.

for $t \in T_j^* \equiv \{T_{j,1}, T_{j,1}+1, \ldots, T_{j,N_j}\}$. In other words, we define the time period for estimation as the span between the earliest and latest calendar dates with non-missing observations. Equation (C.14) is then modified to

$$\hat{\boldsymbol{\Omega}}_j = \frac{1}{N_j} \sum_{k=1}^{N_j} \sum_{s=1}^{N_j} \boldsymbol{x}'_{j,k} \epsilon_{j,k} K_G \left(\frac{T_{j,k} - T_{j,1} + 1}{T_{j,N_j} - T_{j,1} + 1}, \frac{T_{j,s} - T_{j,1} + 1}{T_{j,N_j} - T_{j,1} + 1} \right) (\boldsymbol{x}'_{j,s} \epsilon_{j,s})'. \tag{C.18}$$

This amounts to treating missing observations as non-serially correlated by setting missing residuals to zero. Datta and Du (2012) show that for the Newey-West estimator this approach has better size properties in small samples. Similarly, to estimate $\boldsymbol{S}_{i,j}$, we project both series onto

$$\sqrt{2} \cos \left[\pi b \left(\frac{t - 1/2}{L_{i,j}} \right) \right], \quad b = 1, \ldots, B, \tag{C.19}$$

for $t \in T_{i,j}^* \equiv \{\min(T_{i,1}, T_{j,1}), \min(T_{i,1}, T_{j,1})+1, \ldots, \max(T_{i,N_i}, T_{j,N_j})\}$ and $L_{i,j} \equiv \max(T_{i,N_i}, T_{j,N_j}) - \min(T_{i,1}, T_{j,1}) + 1$. Equation (C.15) is then modified to

$$\hat{\boldsymbol{S}}_{i,j} = \frac{1}{\sqrt{N_i N_j}} \sum_{k=1}^{N_j} \sum_{s=1}^{N_j} \boldsymbol{x}'_{i,k} \epsilon_{i,k} K_G \left(\frac{T_{i,k} - \min(T_{i,1}, T_{j,1}) + 1}{L_{i,j}}, \frac{T_{i,s} - \min(T_{i,1}, T_{j,1}) + 1}{L_{i,j}} \right) (\boldsymbol{x}'_{j,s} \epsilon_{j,s})'. \tag{C.20}$$

D Construction of OOS Forecasts

D.1 Constructing real-time bond excess return predictors

To construct the real-time version of the predictor as in Ludvigson and Ng (2009), we perform the following predictive regressions at time t over expanding windows:

$$\overline{R}_k^e = \gamma_0 + \gamma_1 \hat{F}_{1,k-1} + \gamma_2 \hat{F}_{1,k-1}^3 + \gamma_3 \hat{F}_{2,k-1} + \gamma_4 \hat{F}_{3,k-1} + \gamma_5 \hat{F}_{4,k-1} + \gamma_6 \hat{F}_{8,k-1} + u_k, \quad k = 2, \ldots, t, \tag{D.1}$$

where $\overline{R}_k^e$ denotes the average bond excess returns and $\hat{F}_j$ are the principal components of a broad set of macro variables.[21] The predictor is then calculated as

$$LN_t = \hat{\gamma}_{0,t} + \hat{\gamma}_{1,t} \hat{F}_{1,t} + \hat{\gamma}_{2,t} \hat{F}_{1,t}^3 + \hat{\gamma}_{3,t} \hat{F}_{2,t} + \hat{\gamma}_{4,t} \hat{F}_{3,t} + \hat{\gamma}_{5,t} \hat{F}_{4,t} + \hat{\gamma}_{6,t} \hat{F}_{8,t}. \tag{D.2}$$

To construct the real-time version of the predictor as in Cieslak and Povala (2015), we first perform the following regressions at time t over expanding windows:

$$y_k^n = a_n + b_n \tau_k^{CPI} + \eta_k, \quad k = 1, \ldots, t, \tag{D.3}$$

where y_k^n denotes the n-year nominal yield and τ_k^{CPI} denotes the trend inflation. To have a sufficiently longer time series that allows a burn-in period, here we only use bond yields with

[21] Note that the estimation of $\hat{F}_j$ still uses full-sample time series of macro variables and contains look-ahead information. However, the constructed factor no longer uses look-ahead information in bond returns.

