What Drives the Volatility of Professional Stock Return Forecasts? Causal Evidence from Macro Shocks☆

Martijn Boons* Giorgio Ottonello† Rossen Valkanov‡

Nova SBE Nova SBE UCSD

Abstract

Consensus professional stock return forecasts are three times more volatile than those of non-professionals and econometricians. This volatility stems from professionals' countercyclical responses to macro-shocks, with different shocks each accounting for 20-40% of the variation in forecast differences. We introduce a two-stage procedure to identify the discount rate variation in professional forecasts due to shocks and find it aligns well with realized returns and implications of rational asset pricing models. We conclude that professionals' assessment of the discount rate impact of macro-shocks distinguishes them from other stock return forecasts and forecasts of macro-variables, which challenges models of expectation formation.

Keywords: Macroeconomic Shocks, Stock Return Forecasts, Time-Varying Discount Rates, Real Uncertainty, Subjective versus Objective Forecasts

JEL Classification: E44, G12

*Nova School of Business and Economics, Rua Holanda 1, 2775-405 Carcavelos, Portugal. E-mail: martijn.boons@novasbe.pt Web: https://sites.google.com/site/martijnboons/

†Nova School of Business and Economics, Rua Holanda 1, 2775-405 Carcavelos, Portugal. E-mail: giorgio.ottonello@novasbe.pt Web: https://sites.google.com/view/gi8nello/home/

‡Corresponding author. University of California, San Diego. E-mail: rvalkanov@ucsd.edu Web: https://rady.ucsd.edu/people/faculty/valkanov/

☆ We are grateful to Mikhail Chernov, Jun Liu, Christian Wagner (discussant), Neng Wang, Paul Whelan (discussant), seminar participants at Nova SBE, Tilburg University, SWUFE, Vienna Graduate School of Finance, LUISS, University of Duisburg as well as participants at the CIREQ-CMP Econometrics Conference in Honor of Eric Ghysels, FIRS 2024 in Berlin, and Nova-WU Vienna Finance Workshop for helpful comments. This work was funded by Fundação para a Ciência e a Tecnologia (UIDB/00124/2020, UIDP/00124/2020 and Social Sciences DataLab - PINFRA/22209/2016), POR Lisboa and POR Norte (Social Sciences DataLab, PINFRA/22209/2016) A previous version of the paper was titled "Can Macro Shocks Explain the Excess Volatility in Professional Stock Return Forecasts?"

1. Introduction

There is a growing interest in understanding the macroeconomic drivers of survey-based, subjective economic forecasts. If the beliefs of surveyed forecasters incorporate macroeconomic shocks – representing fundamental, exogenous fluctuations that shape the future state of the economy (Ramey, 2016; Stock and Watson, 2016) – these shocks ought to influence their forecasts of macroeconomic and financial variables. However, the recent literature on survey forecasts of macroeconomic variables, such as GDP growth and inflation, highlights a key stylized fact: the consensus forecast exhibits low volatility and is excessively smooth, such that it responds to shocks more gradually than the realization.[1] This holds true for subjective forecasts made by professionals as much as it applies to those made by non-professionals.[2] Objective forecasts made by econometricians are not significantly more volatile either, primarily because the best real-time predictions rely heavily on shrinkage.[3]

In this paper, we examine forecasts of stock returns and present evidence that challenges this stylized fact: professional forecasts of the aggregate market return are three times more volatile than both non-professional and objective forecasts. Our goal is to explore the volatility differences between professional stock return forecasts and (i) other stock return forecasts and (ii) macroeconomic forecasts, with a particular focus on how fundamental macro shocks contribute to fluctuations in professional forecasters' beliefs about the stock market. Specifically, do professionals update more aggressively in response to macroeconomic news than do other forecasters? If so, is it because professionals better assess the countercyclical impact of macro shocks on risk premia? Is the professionals' forecast therefore relatively more accurate following negative shocks, that is, in bad times when the economic value of a good forecast is plausibly much larger? In answering these questions, we are the first to quantify pro-

[1] As a consequence, forecast errors are positively autocorrelated (see, e.g., Farmer, Nakamura, and Steinsson, 2022) and predictable by forecast revisions (see, e.g., Coibion and Gorodnichenko, 2015).

[2] For instance, Coibion and Gorodnichenko (2012) show that inflation forecasts made by professionals, central bankers and consumers (or non-professionals) respond similarly to a variety of economic shocks. Andre et al. (2022) present experimental evidence confirming the same insight.

[3] For instance, Bianchi, Ludvigson, and Ma (2022) apply state-of-the-art machine learning techniques to construct objective forecasts for inflation and GDP. Their objective forecast is only slightly more volatile, and responds only slightly more strongly to a business cycle shock, than the professional forecast.

fessional forecasters' subjective assessments of the macroeconomic origins of time-variation in expected stock returns and to compare these subjective assessments to more objective forecasts derived from an econometric framework, economic models, and realized returns. At risk of preempting our results, we find that professionals' beliefs respond significantly and countercyclically to macro shocks and thus play an important role in shaping market participants' expectations about equilibrium prices.

To be precise, we study next year's aggregate stock market return in excess of the one year T-bill and three ex ante or real-time forecasts of this quantity: professional, non-professional and objective. Our professional forecast comes from the Livingston survey, which queries respondents with formal training in economic forecasting.[4] Our non-professional forecast comes from Nagel and Xu (2022) and is based on a representative sample of U.S. individuals. The objective forecast is a combination forecast constructed using an extended set of stock return predictors.[5] This combination forecast is interesting as a benchmark for professionals because it obtains an out-of-sample R^2 of 8%, which is on par with state-of-the art approaches (see, e.g., Gu, Kelly, and Xiu, 2020; Liao et al., 2023).[6]

Over a sample period from 1952 to 2020, we find that the volatility of the professional forecast is about 6% (annualized), which is similar to the variation in expected stock returns implied by a full sample regression on the price-dividend ratio (see, e.g., Cochrane, 2011, Table I). In contrast, the volatility of the non-professional and objective forecasts is less than 2%. To distinguish between alternative explanations for the large difference in volatility of these forecasts, we use three well-identified macro shocks to trace out their fundamental drivers: the business cycle shock of Angeletos, Collard, and Dellas (2020), the TFP news shock of Kurmann and Otrok (2013) and Kurmann and Sims (2021), and the oil supply

[4]We study professional forecasts of long-term stock returns from the SPF survey in a robustness check and find that these forecasts agree (both qualitatively and quantitatively) with professional forecasts of short-term returns from the Livingston survey.

[5]Our set of predictors is similar to those studied in Goyal and Welch (2007). We show that our conclusions are robust to alternative objective forecasts, such as Martin (2017) option-implied lower bound or an objective forecast filtered through the present value restriction.

[6]The professionals' out-of-sample R^2 is lower unconditionally at 2.4%. As detailed below, we are most interested in understanding whether the professionals' relative forecasting performance is time-varying (and we find that it is).

shock of Baumeister and Hamilton (2019). Because these shocks are plausibly exogenous to current and past stock market conditions, we can interpret their impact on forecasts as a consequence of their impact on macroeconomic fundamentals. Indeed, these shocks are relevant instruments because they cause significant fluctuations in macroeconomic activity, and as such, they speak directly to the cyclicality of subjective forecasts, an area that is hotly debated since the seminal work of Greenwood and Shleifer (2014).[7]

Our macro shock-based approach differs significantly from the standard predictive regression approach widely used in the literature. These regressions rely on lagged state variables to proxy for time-variation in economic conditions. While these regressions are generally useful for documenting a predictive relationship, they lack clear identification. As a result, linking periods of, for instance, high price-dividend ratios to an underlying economic shock requires additional assumptions.[8] In reality, such periods likely result from a combination of shocks propagating through the economy and financial markets. By contrast, our study focuses on exogenous macroeconomic shocks that hold many confounding factors constant, facilitating a causal inference of their effect on subjective perceptions of risk.

We find using Jordà (2005) local projections that professional stock return forecasts increase by a significant 2 percentage points (at the peak) after a one standard deviation contractionary macro shock of each type. This countercyclical response is economically large. An increase of 2 percentage points represents about a third of what is often argued to be the standard deviation of true expected returns (see, e.g., Cochrane, 2011). In stark contrast to professionals, the response of non-professional and objective forecasts to all three shocks is economically small: it is below 20 basis points in most cases and sometimes even negative. Using the variance decomposition approach of Gorodnichenko and Lee (2020), we find that each shock captures a substantial fraction of the historical variation in professional-minus-non-professional and professional-minus-objective forecast differences. To be precise,

[7] Recent contributions include Møller, Pedersen, and Steffensen (2020); De La O and Myers (2021); Nagel and Xu (2023); Dahlquist and Ibert (2024); Couts, Gonçalves, and Loudis (2023); Gandhi, Gormsen, and Lazarus (2023); Jensen (2023); Bastianello (2022).

[8] There is a long literature on whether a higher than average dividend price ratio captures fluctuations in risk premia or expected growth rates in fundamentals (see Koijen and Van Nieuwerburgh, 2011; Nagel, 2024, for excellent recent reviews).

the variance contribution of the business cycle, TFP, and oil supply shock, respectively, peaks at about 40%, 30%, and 20%, at horizons beyond one year. These results imply that macro shocks are a key determinant of the more volatile stock return forecasts made by professionals.

Because the professionals' strong countercyclical response to macro shocks is consistent with the dynamics of time-varying risk premia in rational asset pricing models, we argue that our results are driven by a discount rate channel. Specifically, professionals incorporate into their return forecasting model real-time information that reflects the impact of business cycle fluctuations, TFP shocks, and oil supply shocks. They recognize that negative macroeconomic news increases aggregate risk, prompting them to forecast higher returns even before fully assessing the adverse impact on the macroeconomy and cash flows. While these shocks are observable to us as econometricians only ex post, they serve as instruments for the true (but unobservable) information used by professional forecasters.

In further support of the discount rate channel, we show that professional forecasts of aggregate cash flows – GDP growth (studied also in Bianchi, Ludvigson, and Ma, 2022) and dividend growth (derived from the present-value identity using a state-space model) – exhibit a relatively muted and, if anything, positive response to contractionary macro shocks. This result obtains even though both realized GDP and dividends drop significantly after each shock. Thus, even if professionals responded to cash flow news only with a delay, a cash flow channel would imply lower, not higher, professional stock return forecasts.

To more precisely quantify our proposed discount rate channel, we develop a two-stage procedure to identify exogenous variation in commonly used discount rate proxies that is due to fundamental business cycle, TFP, and oil supply shocks. In the first stage, we estimate the impact of these shocks on a given proxy for discount rates. In the second stage, we evaluate how forecasts of returns respond to each component of the discount rate proxy, that is, the component instrumented with the shocks versus the residual component. For time-varying discount rates to be a major driver of the stock return forecasts of professionals, the macro shocks must be strong instruments in the first stage. That is indeed what we find.

Our first discount rate proxy is real uncertainty. Because real uncertainty relates directly to the problem of forecasting the macroeconomy and a large literature suggests that real uncertainty is priced, this proxy provides an economically useful example of our approach. We find that business cycle and TFP news shocks, in particular, have a large impact on real uncertainty, which we measure either using the real uncertainty index of Ludvigson, Ma, and Ng (2021) or the dispersion in professional GDP forecasts of Bloom (2009). While existing theory agrees on the sign of the endogenous response of uncertainty to bad macro shocks (and therefore the response of risk premia), the residual (or non-macro) components of uncertainty may relate to returns in complicated ways (Ludvigson, Ma, and Ng, 2021). We find that the first stage is strong: a one standard deviation business cycle or TFP news shock increases uncertainty by about 0.25 standard deviations and together the two shocks comfortably pass the robust test for weak instruments of Montiel Olea and Pflueger (2013). In line with the theory, we find in the second stage that professionals increase their stock return forecasts by a significant 6 to 7 percentage points for a one standard deviation increase in instrumented uncertainty. In contrast, the same increase in the residual component of uncertainty leaves their forecasts largely unaffected.

Running the same two-stage estimation for realized returns, we show that professionals' beliefs about the relation between shocks, real uncertainty, and returns are in line with the ex post truth. Realized returns respond by about the same amount as professional forecasts to instrumented uncertainty. The difference between the two responses (measuring the professionals' forecast error) is small and insignificant at about 1 percentage point. Under the strong assumption that the macro shocks impact (forecasts of) returns only because they impact uncertainty (i.e., the shocks satisfy the exclusion restriction required for 2SLS), one can interpret this result as showing that professionals understand the causal impact of real uncertainty on returns. If the exclusion restriction is strictly violated, this result is interesting nonetheless. It shows that the response of professional forecasts to shocks is qualitatively and quantitatively consistent with the impact of these shocks on real uncertainty and realized stock market returns. In other words, the professionals' forecast is accurate in the most uncertain times caused by large macro shocks.

We next perform the same two-stage estimation for the entire set of 22 stock return predictors used to construct our objective forecast. While eight of these predictors respond strongly to shocks (such as the default spread, dividend yield, and investment-to-capital ratio), the remainder does not. For the eight "shock-sensitive" predictors, we find that professional forecasts load on each predictor with similar sign and magnitude as realized returns. Moreover, this loading nearly doubles when we focus on the component of the predictor due to shocks, again broadly consistent with realized returns. For instance, both the professionals' forecast and realized returns increase by about 3.5% for a one standard deviation increase in the default spread, one of the shock-sensitive predictors. If this increase in the default spread is due to macro shocks as measured by the coefficient on the instrumented default spread in the second stage, both realized and forecasted returns increase by a significant 9 percentage points. In contrast, both are largely insensitive to the residual component of the default spread. Furthermore, both the non-professionals' and objective forecast are largely insensitive to all the predictors we study, even if we focus only on the shock-sensitive ones. As a result, forecast errors are strongly predictable for the non-professional and objective forecasts. In other words, these alternative forecasts are relatively inaccurate in bad times caused by adverse macro shocks.

In sum, professional forecasters are unique in that they understand how returns relate to variation in predictors caused by macro shocks. In turn, professionals do not understand well how returns relate to predictors for which most of the variation is due to things other than our macro shocks. We thus shed new light on the "cyclicality gap" of Nagel and Xu (2023), which refers to their finding that realized returns are much more sensitive ex post to the average stock return predictor than forecasts (see also Dahlquist and Ibert, 2024). We replicate this finding, but show that there is large variation in the cyclicality gap for professional forecasts depending on how strongly a predictor responds to macro shocks. The gap is largest for shock-insensitive predictors, whereas for the eight predictors that are sensitive to macro shocks, there is virtually no gap.

This paper contributes to an expanding body of work on the dynamics of stock returns and their forecasts. In their seminal work, Greenwood and Shleifer (2014) find that stock return forecasts by a variety of agents (such as individual investors, CFOs and consumers) are counterfactually procyclical. Recent work by Møller, Pedersen, and Steffensen (2020), Dahlquist and Ibert (2024), and Couts, Gonçalves, and Loudis (2023) shows that forecasts by more professional agents, such as the analysts in the Livingston survey and institutional investors or money managers, are higher in NBER recessions and when valuation ratios are low.[9] We contribute to this literature by identifying a root cause of the strong, countercyclical variation of professional stock return forecasts: macro shocks. The main advantage of our approach is that it accommodates comparison to a rational benchmark, that is, the impact of a discount rate shock in economic models (see Cochrane, 2011, 2017, for an extensive review). It turns out that the professionals' forecast is unique in that it is surprisingly close to this benchmark: when contractionary shocks hit, their return prediction is well above average at around 10% per year, consistent with realized returns. We conclude that professionals' beliefs have more practical relevance than what existing work on subjective forecasts may seem to suggest (see Greenwood and Shleifer, 2014; Adam, Matveev, and Nagel, 2021; Nagel and Xu, 2023).

Additional advantages of our approach include the ability to revisit the cyclicality gap and to precisely quantify the importance of macro shocks for explaining the time-variation in forecast differences. We argue that the quantitative assessment of the countercyclical discount rate impact of macro shocks reveals a significant wedge between different forecast methods, which contributes to a growing literature on the expectations formation process.[10] This discount rate impact is not reflected in professional forecasts of macro variables or cash flows nor in non-professional or objective forecasts of stock returns. Our benchmark objective stock return forecast conditions only on lower frequency predictors, which definition is consistent with many other stock return forecasts studied in prior literature. Thus, our

[9] Another related paper is Gandhi, Gormsen, and Lazarus (2023), who study "forward" expected returns (derived from option prices and professional surveys) and show that these are countercyclical as well.

[10] See, e.g., Patton and Timmermann (2011); Coibion and Gorodnichenko (2012, 2015); Bordalo et al. (2020); Angeletos, Huo, and Sastry (2021); Farmer, Nakamura, and Steinsson (2022).

findings raise important concerns about using these types of objective forecasts as the rational benchmark for subjective forecasts, which is proposed in recent work by Bianchi, Ludvigson, and Ma (2022) and Nagel and Xu (2023). In a setting with shocks, these objective forecasts actually appear too smooth when compared to predictions of rational models as well as realized returns.[11]

The remainder of the paper is structured as follows. In Section 2, we describe the data and highlight the high volatility in professional stock return forecasts. Section 3 investigates the countercyclical response of these forecasts to macroeconomic shocks and quantifies the proportion of forecast variation attributable to them. In Section 4, we introduce a novel two-stage approach for assessing the impact of macroeconomic shocks on professionals' forecasts through the discount rate channel. Section 5 explores in detail the wedge between objective forecasts, filtered via present-value restrictions, and subjective forecasts. Section 6 concludes.

2. Data and motivating evidence

In this section, we introduce our stock return forecasts and the macroeconomic shocks that are the main focus of our paper. To fix ideas, write the simple return of the aggregate market in excess of the risk free rate as

$$R_{t+1} = \mu_t + \varepsilon_{t+1}, \tag{1}$$

where under the null hypothesis of Full Information Rational Expectations (FIRE), $\mu_t = E(R_{t+1}|I_t)$ is the expected excess return conditional on the information set I_t and ε_{t+1} is mean-zero white noise. μ_t is unobservable and two fundamentally different approaches to generate an ex ante estimate or real-time assessment of this quantity are popular in the literature.

[11]Objective expected returns implied by the Campbell and Shiller (1988a) present-value identity and filtered using a state-space model following Van Binsbergen and Koijen (2010) reflect the information set of the professional forecaster slightly better, because we find that these do respond meaningfully to the shocks. Yet, at about one-third of the response of professional forecasts, the response of these filtered expected returns is still relatively weak.

The first approach is to use an econometric model to construct objective forecasts. This approach is transparent, because the model's assumptions (e.g., functional form of the estimated relationship, parameter space, and set of conditioning variables) are clearly defined. It is the most commonly used approach in the literature and there is a wealth of research available to draw from. Well-known potential drawbacks include model uncertainty, estimation error, and over-fitting.

A more recent approach is to survey market participants about their expectations of future returns, thus capturing the forecasts of professional investors, analysts, economists, and consumers among others. These forecasts, however, are subjective and may depend on survey participants' background, expertise, and biases rather than actual data. For instance, professional forecasters may produce different predictions than consumers who are non-professionals. Moreover, unlike the econometric approach, we do not know the exact model survey participants use to form their forecasts.

2.1 Stock market return forecasts

We study three types of return forecasts: a survey-based professional forecast, $F_t^P(R_{t+1})$, a survey-based non-professional forecast, $F_t^{NP}(R_{t+1})$, and an objective forecast based on an econometric model, $F_t^O(R_{t+1})$. Throughout, $t+1$ refers to t plus one year and $t+\frac{1}{4}$ refers to t plus one quarter. All forecasts are formed in a quarter t for the one-year ahead excess stock market return in $t+1$.

2.1.a Professionals

Forecasts of professionals come from the biannual Federal Reserve Bank of Philadelphia's Livingston Survey, which covers the period from June 1952 to December 2020. The participants in this survey are economists working in commercial banks, the federal reserve, academia, government or consulting; have training in economic theory and forecasting; and, are presumed to use econometric models to generate these same forecasts for their employers or clients. Each second and fourth quarter of the year, participants are asked to forecast the

S&P500 index 6 months and 1 year ahead (i.e, 6 and 12 months after the end of the quarter in which the survey is administered). From 1992 onwards, forecasters are also required to report the "nowcast" of the index. Following Nagel and Xu (2022), we construct a time series of forecasts for the one-year ahead excess stock return from 1952 to 1991 by computing the annualized ratio of the consensus mean 12-month and 6-month ahead index forecasts. From 1992 onwards, we use the ratio of the 12-month ahead index level forecast and nowcast.[12] We thus convert forecasts of levels to growth rates exactly as is done in the macro literature that studies forecasts of GDP and inflation.[13] These stock return forecasts from the Livingston survey are well-populated: the fraction of all panelists that provide a forecast of GDP that also provide a forecast of stock returns is quite constant over time at about 65%. We finally subtract the one-year Treasury yield to obtain the excess return forecast, $F_t^P(R_{t+1})$.

2.1.b Non-Professionals

There are a number of surveys of non-professionals, that is, survey participants for which the likelihood of formal training in economics and forecasting is low. Unfortunately, no single survey spans a long enough period nor offers a large enough coverage of participants. Hence, we use the quarterly time series of consensus non-professional one-year ahead excess (CRSP) index return forecasts from Nagel and Xu (2022, Section 1.3) in excess of the one-year Treasury yield. This 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. Together, these forecasts reflect the expectations of a representative sample of individual investors or consumers in the U.S. The Nagel and Xu (2022) series, $F_t^{NP}(R_{t+1})$, has quarterly continuous coverage from June 1987 and ends in December 2020.

[12]Our results are virtually identical when we continue to use only the 6- and 12-month ahead forecasts post 1992.

[13]Although this does not affect any of our conclusions, we follow previous literature (see, e.g., Nagel and Xu, 2022) and add dividends to these price growth forecasts using the formula $E_t^P(P_{t+1}/P_t) + (D_t/P_t)E(D_{t+1}/D_t) - 1$, setting $E(D_{t+1}/D_t) = 1.064$, which is the sample average of S&P annual dividend growth over the post-WWII sample period.

2.1.c Objective forecasts

The goal is to produce an econometrician's objective forecast of aggregate stock returns. Our main measure is based on the standard linear predictive regression framework that is most popular in the literature (see, e.g., Goyal and Welch, 2007). To start, we run for each predictor X_i:

$$R_{s+1} = \alpha_{i,t} + \beta_{i,t} X_{i,s} + e_{i,s}, \tag{2}$$

over an expanding window of monthly or quarterly observations (depending on availability of the predictor) from $s = 0$ to $t - 1$. We use 10 years as burn-in period and construct out-of-sample forecasts for the remaining sample until December 2020. Rapach, Strauss, and Zhou (2009) show that combination forecasts provide a simple way to aggregate information from many predictors with superior out-of-sample performance. We follow their approach and define the objective forecast as the average of out-of-sample forecasts derived from Eq. (2):

$$F_t^O(R_{t+1}) = \frac{1}{N_t} \sum_{i=1}^{N_t} (\widehat{\alpha_{i,t}} + \widehat{\beta_{i,t}} X_{i,t}) \tag{3}$$

where N_t is the number of predictors for which a forecast can be generated at time t.

For this exercise, we start from the set of predictors studied in (Goyal and Welch, 2007). We add to this set predictors that feature prominently in the recent literature: the real factor from Ludvigson and Ng (2009, $F1$), a long-run exponential average of past per-capita real aggregate dividend growth from Nagel and Xu (2022, $EXPD$), the square of the VIX ($VIX2$), industrial production growth (IP) and GDP growth (GDP). For a detailed description of the predictors and their construction see Online Appendix A.

We also consider two alternative objective forecasts of stock returns. The first is the lower bound for the equity premium of Martin (2017), based on risk-neutral variance derived from S&P500 option prices. Although this forecast does not rely on an econometric model, the exact interpretation of the measure requires several economic assumptions. The second alternative is objective expected returns implied by the Campbell and Shiller (1988a) present-value identity and filtered using a state-space model as in Van Binsbergen and Koijen (2010, see Section 5 below). These objective forecasts are not our main focus because of power considerations: the lower bound is only available for a short sample from 1996 onwards and the state-space model is estimated using annual data, which precludes us from identifying higher frequency variation. In addition, filtered series from a state-space model do not provide a truly out-of-sample forecast.

2.1.d Summary statistics

We present summary statistics for realized and forecasted returns in Table I. We display the time-series of the forecasts in Figure 1, together with NBER recession periods. Three facts stand out.

Fact 1: The professionals' forecast is relatively volatile. The standard deviation for professionals equals 5.9%, which is more than three times the standard deviation of the non-professional and regression-based objective forecast: less than 1.9%. Panel B shows that this large difference in volatility is not due to a scaling effect: the forecasts of professionals are only weakly correlated with those of non-professionals ($corr = 0.20$) as well as the objective forecast ($corr = 0.10$). Professional forecasts are volatile also relative to the alternative objective forecast of Martin (2017). Over their common sample from 1996 to 2012, the volatility of Martin's lower bound for the one year equity premium is 2.4%, compared to 7.1% for the professional forecast. We also note that Martin's measure is highly correlated with our objective forecast (at about 0.67). The relatively low volatility of the non-professional forecast is robust to another popular measure based on the CFO survey.[14] Understanding whether this high volatility is just noise or due to a precise signal incorporated in the professionals' forecast is the goal of our study. Since the cross-section of professional forecasters (about 21 on average) is much smaller than the cross-section of non-professionals (more than 1,000 respondents on average), the volatility of the professional forecast would be higher even if

[14] We view these CFOs as non-professionals, because constructing econometric models to forecast stock market returns is unlikely to be part of their job description. Over their common sample from 2000 to 2020, the volatility of our non-professional forecast is 1.68%, while it is 1.75% for the CFO survey. We thus conclude that our result is not due to specific choices made in the construction of the non-professional series in Nagel and Xu (2022).

they have the exact same model as non-professionals. However, if professional forecasts respond differently to shocks, this would be strong evidence in favor of the hypothesis that their model is different.

Fact 2: The professionals' forecast is strongly countercyclical. Even though the cyclicality of subjective forecasts is hotly debated since the seminal work of Greenwood and Shleifer (2014), this fact clearly distinguishes the professionals from the non-professionals. Indeed, the non-professionals' forecast is procyclical. To summarize the cyclicality, we present in Panel C of Table I average (forecasts of) returns during NBER recessions versus expansions. After an NBER recession quarter, realized returns are about 7% above average at 12.5%. Professional forecasts are closest to this reality and are about 9% above their unconditional mean at 12.0%. Non-professional and objective forecasts do not capture this strong countercyclicality, as the non-professional forecast is about 2% below its mean and the objective forecast is only 1% above the mean in NBER recessions.

Fact 3: The average forecast of professionals is relatively low. While the average professional forecast equals 4.2%, the average non-professional forecast equals 6%. The average objective forecast is even higher at 8.3%, which closely matches the average realized return of 8.5%. Given these differences, it is perhaps unsurprising that the objective forecast outperforms the professionals' forecast in out-of-sample R^2 at 8.0% versus 2.4%. That said, the professionals do outperform the non-professionals, as their out-of-sample R^2 is negative. Although several explanations for a downward bias of professionals exist in the literature, we see in Figure 1 that the bias is mainly due to a number of large negative forecasts before the 2000s.[15] First, note that all return forecasts are in excess of the one-year Treasury yield. Since this yield is not risk-free, a negative forecast does not immediately violate the idea that the aggregate stock market risk premium must be positive. Second, while our objective forecast

[15] Consistent with heterogeneity in priors (Patton and Timmermann, 2010), beliefs of professionals about long-run expected returns may be below non-professionals and, in fact, below realized average returns, leading to shrinkage. Indeed, Fama and French (2002) and many others in the literature have argued that expected returns derived from discounted cash flow techniques are considerably lower than realized average returns. Asymmetry in the professionals objective function may also lead to shrinkage. Elliott, Komunjer, and Timmermann (2008) argue that agents are averse to "bad" outcomes such as lower-than-expected returns and therefore incorporate such loss aversion into their forecasts.

is never negative, previous work has shown that alternative objective forecasts are also negative at times (see, e.g., Campbell and Thompson, 2008; Pettenuzzo, Timmermann, and Valkanov, 2014). Third, we show in Online Appendix B that professionals' negative excess return forecasts are relatively poor in terms of out-of-sample R^2 and relatively disconnected from the macro shocks that we study in this paper. In this appendix, we also show that the professionals' forecasts outside of the left tail provide a much better fit to ex post realized returns than the objective forecast.

2.2 Shocks

To identify the root causes of the cyclicality of the professionals' stock return forecasts, we use well-identified and well-established shocks from the macroeconomics literature. Following the approach of Boons, Ottonello, and Valkanov (2023), we study three different shocks that are available over a relatively long period that best matches the sample of forecasts: business cycle shocks from Angeletos, Collard, and Dellas (2020), TFP news from Kurmann and Otrok (2013), and oil supply shocks from Baumeister and Hamilton (2019). Figure F.1 in the Online Appendix presents their time series. Previous literature finds that these three primitive shocks explain a large fraction of future macroeconomic activity and they are plausibly exogenous to current and lagged conditions in the stock market. The shocks hold many confounding factors constant and thus come closer to the fundamental innovations that drive risk premia in classical asset pricing models. Let us provide more detail about each shock.

2.2.a Business cycle shock

We study the "main business cycle shock" of (Angeletos, Collard, and Dellas, 2020, denoted Cycle), which is derived from a structural VAR with ten variables and identified using the max-share identification strategy of Barsky and Sims (2011). To be precise, the shock we use is the one that maximizes the contribution to business cycle fluctuations in output (GDP) over the sample period from 1955Q3 to 2017Q4. It turns out that this shock is nearly indistinguishable from the shock constructed by maximizing the contribution to any of the following variables: unemployment, employment, investment, and consumption. Because this shock targets variation at business cycle frequency, it is essentially orthogonal to the much longer-term TFP shocks we discuss next.

2.2.b TFP news

Since the influential paper of Beaudry and Portier (2006) it is well established that news about future productivity can be an important driver of macroeconomic activity.[16] We use the TFP shocks derived from a structural VAR as in Kurmann and Otrok (2013), but using the improved identification method proposed in Kurmann and Sims (2021). Thus, our TFP news shock is associated with the maximum forecast error variance contribution to TFP at long horizons of 80 quarters. TFP news shocks are attractive as they appear in major asset pricing models, as a shock to long run growth. The sample period is 1962Q1 to 2019Q4.

2.2.c Oil supply shock

Hamilton (1983) argues that oil supply shocks are a major driver of economic fluctuations and this argument has inspired a large literature studying their macroeconomic effects. For our analysis, we use the oil supply shocks of Baumeister and Hamilton (2019). These shocks are extracted from a Bayesian VAR that relaxes the strong assumption in prior literature that there is no short-run response of oil supply or production to the oil price. Moreover, the setup of Baumeister and Hamilton (2019) allows them to use more data than in previous work (both a longer time series and additional information on supply and demand elasticities). The sample period is 1975Q1 to 2019Q4 and the largest shocks occur during the uncertainty about OPEC production in 1986 and at the start of the Gulf war in 1990Q3. Baumeister and Hamilton (2019) show that the oil shock mainly impacts macroeconomic activity at horizons from one to two years.

Throughout the paper, we sign the shocks so that an increase is bad news and associated

[16] Our conclusions are robust to using a prominent alternative technology shock, that is, the IST news shock of Ben Zeev and Khan (2015). See Figure F.4 in the Appendix.

to lower future economic growth. Important features of our identification strategy are that (i) the shocks are different in nature, (ii) they hit the economy at different times while the economy responds at different horizons, and (iii) they are only observable ex post, because the shocks are structural innovations derived from a VAR that is estimated over the full sample. Requiring a consistently strong, countercylical response of stock return forecasts to these different shocks raises the bar considerably for us to claim that any agent makes forecasts that are in line with predictions from rational models.

3. How do different forecasts respond to macro shocks?

In this section, we study the impulse response function (IRF) of different forecasts to macro shocks. We also ask how important the shocks are quantitatively for capturing historical variation in the forecasts through a forecast error variance decomposition (FEVD).