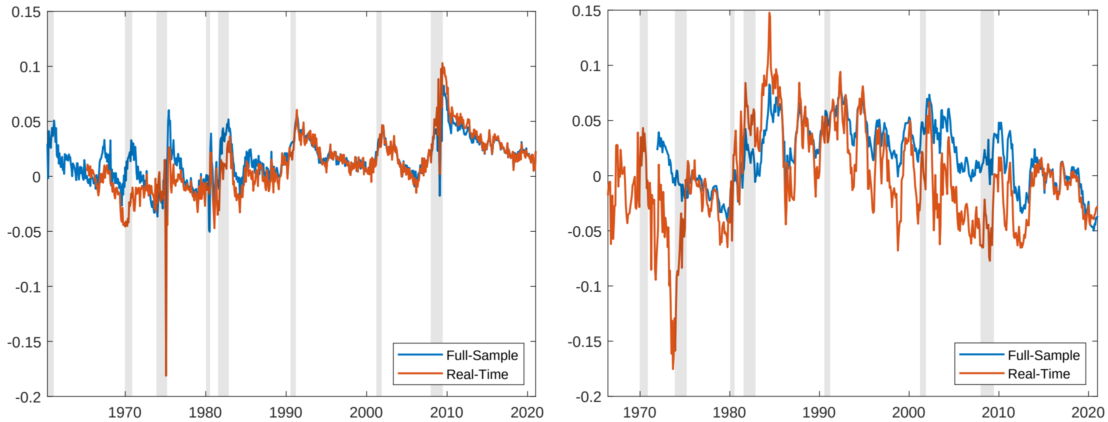

(A) Ludvigson and Ng (2009) (B) Cieslak and Povala (2015)

Figure D.1

Real-Time Bond Excess Return Predictors

The blue line in each panel plots the bond excess return predictors estimated using full-sample data. The red line in each panel plots the real-time versions of these predictors estimated over expanding windows.

maturities between one to seven years.[22] We then run the following predictive regressions:

$$\overline{R}^e_k = \delta_0 + \delta_1 c^1_{t,k-1} + \delta_2 \overline{c}_{t,k-1} + \omega_k, \quad k = 2, \ldots, t, \tag{D.4}$$

with $c^n_{t,k} = y^n_k - \hat{a}_{n,t} + \hat{b}_{n,t}\tau^{CPI}_k$ and $\overline{c}_{t,k} = \frac{1}{6}\sum_{i=2}^{7} c^i_{t,k}$. The real-time cycle factor is then calculated as

$$CYCLE_t = \hat{\delta}_{0,t} + \hat{\delta}_{1,t} c^1_{t,t} + \hat{\delta}_{2,t}\overline{c}_{t,t}. \tag{D.5}$$

We estimate the above regressions with an initial burn-in period of five years. Figure D.1 plots the time series of these two factors in comparison with their full-sample counterparts.

D.2 Individual regression coefficients

Table D.1 and Table D.2 report the detailed results of regressing OOS stock and bond excess return forecasts on asset-specific predictors and cyclical indicators. The OOS forecasts are constructed as described in Section 4.

[22] Longer-maturity yields are only available starting from August 1971.

Table D.1
Regressing OOS Stock Excess Return Forecasts on Predictors

In Panels A and B, dependent variables are the quarterly OOS stock excess returns forecasts from predictive regressions over expanding windows with single predictors. OOS forecasts are shrunk toward the trailing mean excess return with optimal shrinkage estimated based on historical OOS performance in earlier periods. Forecasts in Panel A equally weight historical data in estimating predictive regressions and shrinkage, whereas forecasts in Panel B exponentially downweight past data with a quarterly discount factor of 0.982. The regressions in Panels A and B regress the OOS forecasts on the same predictor variable that was used to construct the OOS forecasts. In Panels C and D, dependent variables are combination OOS forecasts that average the forecasts (without shrinkage) obtained from different predictors. All predictors are standardized to have unit standard deviations in the full sample period. The first row in each block reports the OLS estimates multiplied by 100. EWC p-values following Lazarus et al. (2018) are reported in parentheses.

	Asset-Class Specific				Business-Cycle				
	CAY	D/P	EXPD	NTIS	N-IP	TERM	DEFAULT	F1	VIX^2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
A. Shrinkage, Equally Weighted									
Coeff	2.58	5.29	-0.74	-2.14	1.54	3.41	0.20	0.75	-0.10
(p-value)	(0.00)	(0.00)	(0.01)	(0.09)	(0.00)	(0.00)	(0.33)	(0.01)	(0.10)
Adj. R^2	0.96	0.89	0.27	0.18	0.47	0.78	0.02	0.47	0.02
N	109	210	214	210	214	214	214	76	214
B. Shrinkage, Exponentially Weighted									
Coeff	-0.03	3.76	-0.04	-2.72	1.81	1.47	-0.41	-0.74	0.41
(p-value)	(0.93)	(0.02)	(0.93)	(0.16)	(0.13)	(0.02)	(0.54)	(0.01)	(0.24)
Adj. R^2	-0.01	0.39	-0.00	0.15	0.14	0.21	0.01	0.19	0.02
N	109	210	214	210	214	214	214	76	214
C. Combination Forecasts, Equally Weighted									
Coeff	0.52	1.22	-0.68	-0.50	0.94	0.58	1.48	0.89	0.37
(p-value)	(0.05)	(0.00)	(0.02)	(0.25)	(0.12)	(0.03)	(0.00)	(0.00)	(0.17)
Adj. R^2	0.08	0.32	0.10	0.04	0.07	0.10	0.27	0.23	0.05
N	275	275	275	27	275	275	275	243	275
D. Combination Forecasts, Exponentially Weighted									
Coeff	0.24	0.06	0.57	0.29	-0.61	-0.36	-0.84	-0.02	0.18
(p-value)	(0.49)	(0.90)	(0.19)	(0.62)	(0.26)	(0.36)	(0.21)	(0.94)	(0.58)
Adj. R^2	0.01	-0.00	0.04	0.01	0.02	0.02	0.05	-0.00	0.00
N	275	275	275	275	275	275	275	243	275