Given that the professional forecast is relatively large in recessions (see Table I), it may not seem surprising if professional forecasts respond more to these shocks than the non-professional and objective forecasts. Previous literature suggests a very different null hypothesis, however. Existing work shows that forecasts of macro variables made by a large variety of agents respond similarly and weakly to shocks. For instance, Coibion and Gorodnichenko (2012) arrive at this conclusion studying inflation forecasts of professionals, consumers, firm managers, and central bankers in conjunction with technology (news) and oil shocks. Bianchi, Ludvigson, and Ma (2022) show that professional and machine-learning based objective forecasts of inflation and GDP respond similarly and weakly (relative to the response of realized inflation and GDP) to the business cycle shock. We replicate their findings for professional forecasts and extend the evidence to TFP news and oil supply shocks, as explained in Online Appendix D. We report in Figure 2 the impulse response function for realized GDP and professional forecasts to all three shocks. We consider forecasts from the Livingston survey, which is the focus in our paper, and the SPF survey, which is the focus in Bianchi, Ludvigson, and Ma (2022). In short, we see that professionals strongly underreact to each shock: the response of forecasts is small relative to the large negative response of realized GDP.

These weak responses are often argued to be consistent with information rigidity (Coibion and Gorodnichenko, 2012, 2015) and Bianchi, Ludvigson, and Ma (2022) argue that such information rigidity may derive from the shocks being unobservable in real-time. Thus, any large differences across agents in the response of return forecasts to shocks would be interesting to this literature and would shed important light on the nature of the forecasting problem for stock returns vis-à-vis macro variables. On one hand, realized stock returns are noisy estimates of true expected returns, so one may be inclined to believe that forecasting in real-time the response of returns to shocks is relatively hard. On the other hand, theory suggests that risk aversion or fundamentals, more generally, drive discount rates, and these drivers may respond swiftly to a shock even if the forecaster does not know yet what the precise impact on the macroeconomy will be.

3.1 Impulse response function

We measure the dynamic response of stock return forecasts using the local projection method of Jordà (2005). To this end, we estimate the following linear regression:

$$Y_{t+K+1} = \alpha_K + \beta_K N_t + \sum_l \gamma_l C_{t-l} + \epsilon_{t+K+1}, \tag{4}$$

where Y_{t+K+1} is one of the three stock return forecasts made in $t+K$: $F^P_{t+K}(R_{t+K+1})$, $F^{NP}_{t+K}(R_{t+K+1})$ or $F^O_{t+K}(R_{t+K+1})$, and N_t is one of the three macro shocks: Cycle, TFP, or Oil. The vector C_{t-l} collects the controls and contains one year worth of lags of the dependent variable and the shock. We estimate the impulse response up to 4 years from the time the shock hits. Dictated by the frequency of the forecasts, we use a quarterly frequency for non-professional and objective forecasts (i.e., $K = 0, \frac{1}{4}, \frac{1}{2}, ..., 4$ and $l = \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1$) and a biannual frequency for professional forecasts ($K = 0, \frac{1}{2}, 1, ..., 4$ and $l = \frac{1}{2}, 1$). In the biannual case of professionals, N_t is the sum of two quarterly shocks. To make the IRFs comparable between the quarterly and biannual cases, we multiply the coefficients β_K in both cases by the standard deviation of the biannual shocks. We present these responses in Figure 3 and discuss the evidence for each forecast in turn.

Professionals: The first column of Figure 3 presents the response of professional forecasts to the three shocks. We see that the response of professionals is remarkably consistent: each shock causes an increase in forecasted stock returns that is economically large and significant at about 2 percentage points (or about one-third of the unconditional standard deviation of the professional forecast) at the peak. Given the chosen shock size, this number represents the standard deviation in the professional forecast due to a single macro shock. The response is fast for the Cycle and TFP shocks, with the peak impact occurring at an horizon of zero and two quarters, respectively. The response to Oil shocks is instead delayed, and significant at horizons from 1.5 to 2.5 years after the shock. The timing of these responses aligns quite well with the timing of the response of realized GDP growth (see Figure 2), consistent with the idea that professionals broadly understand when a shock is going to impact the economy (but not necessarily by how much). Further, in Online Appendix C we show that professional forecasts of long-term average stock returns from a different survey respond to the shocks in a way that is qualitatively and quantitatively consistent with what we find here. We conclude that professionals believe that expected returns are persistently high after bad macro-shocks.

Non-professionals and objective forecasts: The second and third column of Figure 3 show the responses of the non-professional and objective forecasts. The major insight is that these responses are small relative to the response of professionals. There is not a single response at any horizon in all six panels that is larger than 70 bps. In fact, at most horizons and for both the non-professional and the objective forecast, the responses are below 20 bps. In addition, some responses have the opposite sign to those of professionals. For instance, non-professionals respond negatively to TFP and Oil shocks. The objective forecast responds negatively to the Cycle shock. Although the responses of the objective and non-professional forecasts are much smaller in magnitude, the peak or trough impacts are sometimes (marginally) significant. This result obtains because the variance of these forecasts is much smaller than that of professionals. The response of professional forecasts is also strong relative to the response of Martin's (2017) lower bound for the equity risk premium, which is an alternative objective forecast. Figure F.6 presents impulse responses for these two forecasts over a shorter sample from 1996 onwards, dictated by availability of the lower bound. Although the lower bound also increases after each shock, the peak response is 3 to 4 times smaller than the response of the professional forecast.

Overall, the IRFs indicate that the response of professional forecasters to shocks is large, significant, and consistent with rational models in which the representative agent desires larger compensation for investing in risky assets when fundamentals worsen. Before discussing this mechanism in more detail, we analyze forecast error variance decompositions to determine the contribution of each shock to the historical variation in forecasts of the different agents.

3.2 Variance decomposition

We follow recent literature to judge the quality of subjective professional forecasts by comparing them to objective forecasts. For instance, Bianchi, Ludvigson, and Ma (2022) argue that an objective forecast "provides an appropriate benchmark for quantifying biases in survey responses" because it is designed to handle large amounts of information and adapts to new information when it becomes available, just like a professional analyst making her forecast. Nagel and Xu (2023) argue that the professional-minus-objective forecast difference is key to determine whether there is hope for asset pricing models to match the volatility of asset prices induced by shocks through time-varying risk premia, or whether models need to introduce a belief wedge. If the null is that forecasts do not respond strongly to shocks, one would also expect these shocks to contribute little to variation in this difference.

Using a standard mean squared error decomposition of professional and objective forecasts, we can decompose the difference between these two forecasts into bias and variance:

$$\underbrace{E((F_t^P(R_{t+1}) - F_t^O(R_{t+1}))^2)}_{MSE} =$$
$$\underbrace{(E(F_t^P(R_{t+1})) - E(F_t^O(R_{t+1})))^2}_{Bias} + \underbrace{Var(F_t^P(R_{t+1}) - F_t^O(R_{t+1}))}_{Variance}. \tag{5}$$

Intuitively, we focus on the difference in forecasts so that the impact of noisy realized returns on the difference in forecast errors cancels out. Over their common sample (from 1952 to 2020), we find that the mean squared error ($\times 10^4$) equals 53.31. The bias of the professional relative to the objective forecast accounts for about one third of the MSE, 17.65. The remaining two-thirds derives from the variance of the difference between the two forecasts, 35.66. Thus, even though professionals are biased, it is time-variation in the difference between the forecasts that contributes about two times as much to the MSE. The IRFs discussed above suggest that macro shocks are an important driver of the time-variation in the difference between these forecasts. We now quantify this contribution precisely using the FEVD approach for local projections developed by Gorodnichenko and Lee (2020).

The FEVD is based on a two step approach which we briefly describe below. First, we run a linear projection like in Eq. 4, but excluding the contemporaneous shock N_t:

$$F_{t+K}^P(R_{t+K+1}) - F_{t+K}^O(R_{t+K+1}) = \alpha_K + \sum_l \gamma_l C_{t-l} + \epsilon_{t+K+1}. \tag{6}$$

We run this regression at the biannual frequency, dictated by the availability of the professional forecast. Second, we estimate the contribution of contemporaneous and future shocks to the forecast error variance of the professional-minus-objective forecast at horizon K as the (uncentered) R^2 from the following regression:

$$\widehat{\epsilon_{t+K+1}} = \sum_{k=0}^{K} \beta_k N_{t+k} + \upsilon_{t+K+1}, \tag{7}$$

This R^2 estimates the fraction of the forecast error variance that is attributed to the exogenous shocks. We correct for small sample bias in the R^2 estimates following the VAR-based bootstrap approach of Gorodnichenko and Lee (2020).

We present the professional-minus-objective FEVD in the first column of Figure 4.[17] Consistent with the IRFs in Figure 3, all three shocks capture a substantial amount of historical variation in the difference between the two forecasts. The variance contribution of Cycle shocks is significant at all horizons and equals about 40% at horizons beyond six quarters. The variance contribution of TFP shocks is always significant as well and peaks at about 25% eight quarters out, after which it decays slowly. Finally, the variance contribution of Oil shocks is significant starting from an horizon of six quarters out and increases to about 20% four years out.

To sum up, each macro shock captures 20 to 40% of the unpredictable variation in the difference between professional and objective stock return forecasts at horizons beyond two years. We conclude that professionals update their forecasts aggressively (both in speed and magnitude) after large macro shocks. This finding implies that professionals feed into their return forecasting model real-time information that either directly or indirectly measures the impact of business cycle fluctuations, technology shocks, and oil supply shocks. This information is not contained or not effectively used in the benchmark objective forecast.

To highlight this unique feature of the model used by professionals, we present in the second column of panels the FEVD for the difference between the forecasts of professionals and non-professionals. By and large, these variance contributions are similar to what we saw for the professional-minus-objective difference. For completeness, we present the FEVD for the non-professional-minus-objective difference in the last column of the figure. Although these contributions are sometimes significant as well, one cannot put these results on the same footing as those for professionals relative to the objective forecasts. The reason is twofold. First, these variance contributions derive from responses of non-professionals to shocks that are quantitatively much smaller than the responses of professionals (see Figure 3). Second, for non-professionals these responses are sometimes procyclical, which is opposite

[17] These FEVDs for differences follow quite straightforwardly from the IRFs for the differences, which we report in Figure F.2. For completeness, we report the FEVDs for the individual forecasts in Figure F.3.

to professionals and to what rational models predict.

3.3 Key takeaways

Overall, professional forecasts of stock returns respond strongly and countercyclically to shocks consistent with rational asset pricing models. If this response was driven by professionals subjective assessment of time-varying discount rates, it would explain why the response of professionals' stock return forecasts is so different from the response of their forecasts of macro variables. Indeed, if professionals understand that shocks cause an increase in discount rates, they will forecast higher returns when a bad shock hits, even if they do not know the true extent of the adverse impact on the macroeconomy yet.

A simple sign restriction already suggests that our evidence is most consistent with this discount rate channel. Since the shocks cause economic growth to slow down, we would expect them to have a negative impact on future cash flows. We confirm this implication in Figure 5. Here we see that annual dividend growth (measured as log real per capita year-on-year growth in dividends and repurchases as in Nagel and Xu (2022)) drops after each shock and troughs at a significant -4% to -2%. Hence, if the professionals' response were due to changing beliefs about future cash flows, we would expect the return forecast to decrease, rather than increase.[18] In the next section, we test additional features of the discount rate channel.

Non-professionals' responses are much smaller and, if anything, slower. This evidence is potentially consistent with inattention models, where non-professional forecasters derive their expectations from news reports that incorporate the view of professional forecasters with a delay (Carroll, 2003). Alternatively, it could be that non-professionals mainly extrapolate past returns, whereas the professionals in the Livingston survey are much closer to the rational investor in benchmark asset pricing models.

Objective forecasts also do not react strongly to macroeconomic shocks. This fact is

[18] In Section 5 below, we impose the present-value restriction to show that professional forecasts of returns and observed dividend-to-price ratios together imply that professional forecasts of cash flows change by only a small amount after each shock.

perhaps unsurprising noting the shocks were not included in the dataset used to construct the objective forecast, consistent with many other objective forecasts studied in previous literature. For instance, for shocks to impact our combination forecast, they must first impact fundamentals as proxied by the predictor variables. However, the predictors are only noisy proxies of the true fundamental that matters for expected returns and, as argued in Nagel and Xu (2023), estimating the relation between predictors and returns in an out-of-sample setting is hard. As discussed also in Bianchi, Ludvigson, and Ma (2022), it is not obvious how to include shocks in the dataset fed to a machine or econometrician making a real-time forecast, because large structural shocks are infrequent and routinely identified ex post. In this light, it is striking that the professional forecast does respond so quickly and so strongly. Ultimately, the interpretation and quantitative assessment of macro shocks drives a large wedge between different forecast methods.

4. Time-varying discount rates

In this section, we present a novel approach to quantify the discount rate channel more precisely and test additional features of this explanation for the strong countercyclical response of professional stock return forecasts. This approach also provides us with a powerful test comparing the response of the professionals' forecast to the ex post truth, that is, realized returns.

To start, Figure 6 presents the impulse response function of realized returns and the professionals' forecast for each macro shock. Consistent with the idea that realized returns are poor proxies for expected returns, we see that the response of realized returns is more noisy and is relatively poorly estimated. That said, we see the first, albeit suggestive, evidence that the professionals' forecast agrees in important dimensions with realized returns. Like the professionals' forecast, realized returns respond quite fast and by a large positive amount to the Cycle and TFP shocks. Realized returns respond to the oil shock as well, but only at horizons beyond three years. Below, we aggregate these responses in an economically informed way to increase power.

4.1 Method

The general approach in this section is inspired by a two-stage least-squares (2SLS) instrumental variables estimation. In the first stage, we estimate the impact of macro shocks on some proxy for discount rates (e.g., uncertainty, the dividend yield, or the default spread), denoted X throughout this section. If the professionals feed into their return forecasting model real-time information that reflects the impact of shocks on fundamentals, we would expect the professionals' forecast to be sensitive to X only if the shocks are strong instruments in the first stage. Intuitively, in this case the professional is likely to understand the origins of X-based return predictability (i.e., the fact that X predicts realized returns in a standard predictive time-series regression). However, if the shocks are weak instruments for X, the professionals' forecast should be less sensitive. With the second stage, we assess the sensitivity of forecasts of returns as well as realized returns to each component of X, that is, the component instrumented with the shocks versus the residual component.

Under the strong assumption that the macro shocks impact (forecasts of) returns only because they impact the proxy for discount rates (i.e., the shocks satisfy the exclusion restriction), one can interpret the evidence from this section as estimating the causal effect of X on returns. If the exclusion restriction is strictly violated, the results of this section are interesting nonetheless. We are the first to quantify the relation between the shock-based component of popular proxies for discount rates and stock returns. Doing so for both realized returns and forecasts, we are also able to assess the quality of the different forecasts conditional on shocks. We use this fact to revisit the "cyclicality gap" introduced in Nagel and Xu (2023). These authors run predictive regressions on a large set of discount rate proxies X and show that ex post realized returns are much more sensitive (so, more countercyclical) than forecasts for the average X. We instead study if there is variation in the cyclicality gap across different discount rate proxies depending on how strongly each X responds to macro shocks.

To accommodate the usual 2SLS setup, we need to pick a single point on the impulse response function of X to shocks for the first stage. To pick this point, we need to consider two features of impulse responses. First, shocks tend to have a persistent impact. Second, the responses to different shocks peak (or trough) at different horizons.[19] For these reasons, we focus on the impact due to one year worth of macro shocks by estimating the following regression:

$$X_{t-\frac{1}{4}} = \delta\left(\sum_{k=\kappa}^{\kappa+1-\frac{1}{4}} N_{t-\frac{1}{4}-k}\right) + \gamma C_{t-\frac{5}{4}-\kappa} + \epsilon_t \tag{8}$$

for $\kappa = 0, \frac{1}{4}, ..., \frac{5}{4}$. By varying κ, we allow for delayed responses to the instrument: the latest shock included in the sum will hit the economy in $t - \frac{1}{4}$, ..., or $t - \frac{3}{2}$. $C_{t-\frac{5}{4}-\kappa}$ is a vector of controls, which includes a constant, lagged X and lagged returns. These controls are realized at $t - \frac{5}{4} - \kappa$ such that there is no overlap with any of the shocks included in Eq. 8. In particular, we predict X at t minus one quarter, such that all relevant information ($X_{t-\frac{1}{4}}$ and the shocks) is available to forecasters when making their forecast in quarter t. We initially run these regressions for one shock (Cycle, TFP, or Oil) at a time and choose the $\kappa^\star$ that maximizes the t-statistic on the instrument, that is, $\left(\sum_{k=\kappa}^{\kappa+1-\frac{1}{4}} N_{t-\frac{1}{4}-k}\right)$. One may be concerned that in this way we overestimate the impact of the shock on X. However, note that $\kappa^\star$ is effectively cross-validated in the second stage: If the first-stage impact of the shocks at this horizon is pure noise, there is no reason why realized or forecasted returns should respond more strongly in the second stage to the component of X instrumented by the shocks.

Since efficiency is increasing in the number of relevant instruments, we also run the regression in Eq. (8) with multiple shocks on the right-hand side.[20] Intuitively, with multiple instruments we better identify the total variation in X that is due to shocks that have relatively large effects on macro-variables, such as GDP, unemployment, and inflation. The residual variation is due to shocks with relatively small macro-effects or other shocks, such as sentiment. We consider two cases for the multiple shock setup. First, we consider the

[19]For instance, we have already seen that stock return forecasts respond faster to Cycle and TFP shocks than to Oil shocks, while responses to all three shocks show substantial persistence. To see that these facts also apply to our discount rate proxies, we present in Figure F.5 the IRF for two measures of real uncertainty that feature prominently in what follows.

[20]For consistency, we fix the lag choices $\kappa^\star$ to what they are in the single shock setup. Our conclusions are unchanged when we re-optimize the lag choices in the multiple regression, because the shocks are largely uncorrelated.

Cycle shock and TFP news, since both are available over a long sample from 1962 until 2017. Second, we add the Oil shock.[21]

Taking a step back, we are interested in how realized and professional forecasts of returns respond to variation in X. The standard approach in the literature is to run the predictive regressions:

$$R_{t+1} = \lambda X_{t-\frac{1}{4}} + \gamma C_{t-\frac{5}{4}-\kappa^\star} + \epsilon_{t+1} \tag{9}$$

$$F_t^P(R_{t+1}) = \lambda^P X_{t-\frac{1}{4}} + \gamma^P C_{t-\frac{5}{4}-\kappa^\star} + \epsilon_{t+1}^P \tag{10}$$

where R_{t+1} is the return over the year starting at the end of the quarter in which the professionals' forecast $F_t^P(R_{t+1})$ is made. To be consistent with Eq. (8) above, $C_{t-\frac{5}{4}-\kappa^\star}$ is a vector of controls that includes a constant, lagged X and lagged realized returns. Our proposed discount rate channel implies that the cyclicality gap $\lambda - \lambda^P$ varies across X. To be precise, for X with a strong first stage (δ significant in Eq. (8)), we expect $\lambda^P > 0$, such that the cyclicality gap $\lambda - \lambda^P$ is smaller. This would mean that the professionals' forecast is similarly countercyclical as realized returns. That said, comparing these reduced form coefficients is tricky, because forecasts of and realized returns can relate quite differently to components of X that are due to macro shocks relative to components that are not. On top of that, the impact these components have on stock returns may be correlated. For instance, the dividend yield will increase if a bad macro shock leads to an increase in discount rates, but it will also increase when investors underreact to the shock's impact on future cash flows. These effects have an opposite impact on realized future returns, but the professional may be more sensitive to the former discount rate shock.

For a better identified test of the discount rate channel, we run the following second stage

[21] Because this shock is available over a relatively short sample (from 1975 onwards), we pad this shock with zeroes so that we do not lose all the information from the Cycle and TFP shocks from 1962 to 1975. Our results are similar if we run the first stage regression using the sample starting in 1975, see Table F.3.

regressions:

$$R_{t+1} = \lambda_S \widehat{X_{t-\frac{1}{4}}^{Shock}} + \lambda_R \widehat{X_{t-\frac{1}{4}}^{Res.}} + \gamma C_{t-\frac{5}{4}-\kappa^\star} + v_{t+1} \tag{11}$$

$$F_t^P(R_{t+1}) = \lambda_S^P \widehat{X_{t-\frac{1}{4}}^{Shock}} + \lambda_R^P \widehat{X_{t-\frac{1}{4}}^{Res.}} + \gamma^P C_{t-\frac{5}{4}-\kappa^\star} + v_{t+1}^P \tag{12}$$

where we include both instrumented X, that is, the fitted value from the first stage that captures the part of X due to macro shocks ($\widehat{X_{t-\frac{1}{4}}^{Shock}}$), and the first-stage residual that captures all remaining variation in X ($\widehat{X_{t-\frac{1}{4}}^{Res.}}$).[22,23]

We broadly test two hypotheses, which should hold for all X with a strong first stage. First, the discount rate channel suggests that $\lambda_S^P > \lambda_R^P$, which means that the professionals' forecast is relatively more sensitive to variation in X that is due to macro shocks. Second, it suggests that $\lambda_S = \lambda_S^P$, which means that professionals' understanding of the relation between macro shocks and discount rates is consistent with reality. In other words, this hypothesis implies that there is no cyclicality gap conditional on shocks. In turn, including the residual in the second stage assesses the sensitivity of (forecasts of) returns to variation in X that may be harder to interpret for professionals, under the assumption that time-varying risk premia drive their forecasting model. We also run these tests for the non-professionals' and objective forecast, such that we can judge how special the professionals' forecast is in these important dimensions.

Because the time-series we study are relatively short and not all our data is observed at the same frequency (e.g., the shocks and some predictors are observed quarterly, but the professionals' forecast is observed biannually), we estimate standard errors using the stationary bootstrap of Politis and Romano (1994) with average block length set to one year. We find that these standard errors are more conservative than the usual GMM-based

[22] This regression is an application of the control function approach first discussed in Hausman (1978) and more recently in Wooldridge (2015).

[23] Recall that in the case of a single shock (instrument), the product of the coefficients δ and λ_S from Eqs. (8) and (11), respectively, is equal to the coefficient from the reduced-form regression of returns on the shock. To stick to the usual 2SLS setup, we use the same control variables in the first and second stage. Since we do not control for lagged forecasts in the first stage (because it is not obvious whose forecasts to control for), we do not control for them in the second stage either. We show in Table F.2 that our main results for the professionals' forecast are robust to additionally controlling for their lagged forecast.

2SLS standard errors.

4.2 Data

Following a large and growing literature, we first study uncertainty as a driver of time-varying risk premia (see, among many others, Nakamura, Sergeyev, and Steinsson, 2017; Bianchi, Ilut, and Schneider, 2017). Uncertainty is uniquely suited to our setting, because it directly relates to the problem of forecasting. As described in Jurado, Ludvigson, and Ng (2015), uncertainty over a variable should econometrically be measured as the volatility of the un-forecastable component of future values of that variable. To measure uncertainty, we use the real uncertainty index of Ludvigson, Ma, and Ng (2021). This index is an aggregation of estimated 12 month-ahead uncertainties for 73 real activity variables constructed using a large panel of conditioning information. Note also that Ludvigson, Ma, and Ng (2021) argue that real uncertainty rises in response to adverse shocks, rather than real uncertainty being a primitive shock. Our use of uncertainty as endogenous variable in the two-stage approach is consistent with this interpretation. We also provide evidence for a second measure of real uncertainty used in Bloom (2009), that is, the dispersion in GDP forecasts from the Livingston survey. The disadvantage of this measure is that it is far from the precise econometric definition of uncertainty mentioned above. The advantage is that the measure, albeit simple, links directly to our main motivation for this section: while the consensus GDP forecast hardly responds to shocks, the shocks lead to increased uncertainty about future GDP and thus to more dispersion. Intuitively, professionals may find it easier to predict second moments, which feed into expected stock returns, than first moments.

Next, we analyze the much broader set of popular proxies for discount rates described in Online Appendix A. Given the large diversity in economic motivation for these predictors, we interpret these results as broadly speaking to the time-varying fundamentals or risk aversion that ultimately determine discount rates.

4.3 Results

Let us turn to the results, which paint a surprisingly uniform picture considering that we are studying a large variety of discount rate proxies and shocks with distinct economic origins.

4.3.a Uncertainty

We report estimates from the first stage Eq. (8) in Table II for both the real uncertainty index of Ludvigson, Ma, and Ng (2021, denoted Unc in Panel A) and the dispersion in Livingston GDP forecasts (denoted $Disp$ in Panel B). Across columns we consider each shock in isolation as well as two combinations of shocks (Cycle and TFP in column four and all three shocks in column five). We also report the first stage effective F-statistic (from the robust test for weak instruments of Montiel Olea and Pflueger, 2013) and the optimal lag choice κ. In the multiple shock-setting, we report the largest κ across shocks. That κ is leading in the definition of the controls and in this way we avoid overlap between controls and shocks.

Let us focus first on the evidence for Unc in Panel A. We see that Unc increases by 0.28, 0.40, and 0.12 standard deviations for a one standard deviation (annual) Cycle, TFP, and Oil shock, respectively. Thus, bad shocks lead an increase in real uncertainty, as reported also in Ludvigson, Ma, and Ng (2021). While the coefficients for the Cycle and TFP shock are strongly significant ($t > 3$), the coefficient for Oil shocks is marginally insignificant ($t = 1.5$). Hence, Oil shocks are a weak instrument in isolation. In contrast, the TFP shock is strong and the effective F-statistic of 32.45 is significant at the 1% level. Although the effective F-statistic for the Cycle shock is not significant, the value of 9.65 versus 2.42 for the Oil shock suggests that the Cycle shock is relatively much stronger.

We are most interested in the tests with multiple shocks, which tests are more efficient when the shocks are strong enough. Economically, the model with multiple shocks also gets closer to decomposing real uncertainty in the total variation coming from primitive macro shocks and residual variation. In the multiple regression with Cycle and TFP, we find that the coefficients on both shocks are significantly positive, with similar coefficient estimates as before (0.18 and 0.33, respectively). The first stage with these two shocks is strong, because the effective F-statistic of 18.03 is significant at the 1% level. In the last column, we see that adding to these two shocks the Oil shock does not improve the first stage much. Oil shocks do not have a significant impact on uncertainty and the effective F-statistic drops and is only significant at the 10% level.

In Panel B, where *Disp* is used to proxy for real uncertainty, we find largely similar results. Both Cycle and TFP have an economically large and statistically significant impact on *Disp*, with a one standard deviation shock leading to an increase in *Disp* of about 0.3 standard deviations ($t > 3$). Consistent with this finding, the model that joins these two shocks has the strongest first stage (the effective F-statistic is 10.61 and is significant at the 1% level). In contrast, the impact of the Oil shock on *Disp* is negative and relatively small, both in isolation and controlling for the other two shocks.

Looking at the optimal lag lengths, we see that the response of both measures of real uncertainty to Cycle shocks peaks in the first year after the shock, whereas the response peaks between one and two years out for the TFP and Oil shocks. Consistent with the idea that it is uncertainty that mediates the impact of the shocks on professional forecasts of returns, we have also seen in Figure 3 that professional forecasts respond relatively faster to Cycle shocks. We formalize the link between shocks, uncertainty, and (professional) forecasts next using evidence from the second stage. Because Cycle and TFP shocks jointly provide for the strongest first stage, the second stage regressions reported in Table III focus on the specification with these two shocks. As a robustness check, we present results for the case of three shocks in Table F.1.

Table III presents evidence from the second stage for realized returns as well as professional, non-professional and objective forecasts. To put this evidence in perspective, we also report the results from simple regressions of (forecasts of) returns on the two measures of real uncertainty.

Let us focus first on the reduced form evidence for the real uncertainty index of Ludvigson, Ma, and Ng (2021) in Panel A. In the first column, we see that realized returns are largely insensitive to Unc: a one standard deviation increase in Unc leads to an increase in realized annual excess returns of only 71 basis points. This finding is surprising from the perspective of rational asset pricing models where uncertainty is priced. Forecasts of professionals are more consistent with these models, because they predict that returns increase by a large and significant 3.60% (t = 2.92). Thus, the professionals' forecast is more countercyclical than realized returns, suggesting that the cyclicality gap can even switch sign for a discount rate proxy that is strongly impacted by shocks. Professionals are alone in their assessment, because the coefficients for non-professionals and the objective forecast are comparatively small at -3 and +34 basis points, respectively. In Panel B, we see that professionals do not predict as large an increase in return for a one standard deviation increase in our second measure of real uncertainty, $Disp$. In fact, for this measure, the response of realized returns is relatively large at 2.28% versus 0.64% for the professional.

As discussed above, one cannot draw strong conclusions from comparing these reduced form coefficients. As argued in recent literature, comparing forecasts to realizations sets a really high bar for any forecaster. This is especially true for returns, which are very noisy. On top of that, variation in real uncertainty is hard to interpret. On one hand, a wide variety of mechanisms can generate an endogenous response of uncertainty to economic shocks.[24] On the other hand, uncertainty shocks may also be a root cause of time-variation in fundamentals.[25] Given these different roles of uncertainty, returns may relate to uncertainty in complicated ways. This is exactly why we are interested in instrumenting uncertainty using macro shocks. Existing models agree that the uncertainty induced by these shocks should lead to higher expected returns. With our second stage regressions, we ask if professionals get this dynamic right.

[24]Some theories suggest that bad times incentivize risky behavior (Bachmann and Moscarini, 2012; Fostel and Geanakoplos, 2012), or reduce information and hence the predictability of future outcomes (Van Nieuwerburgh and Veldkamp, 2006; Fajgelbaum, Schaal, and Taschereau-Dumouchel, 2017), or provoke new and highly uncertain economic policies (Pastor and Veronesi, 2013), or generate endogenous countercyclical uncertainty in consumption growth as investment is costly to reverse (Gomes, Jermann, and Schmid, 2016).

[25]This includes models of the real options effects of uncertainty (Bernanke, 1983; McDonald and Siegel, 1986), or where uncertainty influences financing constraints (Gilchrist, Sim, and Zakrajšek, 2014; Arellano, Bai, and Kehoe, 2019), or precautionary saving (Basu and Bundick, 2018; Leduc and Liu, 2016; Fernández-Villaverde et al., 2011). These theories presume that uncertainty is an exogenous shock to underlying fundamentals.

In the second column of Panel A, we see that the macro shock-component of real uncertainty ($\widehat{Unc^{Shock}}$) has a large positive impact on realized returns. For a one standard deviation increase in Unc that is solely due to Cycle and TFP shocks, realized returns increase by an economically large 8.16%. For the same increase in Unc, professional forecasts increase by roughly the same magnitude, that is, a significant 6.79%. Further, we see that the difference between these two coefficients is small and insignificant at 1.37% ($t = 0.54$). This formally means that the forecast error (realized minus professional forecast) is not predictable, consistent with our proposed discount rate channel (in particular, the hypothesis that $\lambda_S = \lambda_S^P$, see Eqs. (11) and (12)). We conclude that professionals correctly understand the response of returns to the macro shock-component of uncertainty or, following the stronger 2SLS interpretation, the causal impact of real uncertainty on stock returns. Professionals are unique in their correct assessment. Although the response of the objective forecast is positive and significant, it is small. For both non-professionals and the objective forecast, the difference between the impact of $\widehat{Unc^{Shock}}$ on realized returns and the forecast is large and significant at over 7.5%.

Also consistent with the discount rate channel (in particular, the hypothesis that $\lambda_S^P > \lambda_R^P$ in Eq. (12)), we see that the professionals' forecast is much more sensitive to the shock-based than the residual component of uncertainty, with a coefficient estimate for $\widehat{Unc^{Shock}}$ that is almost three times larger than the estimate for $\widehat{Unc^{Res.}}$: $\lambda_S^P = 6.79\%$ versus $\lambda_R^P = 2.02\%$. Indeed, what professionals get wrong is how the residual variation in Unc (the endogenous variation under the stronger 2SLS interpretation) relates to returns. Indeed, while this residual component has a small negative impact on realized returns, professionals forecast it has a small positive impact. Thus, we conclude that the divergence in the coefficients from realized returns and professional forecasts on Unc is driven by the residual component. Because the two component of Unc have such different effects on realized returns (-2.69% vs 8.16%), we also find that the R^2 in the regression with decomposed Unc is much higher than in the regression that does not decompose Unc (4% vs -1%).