Table D.2

Regressing OOS Treasury Bond Average Excess Return Forecasts on Predictors

In Panels A and B, dependent variables are the monthly OOS Treasury bond excess returns forecasts from predictive regressions over expanding windows with single predictors. OOS forecasts are shrunk toward the trailing mean excess return with optimal shrinkage estimated based on historical OOS performance in earlier periods. Forecasts in Panel A equally weight historical data in estimating predictive regressions and shrinkage, whereas forecasts in Panel B exponentially downweight past data with a monthly discount factor of 0.994. The regressions in Panels A and B regress the OOS forecasts on the same predictor variable that was used to construct the OOS forecasts. In Panels C and D, dependent variables are combination OOS forecasts that average the forecasts (without shrinkage) obtained from different predictors. All predictors are standardized to have unit standard deviations in the full sample period. The first row in each block reports the OLS estimates multiplied by 100. EWC p-values following Lazarus et al. (2018) are reported in parentheses.

	Asset-Class Specific		Business-Cycle				
	LN	CYCLE	N-IP	TERM	DEFAULT	F1	VIX^2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Shrinkage, Equally Weighted							
Coeff	0.93	1.55	0.21	1.44	0.09	0.44	0.01
(p-value)	(0.00)	(0.00)	(0.04)	(0.00)	(0.14)	(0.00)	(0.77)
Adj. R^2	0.43	0.43	0.07	0.94	0.02	0.41	-0.00
N	397	393	397	397	397	397	397
B. Shrinkage, Exponentially Weighted							
Coeff	0.98	0.97	0.21	1.36	0.08	0.28	0.09
(p-value)	(0.01)	(0.03)	(0.12)	(0.00)	(0.32)	(0.13)	(0.33)
Adj. R^2	0.41	0.16	0.03	0.83	0.00	0.14	0.02
N	397	393	397	397	397	397	397
C. Combination Forecasts, Equally Weighted							
Coeff	0.54	0.09	0.80	0.55	0.67	0.39	0.15
(p-value)	(0.00)	(0.22)	(0.01)	(0.00)	(0.00)	(0.00)	(0.03)
Adj. R^2	0.45	0.02	0.27	0.63	0.31	0.23	0.12
N	397	397	397	397	397	397	397
D. Combination Forecasts, Exponentially Weighted							
Coeff	0.64	0.05	0.89	0.67	0.72	0.39	0.14
(p-value)	(0.00)	(0.59)	(0.01)	(0.00)	(0.00)	(0.03)	(0.10)
Adj. R^2	0.42	0.00	0.23	0.64	0.24	0.15	0.07
N	397	397	397	397	397	397	397

References

Adam, Klaus, Albert Marcet, and Johannes Beutel, 2017, Stock price booms and expected capital gains, *American Economic Review* 107, 2352–2408.

Amihud, Yakov, Clifford M. Hurvich, and Yi Wang, 2009, Multiple-Predictor Regressions: Hypothesis Testing, *The Review of Financial Studies* 22, 413–434.

Bacchetta, Philippe, Elmar Mertens, and Eric Van Wincoop, 2009, Predictability in financial markets: What do survey expectations tell us?, *Journal of International Money and Finance* 28, 406–426.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles, *The Journal of Finance* 59, 1481–1509.

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

Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey, 2013, Managerial Miscalibration, *The Quarterly Journal of Economics* 128, 1547–1584.

Bordalo, Pedro, Nicola Gennaioli, Spencer Yongwook Kwon, and Andrei Shleifer, 2021, Diagnostic bubbles, *Journal of Financial Economics* 141, 1060–1077.

Buraschi, Andrea, Ilaria Piatti, and Paul Whelan, 2021, Subjective bond returns and belief aggregation, *Review of Financial Studies* forthcoming.

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, 205–251.

Cieslak, Anna, and Pavol Povala, 2015, Expected Returns in Treasury Bonds, *The Review of Financial Studies* 28, 2859–2901.