These conclusions are robust in various important dimensions. First, we see in Panel B of Table III that professionals correctly assess the impact of the macro shock-component of real uncertainty also when it is measured using *Disp*. In this case, the relevant coefficients equal 5.58% for realized returns and 6.15% for professional forecasts. The difference of -0.57% capturing the professionals' forecast error is again insignificant. For both realized returns and professional forecasts, these coefficients are larger than the response to not-decomposed uncertainty as well as the residual component of uncertainty, consistent with the evidence in Panel A and the hypotheses derived from our proposed discount rate channel. Second, in Table F.1, we see similar results when we use all three shocks as instruments for real uncertainty. Third, we present evidence from reduced-form regressions of (forecasts of) returns on the shocks in Table F.4 of the Online Appendix. We find that professional forecasts responds to all three shocks with a positive coefficient. This response is economically large and significant for Cycle and TFP shocks at 3 and 2%, respectively. The response of realized returns to these two shocks is similarly large at about 3% and marginally significant. Realized returns respond with a small negative coefficient to Oil shocks, however. Neither non-professional nor objective forecasts respond by an economically large amount to any of the three shocks. If anything, non-professional forecasts respond with the wrong sign (when compared to predictions from rational asset pricing models and realized returns).

In conclusion, our evidence is broadly consistent with the discount rate channel, which says that the professionals forecast is sensitive to shocks because they correctly understand that shocks lead to real uncertainty that is priced. Consequently, there is little evidence of a cyclicality gap when we consider the response of returns to uncertainty instrumented using macro shocks. Consistent with the idea that non-professionals are more prone to extrapolation and the impact of shocks on realized returns is relatively small because these returns are noisy, we find that the shocks have little impact on non-professional forecasts. Objective forecasts do respond significantly and with the same sign as professionals to shocks. However, this response is economically very small in all of our tests. This result is perhaps unsurprising noting that the shocks are not included in the objective forecasting algorithm and it is not obvious how to do so. Hence, in real-time, the objective forecast can only learn about the impact of shocks on returns through their impact on the predictors used in the algorithm. If this impact is hard-to-estimate and/or small for some of these predictors, the implied response of the objective forecast to shocks will also be small. We analyze such variation across predictors in the next subsection.

4.3.b Alternative discount rate proxies

The hypotheses derived from our discount rate channel should hold more generally than just for real uncertainty. Indeed, we are assuming that professional forecasters understand the impact of macro shocks on discount rates. To test this idea more generally, we run the estimation outlined in Section 4.1 for the extended set of 22 stock market predictors used in the objective forecast. Our main interest is in asking how (forecasts of) returns respond to these predictors, taking into account the differential impact of the shocks on these predictors. Indeed, under the null that the impact of macro shocks on discount rates is captured in real-time in the professionals' model, in line with the implications of rational asset pricing models, professionals may be ill-equipped to understand the variation in predictors that do not strongly respond to macro shocks (i.e., predictors for which the the first stage regression (8) is weak).

We start by selecting the subset of predictors for which macro shocks have a large and consistent impact. We call such predictors "shock-sensitive" throughout and we capture this sensitivity using two criteria. As before, we focus on the case with two shocks (Cycle and TFP). First, the predictor must respond to the two shocks with the same sign. By definition, some predictors are procyclical, while others are countercyclical. To accommodate interpretation in what follows, we sign the predictors so that all of them predict realized returns with a positive sign.[26] Thus, a positive response to the two shocks can be interpreted as countercyclical variation in discount rates. Second, the first stage must be strong (the effective F-statistic must be significant at the 10%-level). We present point estimates from the first stage for all predictors in Table IV. According to these criteria, the following eight

[26]The following predictors are multiplied by −1: CSP, DFR, $EXPD$, GDP, IK, $INFL$, IP, and $NTIS$.

predictors are shock-sensitive: book-to-market (BM), default spread (DEF), dividend yield (DY), dividend-to-price ratio (DP), real factor ($F1$), GDP growth (GDP), investment-to-capital ratio (IK), and industrial production growth (IP).

With this classification in hand, we move to the relation between predictors and (forecasts of) returns, both in reduced form and in the two-stage setup. Let us start with the reduced form regression of (forecasts of) returns on each predictor (see Eq. (9)). Figure 7 plots the coefficient estimates for both subsets of predictors. Point estimates for the shock-sensitive predictors are reported in Table V (and those for the shock-insensitive predictors in Table F.5).

For the shock-insensitive predictors in Panel A, we see that there is a large divergence between the predictive coefficients for realized returns and professional forecasts. The former are generally larger and in various cases the sign of the coefficient for professionals is opposite to the one for realizations. Out of 14 predictors, two are statistically significant for realized returns (CAY, CSP), while three others are significant for professional forecasts ($EXPD$, $SVAR$, and $VIX2$). This evidence suggests that professionals do not understand well the relation between returns and those predictors that are only weakly affected by macro shocks. In the words of Nagel and Xu (2023), shock-insensitive predictors present a large cyclicality gap.

This finding stands in stark contrast to the shock-sensitive predictors in Panel B. For all these predictors, the coefficients for realizations and forecasts have the same sign and are close in magnitude. In other words, there is little evidence of a cyclicality gap for shock-sensitive predictors. In fact, in Table V we see that these coefficients are always significant for professionals and often at least marginally significant for realized returns. The fact that professional forecasters respond strongly and countercyclically to these predictors alone is consistent with the idea that a substantial part of the variation in these predictors captures discount rate variation driven by macro shocks and, therefore, our proposed discount rate channel.

Panel C of Figure 7 presents the coefficient estimates from the second stage regression (see Eq. 11) for realized returns and professional forecasts for shock-sensitive predictors. To be precise, we plot the coefficient on the instrumented component of each predictor, where the instruments are the Cycle and TFP shocks. We present the full set of point estimates for the second stage regressions in Table V. Similar to Panel B, we see in Panel C that coefficients on the instrumented predictors are close for realized returns and professional forecasts. The difference between these two coefficients (measuring whether the forecast error is predictable) is economically small and insignificant in all eight cases (see Table V). Furthermore, the magnitude of the coefficients in Panel C is on average about twice that in Panel B (4.04% versus 7.02% for realized returns and 4.62% versus 8.22% for professionals). This result means that (forecasts of) returns are much more sensitive to the exogenous variation in the predictor due to macro shocks than the residual (endogenous) variation. This finding is consistent with the idea that the causal impact of the shocks on discount rates is large and professionals correctly assess this impact.

We conclude that professionals are unique in that they update their stock return forecasts aggressively when a shock hits and they do so consistent with the impact of the shock on discount rates. An important element of our claim is that professional forecasts are unbiased conditional on shocks. Let us test this directly to connect to the summary statistics reported in Table I. Here, we saw that professional forecasts are unconditionally downward biased, but not in bad times measured as NBER recessions. To see whether professional forecasts (made in quarter t) are unbiased conditional on shocks, we define a bad shock dummy that is equal to one in quarter t when one or more of the annualized macro shocks are above their 90th percentile. As in the previous subsections, we consider separately the case of Cycle and TFP shocks and the case with all three shocks.[27] For the case with two (three) shocks, this approach implies that about 15% (20%) of our sample is a bad shock quarter.

We find that realized returns in the year after a bad shock are well above average at 14.5% and 8.7% in the two and three shock case, respectively, which is relative to less than

[27]Consistent with the first-stage regressions for real uncertainty in Panel A of Table II, we cumulate Cycle shocks over quarters $t-3/4:t$ and TFP and oil supply shocks over quarters $t-6/4:t-3/4$. This definition is closely linked to NBER recessions. For the case with two (three) shocks, 56% (45%) of bad shock quarters are NBER recession quarters, whereas only 5% (4%) of other quarters are NBER recession quarters.

6% in quarters without shocks. Professional forecasts are well above average as well after bad shocks at 10.8% and 9.3% in the two and three shock case, respectively, which is relative to about 3% in other quarters. This result confirms again that the professionals' forecast is largely unbiased in bad times. While the objective forecast is also countercyclical, it does not vary much between quarters with and without bad shocks at 8.5% versus 7.5%. Non-professional forecasts do not respond to the shocks at all and are roughly 6% in all cases. At the moment a large macro shock hits, we should turn to professionals to understand the impact of the shock on future stock returns.

5. Objective and subjective forecasts using the present-value restriction

Campbell and Shiller's (1988b) framework to produce objective expectations of returns and cash flows disciplined by a present-value relation has been a workhorse in empirical finance. In this section, we build on the approach in Van Binsbergen and Koijen (2010) to compare the response to macro-shocks of objective and subjective expectations. We will impose present-value restrictions with expected returns and dividend growth modeled as unobservable quantities with flexible dynamics in a state-space system. We summarize our approach and main conclusions here and present further details and all results in Online Appendix E.

5.1 Objective expectations

In a first exercise, we use the exact specification of Van Binsbergen and Koijen (2010) to filter the objective expectations of returns, μ_t, and dividend growth, g_t from the log price-to-dividend ratio pd_t and realized dividend growth Δd_t. We assume the following dynamics for these quantities:

$$\mu_{t+1} = \delta_0 + \delta_1(\mu_t - \delta_0) + \epsilon_{\mu,t+1} \tag{13}$$

$$g_{t+1} = \gamma_0 + \gamma_1(g_t - \gamma_0) + \epsilon_{g,t+1} \tag{14}$$

$$\Delta d_{t+1} = g_t + \epsilon_{d,t+1} \tag{15}$$

$$pd_t = A - B_1(\mu_t - \delta_0) + B_2(g_t - \gamma_0) \tag{16}$$

$$Var(\epsilon_{t+1}) = Var([\epsilon_{d,t+1}, \epsilon_{g,t+1}, \epsilon_{\mu,t+1}]) = \begin{pmatrix} \sigma_d^2 & 0 & \sigma_{d\mu} \\ 0 & \sigma_g^2 & \sigma_{g\mu} \\ \sigma_{d\mu} & \sigma_{g\mu} & \sigma_\mu^2 \end{pmatrix} \tag{17}$$

where $A = (\kappa + \gamma_0 - \delta_0)/(1-\rho), B_1 = 1/(1-\rho\delta_1),$ and $B_2 = 1/(1-\rho\gamma_1)$ using $\kappa = ln(1+e^{\bar{pd}}) - \rho\bar{pd}$ and $\rho = e^{\bar{pd}}/(1+e^{\bar{pd}})$. The state-space system uses as observables the annual (end-of-year) repurchase-adjusted dividend growth and dividend-to-price series from Nagel and Xu (2022).

Our main interest is in the responses of filtered expected returns, μ_t, and dividend growth, g_t, to macro shocks. In Figure E.1, we display the estimated responses, which are obtained by projecting the filtered series on the shocks, with each shock cumulated over the four quarters in a year, and controlling for lags of the shock, expected return, and expected growth. Considering first the response of objective expected returns after a Cycle and TFP shock, we observe an increase that peaks at about 60 bps to 70 bps.[28] The responses are larger in economic magnitude than those of the objective forecasts studied in Section 3, which peak at less than 20 bps (Figure 3). This finding suggests that even though a combination forecast may be useful to predict returns out-of-sample, it does not adequately reflect the true discount rate shock incorporated in realized returns. More importantly, the responses of the filtered objective forecasts are still much smaller in magnitude than those of professionals, which are in the range of 200 bps, as shown in Figure 3.

Objectively expected dividend growth experiences a marginally significant decrease in the years after all three shocks, with the trough impact ranging between -65 bps and -21

[28] Consistent with previous evidence that the Oil shock is a relatively weaker instrument for discount rates, objective expected returns do not increase after this shock.

bps. This finding suggests that the drop in realized growth seen in Figures 2 and 5 for GDP and dividends, respectively, is partially predictable ex post. We now ask if the professional's subjective growth forecast reflects this predictability.

5.2 Subjective professional expectations

Since the present value identity must hold for any set of expectations, we can rewrite Eq. (16) as $pd_t = A^P - B_1^P(\mu_t^P - \delta_0^P) + B_2^P(g_t^P - \gamma_0^P)$, where the superscript "^{P}" refers to the subjective expectation of the professional forecaster. This fact implies that we can infer the professional's dividend growth forecast from his/her return forecast under the assumption that he/she is using the present value model and believes expected returns and growth follow the dynamics in Eqs. (13) and (14). Using semi-annual observations of the professional's return forecasts and the repurchase-adjusted dividend-to-price series from Nagel and Xu (2022), we filter the professional's forecast of dividend growth.

Our main interest is in how this growth forecast responds to shocks. To this end, we project the filtered series (μ_t^P and g_t^P) on the shocks, with each shock cumulated over two quarters. The projections control for lags of the shock, expected return, and expected growth. We report these responses in Figure E.2. First, the responses for μ_t^P are large, significant, and almost indistinguishable from what we reported in Figure 3. Second, the response of subjective dividend growth expectations is relatively much smaller. If anything, subjectively expected dividend growth increases after each shock (by about 50 bps), which response is opposite in sign to that of realized dividend growth.

Overall, the evidence from these present-value systems confirms our earlier conclusions: macro shocks have a large impact on subjective return expectations, just like rational models would predict, and this impact is consistent with realized returns. In contrast, the present-value identity implies that shocks have a much smaller impact on subjectively expected dividend growth. An interesting question is what can explain these subjective growth forecasts. As discussed in Online Appendix E, we find that subjective growth expectations are persistent and explain a relatively large share of the unconditional variation in the price-dividend ratio. This persistence is consistent with sticky cash flow expectations as assumed in Gomez-Cram (2022), for instance. That said, this assumption alone cannot explain that subjective cash flow expectations respond, if anything, with the wrong sign to macro shocks.

6. Conclusion

Survey-based professional forecasts of aggregate stock market returns are three times more volatile than non-professional and objective forecasts. This relative volatility is surprising because professional forecasts of macro variables, like GDP and inflation, are similarly smooth as non-professional and objective forecasts of these quantities. We attribute a large share of this relative volatility to macro shocks, including business cycle, TFP news, and oil supply shocks. In contrast to non-professional and objective stock return forecasts, professional forecasts are strongly countercyclical, consistent with implications from rational asset pricing models.

To shed light on the economic mechanism behind these results, we employ a novel two-stage least squares procedure to identify the discount rate variation in professional forecasts due to macro shocks. We show this variation aligns qualitatively and quantitatively with ex post realized returns. We conclude that professionals correctly assess the countercylical discount rate impact of macro shocks, while other non-professional and objective forecasts of stock returns do not. This discount rate channel can also explain why forecasts of macro variables by the same professionals do not respond strongly to the same shocks. When bad shocks hit, the true extent of the adverse impact on the macroeconomy may be hard to assess (e.g., because the exact size of the shock is unknown in real-time or because signals contain more noise in times of large shocks), while the impact on uncertainty or risk, more generally, may be quite evident to professionals.

Overall, this paper quantifies forecasters' subjective assessments of the macroeconomic origins of time-variation in expected stock returns and demonstrates that professionals' beliefs offer valuable insights. More broadly, our findings reveal a tension that should be considered when modelling these beliefs. On one hand, professionals' forecasts respond more strongly to macro shocks via discount rate news than cash flow news. On the other hand, when imposing the present-value identity, we observe that the professionals' return forecasts are less persistent than what is typically assumed for expected returns in benchmark calibrations of classical asset pricing models. As a result, most of the time-variation in the price-dividend ratio is attributed to subjective expectations of cash flow growth. To reconcile these facts within a rational framework, macroeconomic shocks ought to drive discount rate disturbances that are more volatile and less persistent than other discount rate shocks, while frictions prevent forecasters from fully updating their cash flow expectations when new information arises. Although the strong response of professionals to discount rate news challenges popular behavioral explanations, recent models of agents with sticky expectations Gomez-Cram (2022), fading memory Nagel and Xu (2022), or diagnostic expectations Bordalo et al. (2024) feature a weak response to cash flow news. As argued in Nagel (2024), such explanations imply that most of the variation in valuation ratios stems from long-term cash flow expectations. With either modelling approach, we believe exploring this tension is a promising direction for future research that can provide deeper insights into how aggregate shocks shape expectations.

References

Adam, K., D. Matveev, and S. Nagel, 2021, Do survey expectations of stock returns reflect risk adjustments?, *Journal of Monetary Economics* 117, 723–740.

Andre, P., C. Pizzinelli, C. Roth, and J. Wohlfart, 2022, Subjective models of the macroeconomy: Evidence from experts and representative samples, *The Review of Economic Studies* 89, 2958–2991.

Angeletos, G.-M., F. Collard, and H. Dellas, 2020, Business-cycle anatomy, *American Economic Review* 110, 3030–70.

Angeletos, G.-M., Z. Huo, and K. A. Sastry, 2021, Imperfect macroeconomic expectations: Evidence and theory, *NBER Macroeconomics Annual* 35, 1–86.

Arellano, C., Y. Bai, and P. J. Kehoe, 2019, Financial frictions and fluctuations in volatility, *Journal of Political Economy* 127, 2049–2103.

Bachmann, R., and G. Moscarini, 2012, Business Cycles and Endogenous Uncertainty, Working Yalo University.

Barsky, R. B., and E. R. Sims, 2011, News shocks and business cycles, *Journal of Monetary Economics* 58, 273–289.

Bastianello, F., 2022, Time-series and cross-section of risk premia expectations: A bottom-up approach, *Available at SSRN 4204968* .

Basu, S., and B. Bundick, 2018, Uncertainty shocks in a model of effective demand: Reply, *Econometrica* 86, 1527–1531.

Baumeister, C., and J. D. Hamilton, 2019, Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks, *American Economic Review* 109, 1873–1910.

Beaudry, P., and F. Portier, 2006, Stock prices, news, and economic fluctuations, *American Economic Review* 96, 1293–1307.

Ben Zeev, N., and H. Khan, 2015, Investment-specific news shocks and us business cycles, *Journal of Money, Credit and Banking* 47, 1443–1464.

Bernanke, B. S., 1983, Irreversibility, Uncertainty, and Cyclical Investment*, *The Quarterly Journal of Economics* 98, 85–106.

Bianchi, F., C. L. Ilut, and M. Schneider, 2017, Uncertainty Shocks, Asset Supply and Pricing over the Business Cycle, *The Review of Economic Studies* 85, 810–854.

Bianchi, F., S. C. Ludvigson, and S. Ma, 2022, Belief distortions and macroeconomic fluctuations, *American Economic Review* 112, 2269–2315.

Bloom, N., 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.

Boons, M., G. Ottonello, and R. Valkanov, 2023, Do credit markets respond to macroeconomic shocks? the case for reverse causality, *The Journal of Finance* 78, 2901–2943.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer, 2020, Overreaction in macroeconomic expectations, *American Economic Review* 110, 2748–82.

Bordalo, P., N. Gennaioli, R. L. Porta, and A. Shleifer, 2024, Belief overreaction and stock market puzzles, *Journal of Political Economy* 132, 1450–1484.

Campbell, J. Y., and R. J. Shiller, 1988a, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

Campbell, J. Y., and R. J. Shiller, 1988b, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661–676.

Campbell, J. Y., and S. B. Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *The Review of Financial Studies* 21, 1509–1531.

Carroll, C. D., 2003, Macroeconomic expectations of households and professional forecasters, *The Quarterly Journal of Economics* 118, 269–298.

Cochrane, J. H., 1991, Production-based asset pricing and the link between stock returns and economic fluctuations, *The Journal of Finance* 46, 209–237.

Cochrane, J. H., 2011, Presidential address: Discount rates, *Journal of Finance* 66, 1047–1108.

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

Coibion, O., and Y. Gorodnichenko, 2012, What can survey forecasts tell us about information rigidities?, *Journal of Political Economy* 120, 116–159.

Coibion, O., and Y. Gorodnichenko, 2015, Information rigidity and the expectations formation process: A simple framework and new facts, *American Economic Review* 105, 2644–78.

Couts, S. J., A. Gonçalves, and J. Loudis, 2023, The subjective risk and return expectations of institutional investors, *Fisher College of Business Working Paper* 014.

Dahlquist, M., and M. Ibert, 2024, Equity Return Expectations and Portfolios: Evidence from Large Asset Managers, *The Review of Financial Studies* 37, 1887–1928.

De La O, R., and S. Myers, 2021, Subjective cash flow and discount rate expectations, *The Journal of Finance* 76, 1339–1387.

Elliott, G., I. Komunjer, and A. Timmermann, 2008, Biases in macroeconomic forecasts: Irrationality or asymmetric loss?, *Journal of the European Economic Association* 6, 122–157.

Fajgelbaum, P. D., E. Schaal, and M. Taschereau-Dumouchel, 2017, Uncertainty Traps*, *The Quarterly Journal of Economics* 132, 1641–1692.

Fama, E. F., and K. R. French, 2002, Testing Trade-Off and Pecking Order Predictions About Dividends and Debt, *The Review of Financial Studies* 15, 1–33.

Farmer, L., E. Nakamura, and J. Steinsson, 2022, Learning about the long run, NBER Working Paper No. 29495.

Fernández-Villaverde, J., P. Guerrón-Quintana, J. F. Rubio-Ramírez, and M. Uribe, 2011, Risk matters: The real effects of volatility shocks, *American Economic Review* 101, 2530–61.

Fostel, A., and J. Geanakoplos, 2012, Why does bad news increase volatility and decrease leverage?, *Journal of Economic Theory* 147, 501–525.

Gandhi, M., N. J. Gormsen, and E. Lazarus, 2023, Excess persistence in return expectations, Working Paper Chicago Booth.

Gilchrist, S., J. W. Sim, and E. Zakrajšek, 2014, Uncertainty, financial frictions, and investment dynamics, Working Paper 20038, National Bureau of Economic Research.

Gomes, J., U. Jermann, and L. Schmid, 2016, Sticky leverage, *American Economic Review* 106, 3800–3828.

Gomez-Cram, R., 2022, Late to recessions: Stocks and the business cycle, *The Journal of Finance* 77, 923–966.

Gorodnichenko, Y., and B. Lee, 2020, Forecast error variance decompositions with local projections, *Journal of Business & Economic Statistics* 38, 921–933.

Goyal, A., and I. Welch, 2007, A Comprehensive Look at The Empirical Performance of Equity Premium Prediction, *The Review of Financial Studies* 21, 1455–1508.

Greenwood, R., and A. Shleifer, 2014, Expectations of Returns and Expected Returns, *Review of Financial Studies* hht082, 00087.

Gu, S., B. Kelly, and D. Xiu, 2020, Empirical asset pricing via machine learning, *The Review of Financial Studies* 33, 2223–2273.

Hamilton, J. D., 1983, Oil and the macroeconomy since world war ii, *Journal of Political Economy* 91, 228–248.

Hausman, J. A., 1978, Specification tests in econometrics, *Econometrica* 46, 1251–1271.

Jensen, T. I., 2023, Subjective risk and return, *Available at SSRN 4276760* .

Jordà, Ò., 2005, Estimation and inference of impulse responses by local projections, *American Economic Review* 95, 161–182.

Jurado, K., S. C. Ludvigson, and S. Ng, 2015, Measuring uncertainty, *American Economic Review* 105, 1177–1216.

Koijen, R. S., and S. Van Nieuwerburgh, 2011, Predictability of returns and cash flows, *Annu. Rev. Financ. Econ.* 3, 467–491.

Kurmann, A., and C. Otrok, 2013, News shocks and the slope of the term structure of interest rates, *American Economic Review* 103, 2612–32.

Kurmann, A., and E. Sims, 2021, Revisions in utilization-adjusted tfp and robust identification of news shocks, *Review of Economics and Statistics* 103, 216–235.

Leduc, S., and Z. Liu, 2016, Uncertainty shocks are aggregate demand shocks, *Journal of Monetary Economics* 82, 20–35.

Lettau, M., and S. Ludvigson, 2001, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238–1287.

Liao, Y., X. Ma, A. Neuhierl, and Z. Shi, 2023, Economic forecasts using many noises, *arXiv preprint arXiv:2312.05593* .

Ludvigson, S. C., S. Ma, and S. Ng, 2021, Uncertainty and business cycles: Exogenous impulse or endogenous response?, *American Economic Journal: Macroeconomics* 13, 369–410.

Ludvigson, S. C., and S. Ng, 2009, Macro factors in bond risk premia, *Review of Financial Studies* 22, 5027–5067.

Manela, A., and A. Moreira, 2017, News implied volatility and disaster concerns, *Journal of Financial Economics* 123, 137–162.

Martin, I., 2017, What is the Expected Return on the Market?*, *The Quarterly Journal of Economics* 132, 367–433.

McDonald, R., and D. Siegel, 1986, The value of waiting to invest, *The Quarterly Journal of Economics* 101, 707–728.

Møller, S. V., T. Q. Pedersen, and S. Steffensen, 2020, Countercyclical expected returns, *Available at SSRN 3555481* .

Montiel Olea, J. L., and C. Pflueger, 2013, A robust test for weak instruments, *Journal of Business & Economic Statistics* 31, 358–369.

Nagel, S., 2024, Stock market valuation: Explanations, non-explanations, and some open questions, Technical report, Technical Report, University of Chicago.

Nagel, S., and Z. Xu, 2022, Asset Pricing with Fading Memory, *The Review of Financial Studies* 35, 2190–2245.

Nagel, S., and Z. Xu, 2023, The Dynamics of Subjective Risk Premia, *Journal of Financial Economics* Forthcoming.

Nakamura, E., D. Sergeyev, and J. Steinsson, 2017, Growth-rate and uncertainty shocks in consumption: Cross-country evidence, *American Economic Journal: Macroeconomics* 9, 1–39.

Pastor, L., and P. Veronesi, 2013, Political uncertainty and risk premia, *Journal of Financial Economics* 110, 520–545.

Patton, A. J., and A. Timmermann, 2010, Monotonicity in asset returns: new tests with applications to the term structure, the capm, and portfolio sorts, *Journal of Financial Economics* 98, 605–625.

Patton, A. J., and A. Timmermann, 2011, Predictability of output growth and inflation: A multi-horizon survey approach, *Journal of Business & Economic Statistics* 29, 397–410.

Pettenuzzo, D., A. Timmermann, and R. Valkanov, 2014, Forecasting stock returns under economic constraints, *Journal of Financial Economics* 114, 517–553.

Politis, D. N., and J. P. Romano, 1994, The stationary bootstrap, *Journal of the American Statistical Association* 89, 1303–1313.

Polk, C., S. Thompson, and T. Vuolteenaho, 2006, Cross-sectional forecasts of the equity premium, *Journal of Financial Economics* 81, 101–141.

Ramey, V. A., 2016, Macroeconomic shocks and their propagation, in *Handbook of macroeconomics*, volume 2, 71–162 (Elsevier).

Rapach, D. E., J. K. Strauss, and G. Zhou, 2009, Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy, *The Review of Financial Studies* 23, 821–862.

Stock, J., and M. Watson, 2016, Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics, in *Handbook of macroeconomics*, volume 2, 418–517 (Elsevier).

Van Binsbergen, J., and R. Koijen, 2010, Predictive regression: A present-value approach, *Journal of Finance* 65, 1439–1471.

Van Nieuwerburgh, S., and L. Veldkamp, 2006, Learning asymmetries in real business cycles, *Journal of Monetary Economics* 53, 753–772.

Wooldridge, J. M., 2015, Control function methods in applied econometrics, *The Journal of Human Resources* 50, 420–445.

(A) Realized returns

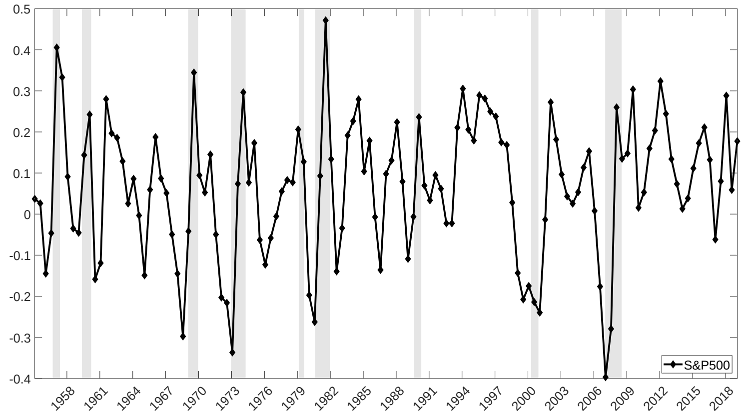

(B) Forecasts of returns

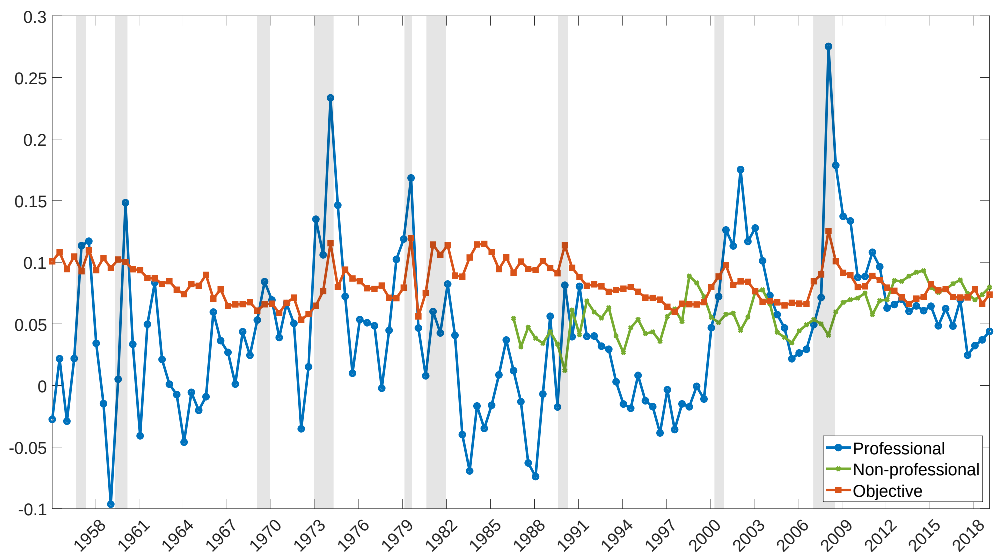

FIGURE 1: **Realized and Forecasted Stock Returns**
Panel A of this figure presents the quarterly time series of one-year ahead returns on the S&P500 index in excess of the one year T-bill rate. Panel B presents the forecasts we study, that is, the professional forecast (consensus average from the Livingston survey, sampled biannually), non-professional forecast (from Nagel and Xu, 2022, sampled quarterly) and the objective forecast (a combination forecast constructed as an equal-weighted average of out-of-sample predictions from 22 individual predictors, sampled quarterly). Shaded areas are NBER recessions.