Cochrane, John H, 2017, Macro-finance, *Review of Finance* 21, 945–985.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A Lochstoer, 2017, Asset pricing when 'this time is different', *Review of Financial Studies* 30, 505–535.

Dahlquist, Magnus, and Markus Ibert, 2021, How cyclical are stock market return expectations? evidence from capital market assumptions, Technical report, Stockholm School of Economics.

Datta, Deepa, and Wenxin Du, 2012, Nonparametric HAC Estimation for Time Series Data with Missing Observations, *International Finance Discussion Paper No.1060, Board of Governors of the Federal Reserve System* .

De La O, Ricardo, and Sean Myers, 2021, Subjective cash flows and discount rates, *Journal of Finance* 76, 1339–1387.

Diebold, Francis X, and Canlin Li, 2006, Forecasting the term structure of government bond yields, *Journal of econometrics* 130, 337–364.

Farmer, Leland, Emi Nakamura, and Jón Steinsson, 2021, Learning about the long run, Technical report, National Bureau of Economic Research.

Goetzmann, William N, Dasol Kim, and Robert J Shiller, 2016, Crash beliefs from investor surveys, Technical report, National Bureau of Economic Research.

Gourinchas, Pierre-Olivier, and Aaron Tornell, 2004, Exchange rate puzzles and distorted beliefs, *Journal of International Economics* 64, 303–333.

Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of returns and expected returns, *Review of Financial Studies* 27, 714–746.

Hassan, Tarek A, and Rui C Mano, 2019, Forward and Spot Exchange Rates in a Multi-Currency World, *The Quarterly Journal of Economics* 134, 397–450.

Hong, Harrison, and Motohiro Yogo, 2012, What Does Futures Market Interest Tell Us about the Macroeconomy and Asset Prices?, *Journal of Financial Economics* 105, 473 – 490.

Jin, Lawrence J, and Pengfei Sui, 2021, Asset pricing with return extrapolation, *Journal of Financial Economics* forthcoming.

Kim, Don H, and Athanasios Orphanides, 2012, Term structure estimation with survey data on interest rate forecasts, *Journal of Financial and Quantitative Analysis* 47, 241–272.

Koijen, Ralph S.J., Tobias J. Moskowitz, Lasse Heje Pedersen, and Evert B. Vrugt, 2018, Carry, *Journal of Financial Economics* 127, 197 – 225.

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

Lettau, Martin, and Sydney Ludvigson, 2001, Consumption, Aggregate Wealth, and Expected Stock Returns, *The Journal of Finance* 56, 815–849.

Lettau, Martin, and Sydney C. Ludvigson, 2010, Chapter 11-measuring and modeling variation in the risk-return trade-off, in Yacine Ait-Sahalia, and Lars Peter Hansen, eds., *Handbook of Financial Econometrics: Tools and Techniques*, volume 1 of *Handbooks in Finance*, 617–690 (North-Holland, San Diego).

Liu, Yan, and Jing Cynthia Wu, 2021, Reconstructing the Yield Curve, *Journal of Financial Economics* 142, 1395–1425.

Lochstoer, Lars A, and Tyler Muir, 2019, Volatility expectations and returns, *Journal of Finance* forthcoming.

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

Lustig, Hanno, Nikolai Roussanov, and Adrien Verdelhan, 2014, Countercyclical Currency Risk Premia, *Journal of Financial Economics* 111, 527–553.

Manela, Asaf, and Alan Moreira, 2017, News Implied Volatility and Disaster Concerns, *Journal of Financial Economics* 123, 137–162.

Nagel, Stefan, and Zhengyang Xu, 2022, Asset pricing with fading memory, *Review of Financial Studies* forthcoming.

Nicholls, D. F., and A. L. Pope, 1988, Bias in the Estimation of Multivariate Autoregressions, *Australian Journal of Statistics* 30A, 296–309.

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

Politis, Dimitris N., and Halbert White, 2004, Automatic Block-Length Selection for the Dependent Bootstrap, *Econometric Reviews* 23, 53–70.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2010, Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy, *The Review of Financial Studies* 23, 821–862.

Shiller, Robert J., 2000, Measuring Bubble Expectations and Investor Confidence, *Journal of Psychology and Financial Markets* 1, 49–60.

Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.

Sun, Yixiao, 2013, A Heteroskedasticity and Autocorrelation Robust F Test using an Orthonormal Series Variance Estimator, *The Econometrics Journal* 16, 1–26.

Wang, Renxuan, 2021, Subjective return expectations, Technical report, Columbia University.

Welch, Ivo, and Amit Goyal, 2008, A Comprehensive Look at The Empirical Performance of Equity Premium Prediction, *The Review of Financial Studies* 21, 1455–1508.

Wu, Liuren, 2018, Estimating Risk-Return Relations with Analysts Price Targets, *Journal of Banking & Finance* 93, 183–197.