(A) Cycle

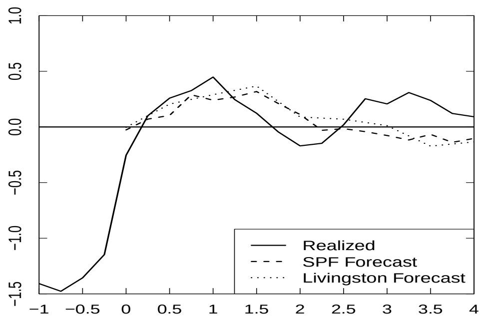

(B) TFP

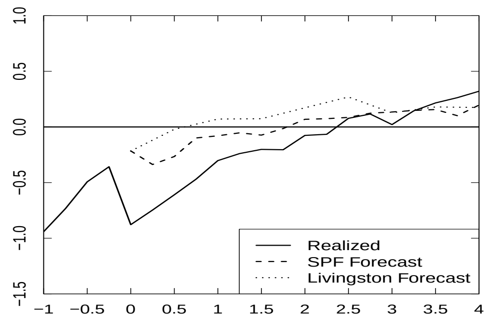

(C) Oil

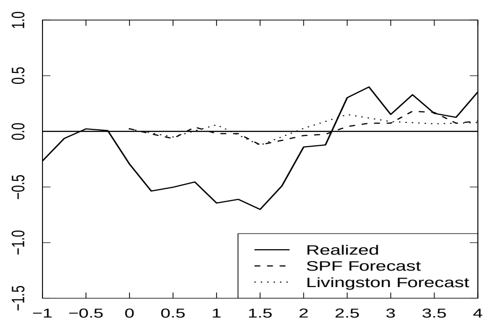

FIGURE 2: **Impulse Responses of GDP and Forecasts to Macro Shocks.**
Analogous to Figure 3, this figure presents the impulse response function for realized one-year ahead real GDP growth as well as consensus forecasts of this quantity from two different surveys: SPF and Livingston. The responses are estimated as described in detail in Appendix D. Following Bianchi, Ludvigson, and Ma (2022), the responses for realized GDP have four additional points (indexed $t-1$ to $t-\frac{1}{4}$) and serve to show the response of realized GDP growth over annual periods that include the quarter the shock hits (t), but for which forecasters were not aware of the shock yet.

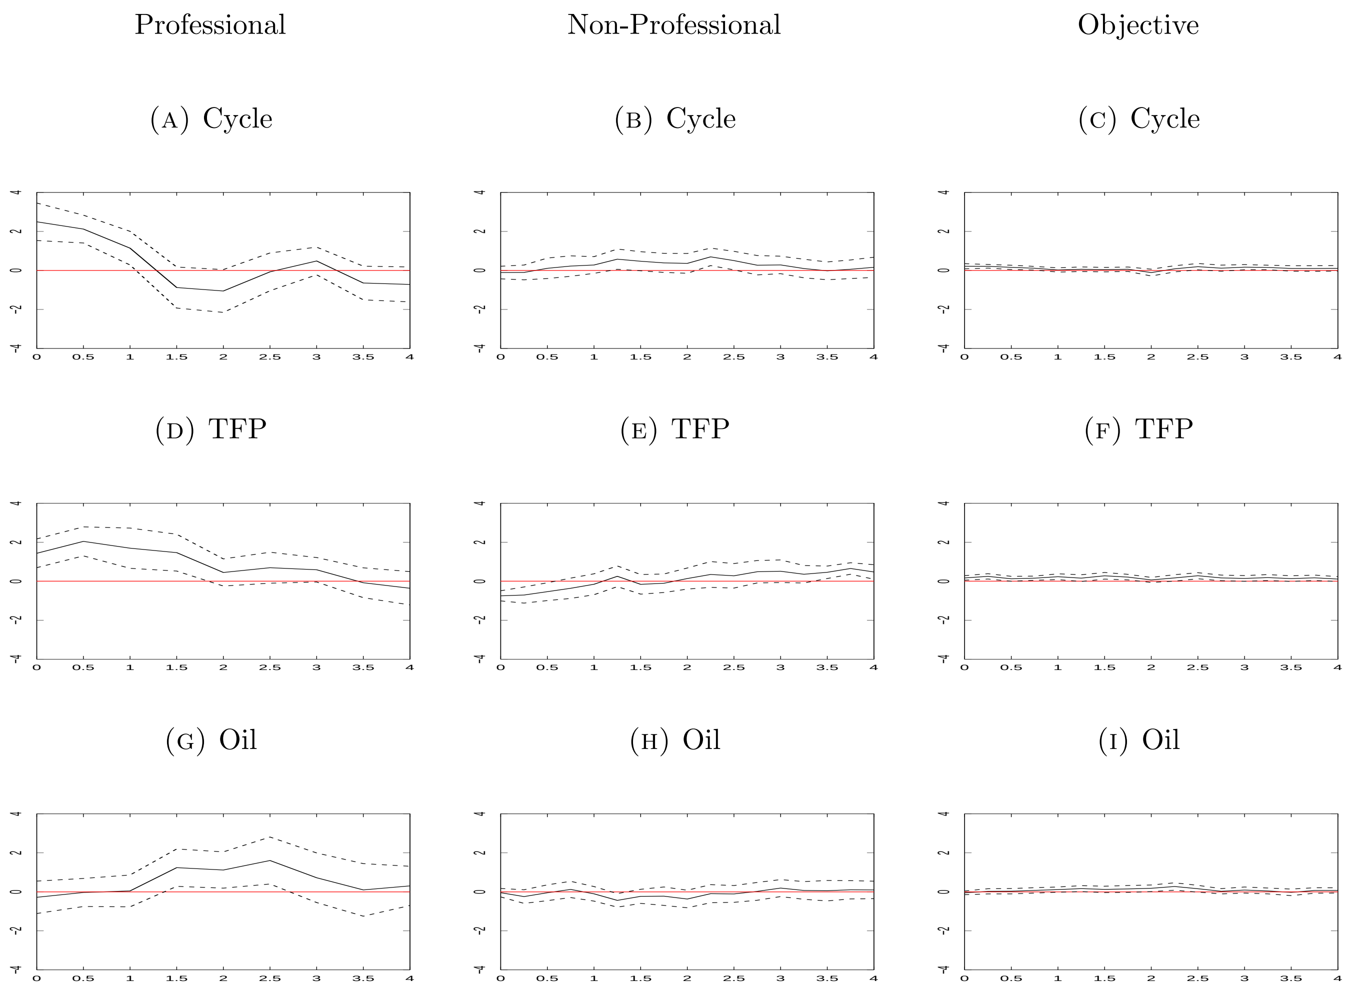

Figure 3: **Impulse Responses of Stock Return Forecasts to Macro Shocks.**
This figure presents the impulse response function (with 90%-confidence interval constructed using Newey-West standard errors with lag length set to one year) of three types of stock return forecasts (subjective forecasts by professionals and non-professionals as well as an objective forecast) to three different macro shocks (Cycle, TFP and Oil). The responses are estimated using the local linear projections of Eq. (4) with biannual data for professionals and quarterly data in the other two cases. The projections control for one year worth of lags of stock returns, the forecast and the shock. For the case of professionals, each biannual shock sums two quarterly shocks. To make the responses comparable across the biannual and quarterly frequency, we use the standard deviation of the biannual shocks to scale all responses. For the business cycle shock, the sample runs from 1955Q3 to 2017Q4; for TFP news, the sample runs from 1962Q1 to 2019Q4; for the oil supply shock, the sample runs from 1975Q1 to 2019Q4.

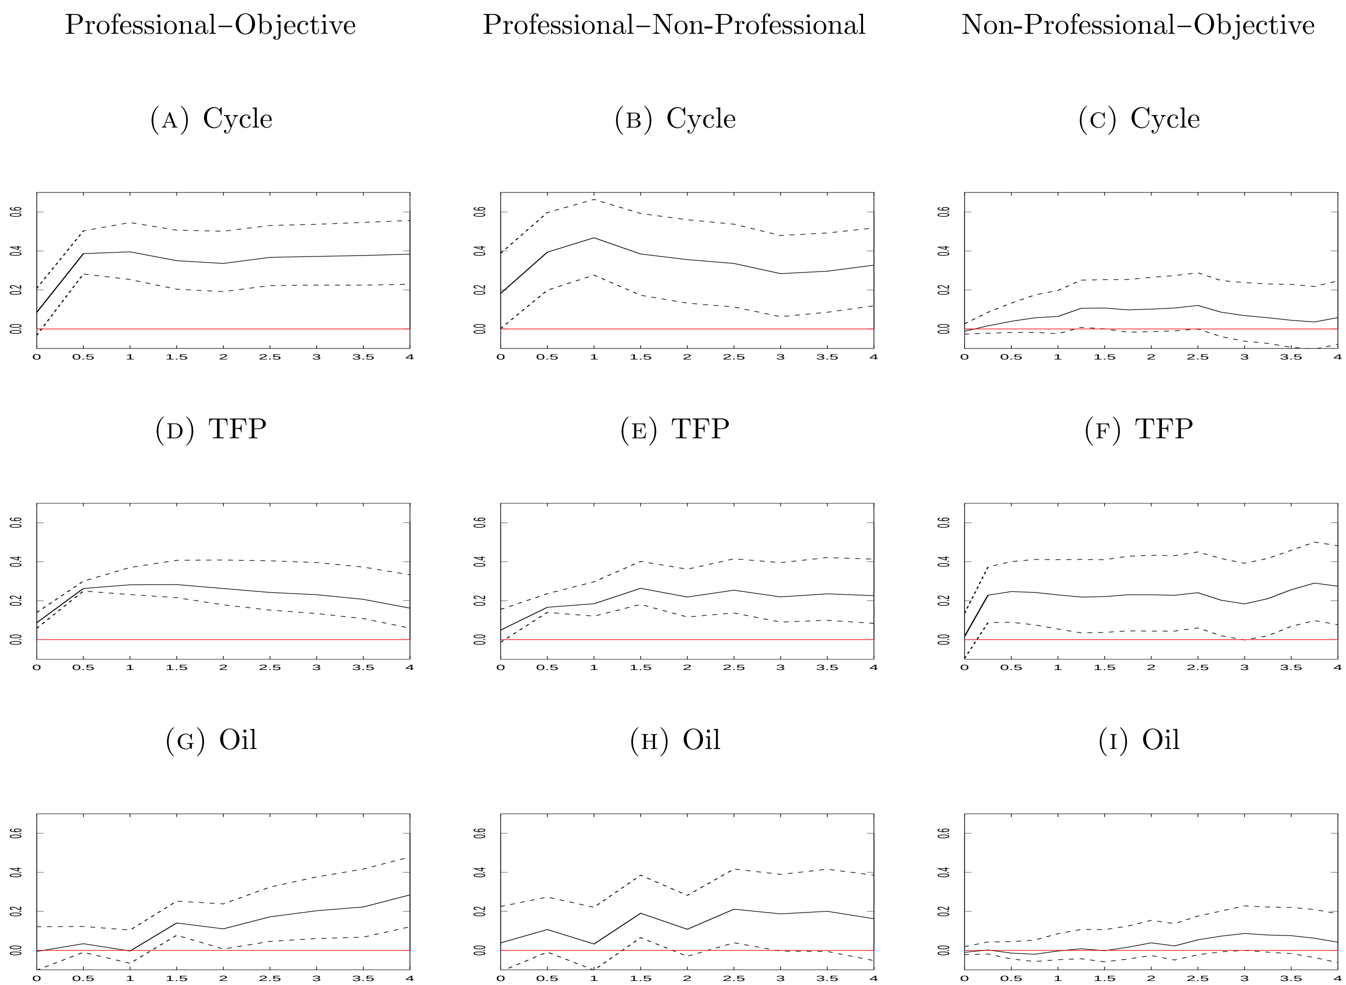

FIGURE 4: **Forecast Error Variance Decomposition for Differences in Forecasts.**
This figure presents forecast error variance decompositions based on local projections of the differences in stock return forecasts (professional versus objective, professional versus non-professional, and non-professional versus objective forecasts) on three different macro shocks (Cycle, TFP and Oil). Following Gorodnichenko and Lee (2020), we compute the 90%-confidence interval and bias-adjustment using a VAR-based bootstrap. For the business cycle shock, the sample runs from 1955Q3 to 2017Q4; for TFP news, the sample runs from 1962Q1 to 2019Q4; for the oil supply shock, the sample runs from 1975Q1 to 2019Q4.

(A) Cycle

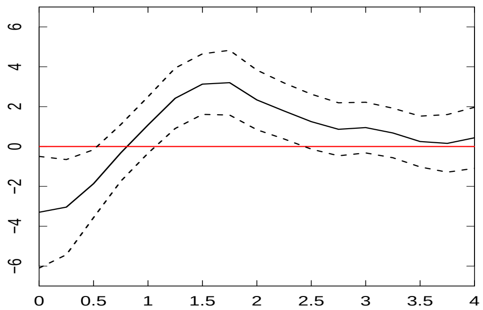

(B) TFP

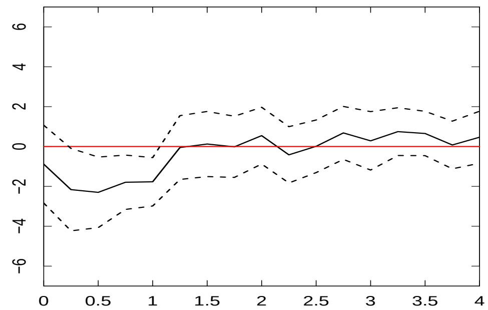

(C) Oil

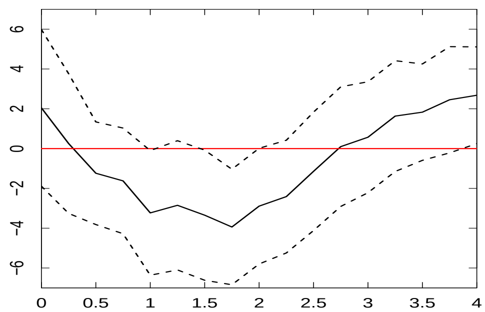

FIGURE 5: **Impulse Responses of Dividend Growth to Macro Shocks**
This figure is analogous to Figure 3 and presents the impulse response function of realized one-year ahead repurchase-adjusted dividend growth to three macro shocks (Cycle, TFP and Oil).

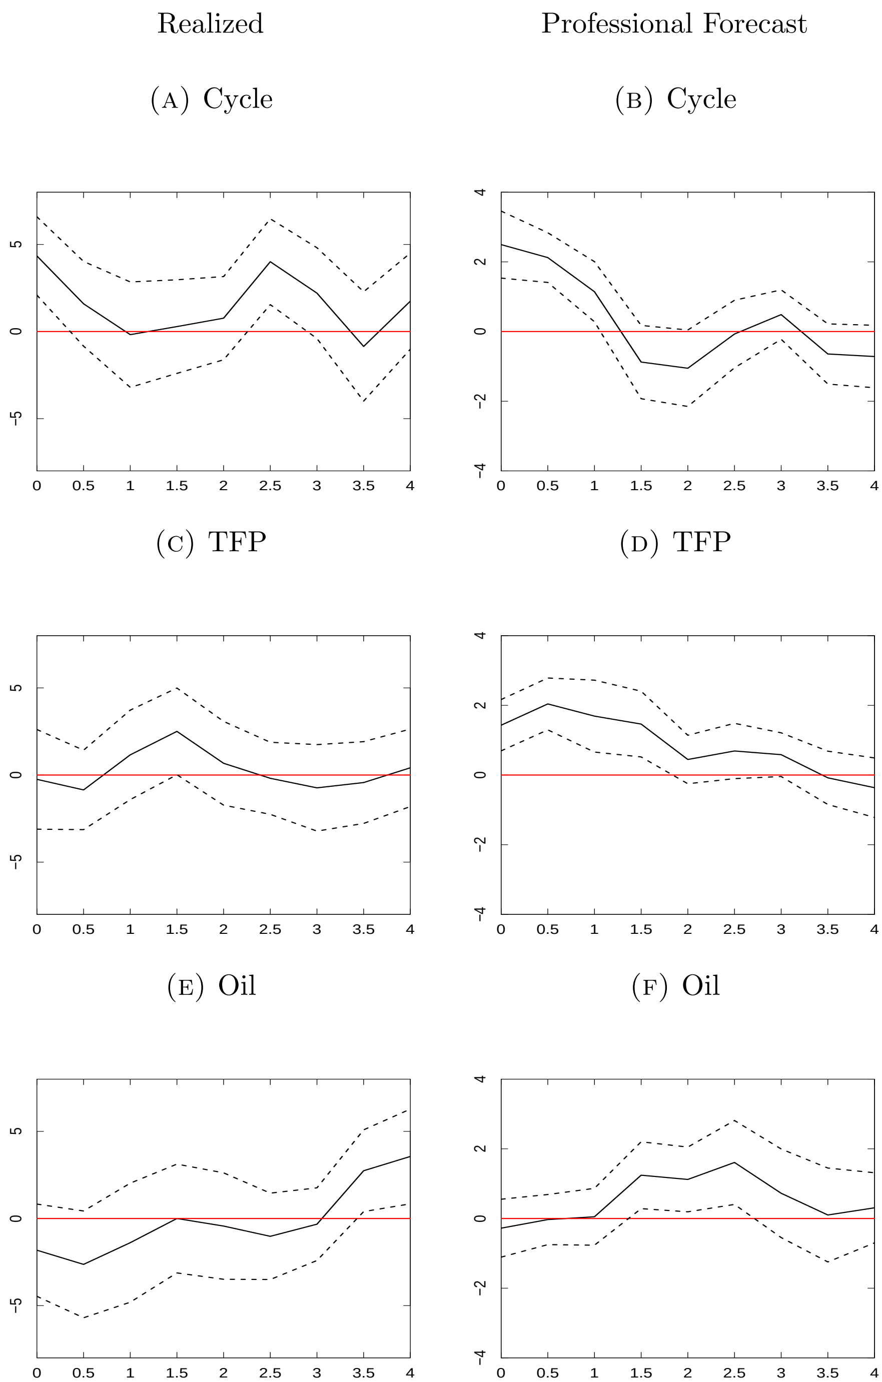

Figure 6: **Impulse Responses for Professional Forecasts and Realized Returns** This figure presents the impulse response function (with 90%-confidence interval constructed using Newey-West standard errors with lag length set to one year) of realized one-year ahead excess stock returns to three macro shocks (Cycle, TFP and Oil). The responses of professional forecasts on the right are reproduced from Figure 3. The responses are estimated using the local linear projections of Eq. (4) with biannual data. The projections control for one year worth of lags of stock returns, the forecast and the shock. Each biannual shock sums two quarterly shocks.

(A) Returns on Shock-Insensitive Predictors

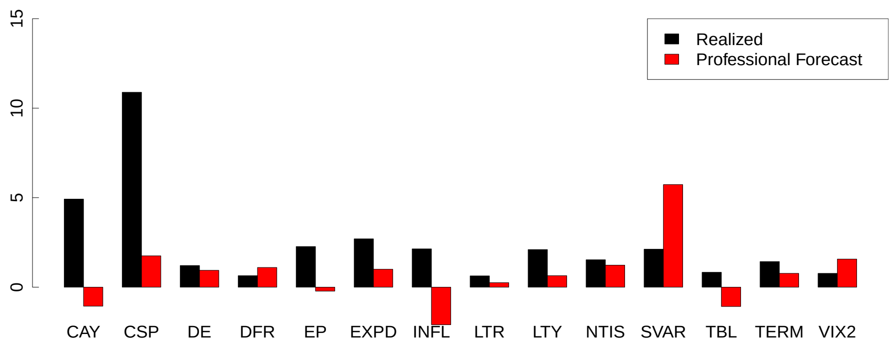

(B) Returns on Shock-Sensitive Predictors

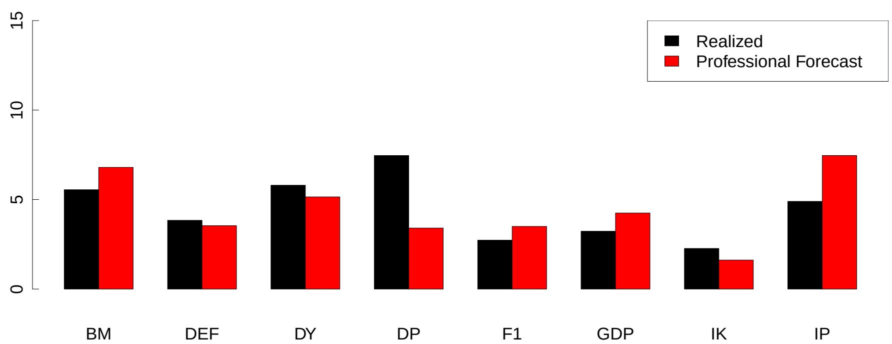

(C) Returns on Instrumented Shock-Sensitive Predictors

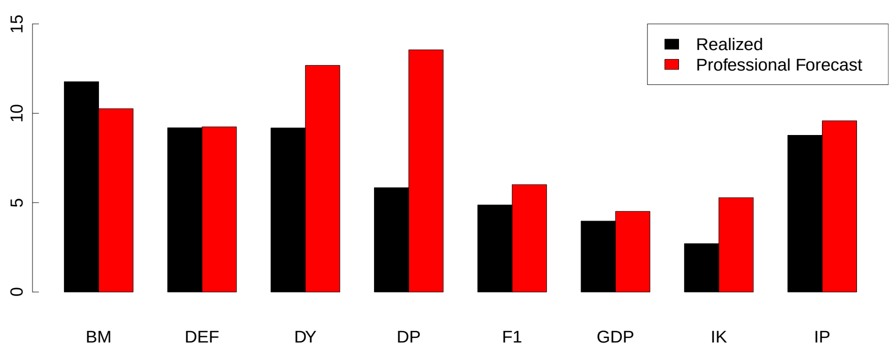

FIGURE 7: **(Forecasts of) Returns on Predictors.**
This figure summarizes the main insights from two types of regressions of returns on a large set of stock return predictors (see Appendix A). We differentiate between those predictors that are "shock-insensitive" ("shock-sensitive"), for which the first-stage regression in Eq. (8) is weak (strong), meaning the effective F-statistic of Montiel Olea and Pflueger (2013) is not (is) significant at the 10%-level. In Panels A and B, respectively, we report the coefficients from the regression in Eq. (9) of realized one year ahead excess stock returns and professional forecasts thereof on these two sets of predictors. In Panel C, we focus on the shock-sensitive predictors and report the coefficient from the second-stage regression in Eq. (11) of returns on the predictor instrumented by Cycle and TFP shocks. The predictors are standardized so that all coefficients measure percentage points.

Table I: **Summary Statistics**

This table presents summary statistics of realized one year ahead excess stock returns as well as three different forecasts of that quantity (professional, non-professional and objective). Panel A reports the mean, standard deviation, sample period and frequency for each series. We report also the OOS R^2 for each forecast, which is calculated as $1 - \frac{\sum_t (R_{t+1} - F_t(R_{t+1}))^2}{\sum_t (R_{t+1} - \overline{R_{t+1}})^2}$. Here, $F_t(R_{t+1})$ is a forecast and $\overline{R_{t+1}}$ is the historical expanding window mean of realized returns R_{t+1}. Panel B reports the correlation matrix. Panel C reports the differences in (forecasts of) returns between NBER recessions and expansions.

Panel A: Summary Statistics

	Professional	Non-Prof.	Objective	Realized
Average	4.187	6.048	8.276	8.542
Standard Dev.	5.942	1.817	1.897	21.389
OOS R^2	2.425	-3.818	8.002	
Sample Period	1952Q2 - 2020Q4	1987Q2 - 2020Q4	1937Q4 - 2020Q4	1926Q4 - 2020Q4
Frequency	Biannual	Quarterly	Quarterly	Monthly

Panel B: Correlations

	Professional	Non-Prof.	Objective	Realized
Professional	1.000	0.199	0.103	0.167
Non-Prof.		1.000	0.012	-0.025
Objective			1.000	0.294
Realized				1.000

Panel C: Recessions versus Expansions

	Professional	Non-Prof.	Objective	Realized
Recession	11.946	4.468	8.512	12.497
Expansion	3.155	6.103	7.362	5.659
Difference	8.790	-1.635	1.150	6.838
t-statistic	(5.313)	(-3.805)	(3.636)	(1.331)

Table II: **Real Uncertainty on Macro Shocks (First Stage)**

This table presents results from the first-stage regression in Eq. (8) of two measures of real uncertainty on macro shocks. The measures are the real uncertainty index of Ludvigson, Ma, and Ng (2021, *Unc* in Panel A) and the dispersion in Livingston GDP forecasts (*Disp* in Panel B). Across columns we consider each shock in isolation as well as two combinations of shocks (Cycle and TFP in column four and all three shocks in column five). We report the estimated coefficients and t-statistics based on Newey-West standard errors with 4 lags. We also report the R^2, the number of years used in the regression, the first-stage effective F-statistic of Montiel Olea and Pflueger (2013, with ***,**,* denoting significance at the 1, 5, and 10-% level, respectively), and the optimal lag choice κ, which is used to select the four consecutive quarters over which the shocks are cumulated. Each regression controls for lagged real uncertainty and realized stock returns.

Panel A: *Unc*

	Cycle	TFP	Oil	Cycle+TFP	Cycle+TFP+Oil
Cycle	0.28			0.18	0.18
	(3.12)			(2.07)	(2.10)
TFP		0.40		0.33	0.32
		(5.58)		(4.60)	(4.28)
Oil			0.12		0.05
			(1.52)		(0.64)
Adj.R^2	0.45	0.37	0.20	0.40	0.40
# years	58	56	43	55	55
κ^*	0/4	3/4	3/4	3/4	3/4
Eff. F-stat.	9.65	32.45***	2.42	18.03***	13.08*

Panel B: *Disp*

	Cycle	TFP	Oil	Cycle +TFP	Cycle+TFP+Oil
Cycle	0.30			0.23	0.23
	(3.32)			(2.51)	(2.51)
TFP		0.32		0.25	0.28
		(3.81)		(2.78)	(2.99)
Oil			-0.09		-0.15
			(-0.91)		(-1.62)
Adj.R^2	0.10	0.11	0.02	0.15	0.17
# years	58	56	44	55	58
κ^*	1/4	5/4	2/4	2/4	2/4
Eff. F-stat.	10.79	14.89	0.86	10.61***	8.14

Table III: **(Forecasts of) Returns on Instrumented Real Uncertainty (Second Stage).**

This table presents results from regressing (forecasts of) returns on the two measures of real uncertainty (see Eq. (9), first column in each block of two columns) as well as on real uncertainty instrumented by the Cycle and TFP shock (the second stage regression reported in Eq. (11) second column in each block). In this second stage regression, we also include the residual component of real uncertainty that is uncorrelated to the shocks. The shocks are cumulated over 4 consecutive quarters as explained in Section 4.1. We also report the coefficients of a regression of forecast errors on instrumented real uncertainty (λ_S in Eq. (11)), which is simply the difference in the coefficients between forecasts (e.g., $F_t^P(R_{t+1})$ for professionals) and realized returns (R_{t+1}). In the regression of (forecasts) of returns on real uncertainty, the t-statistics that are presented below each estimated coefficient are based on Newey-West standard errors with four lags. For the second stage regression, t-statistics are based on standard errors derived from a stationary bootstrap. We also report the adjusted R^2 and the number of years in the sample. Each regression controls for lagged real uncertainty and realized stock returns.

Panel A: *Unc*

	R_{t+1}		$F_t^P(R_{t+1})$		$F_t^{NP}(R_{t+1})$		$F_t^O(R_{t+1})$	
Unc	0.71		3.60		-0.03		0.34	
	(0.32)		(2.92)		(-0.09)		(3.46)	
$\widehat{Unc}^{Shock}$		8.16		6.79		-0.38		0.64
		(1.50)		(2.54)		(-0.40)		(2.37)
$\widehat{Unc}^{Res.}$		-2.69		2.08		0.03		0.21
		(-1.05)		(1.83)		(0.05)		(1.84)
Forecast Error on $\widehat{Unc}^{Shock}$				1.37		8.54		7.52
				(0.54)		(3.11)		(3.04)
Adj. R^2	-0.01	0.04	0.24	0.29	0.01	0.00	0.30	0.32
# years	55	55	54	54	32	32	55	55

Panel B: *Disp*

	R_{t+1}		$F_t^P(R_{t+1})$		$F_t^{NP}(R_{t+1})$		$F_t^O(R_{t+1})$	
$Disp$	2.28		0.64		0.04		0.38	
	(1.32)		(1.10)		(0.14)		(3.35)	
$\widehat{Disp}^{Shock}$		5.58		6.15		0.34		0.77
		(1.22)		(2.11)		(0.35)		(3.34)
$\widehat{Disp}^{Res.}$		1.61		-0.45		0.00		0.31
		(0.90)		(-0.71)		(0.00)		(2.78)
Forecast Error on $\widehat{Disp}^{Shock}$				-0.57		5.24		4.81
				(-0.29)		(2.87)		(2.75)
R^2	0.00	0.01	56 0.03	0.17	-0.02	-0.03	0.19	0.22
# years	55	55	55	55	32	32	55	55

Table IV: **Stock Return Predictors on Macro Shocks (First Stage).**
This table presents results from the first-stage regression in Eq. (8). We regress a large set of 22 stock return predictors on Cycle and TFP shocks. We report the estimated coefficients and t-statistics based on Newey-West standard errors with 4 lags. We also report the R^2, the number of years used in the regression, the first-stage effective F-statistic of Montiel Olea and Pflueger (2013, with ***,**,* denoting significance at the 1, 5, and 10-% level, respectively), and the optimal lag choice κ, which is used to select the four consecutive quarters over which the shocks are cumulated. Each regression controls for the lag of the predictor and realized stock returns. The last row indicates which predictors have been signed to accommodate interpretation. After signing, all predictors predict stock market returns with a positive sign.

	BM	CAY	CSP	DE	DEF	DFR	DP	DY
Cycle	0.14	0.01	0.14	0.23	0.29	-0.16	0.11	0.13
	(2.45)	(0.14)	(2.07)	(1.61)	(-3.19)	(2.00)	(1.54)	(3.36)
TFP	0.19	-0.22	-0.09	-0.05	0.17	0.07	0.19	0.07
	(4.69)	(-2.87)	(-1.86)	(-0.86)	(2.55)	(1.02)	(3.39)	(2.43)
Adj.R^2	0.82	0.54	0.25	0.03	0.33	0.01	0.58	0.88
# years	55	55	40	55	55	57	56	55
κ^*	0/4	4/4	1/4	1/4	5/4	1/4	4/4	5/4
Eff. F-stat.	14.67***	5.33	2.69	1.78	11.44**	2.96	10.15***	16.27***
Signed			✓			✓		

	EP	EXPD	F1	GDP	IK	INFL	LTR	LTY
Cycle	-0.15	0.39	0.50	0.75	0.60	0.04	-0.14	-0.09
	(-1.21)	(3.27)	(5.33)	(18.54)	(6.10)	(0.89)	(-1.86)	(-1.57)
TFP	0.19	0.03	0.21	0.17	0.08	-0.26	-0.18	0.20
	(2.08)	(0.48)	(2.81)	(5.33)	(1.30)	(-3.45)	(-2.83)	(2.95)
Adj.R^2	0.37	0.21	0.34	0.84	0.45	0.27	0.05	0.77
# years	56	56	56	55	55	57	57	55
κ^*	5/4	0/4	5/4	5/4	1/4	0/4	0/4	0/4
Eff. F-stat.	2.37	10.12	27.12***	343.47***	29.04***	10.37	5.95**	6.15
Signed		✓		✓	✓	✓		

	IP	NTIS	SVAR	TBL	TERM	VIX2
Cycle	0.35	-0.22	-0.02	-0.34	0.61	0.20
	(8.05)	(-2.72)	(-0.80)	(-8.71)	(7.97)	(1.49)
TFP	0.09	-0.02	0.10	0.30	-0.56	0.05
	(2.87)	(-0.29)	(1.91)	(6.43)	(-7.36)	(0.83)
Adj.R^2	0.63	0.18	0.04	0.83	0.56	0.11
# years	55	56	57	56	56	56
κ^*	5/4	4/4	2/4	0/4	3/4	2/4
Eff. F-stat.	57.87***	4.91	3.41	35.86***	42.30***	2.09
Signed	✓	✓				

Table V: **(Forecasts of) Returns on Instrumented Predictors (Second Stage).** This table presents results from regressing (forecasts of) returns on shock-sensitive predictors (see Eq. (9), first column in each block of two columns) as well as on these predictors instrumented by the Cycle and TFP shock (the second stage regression reported in Eq. (11), second column in each block). In this second stage regression, we also include the residual component of the predictor that is uncorrelated to the shocks. The shocks are cumulated over 4 consecutive quarters as explained in Section 4.1. Each regression controls for lags of the predictor and realized stock returns. We also report the coefficients of a regression of forecast errors on instrumented real uncertainty (λ_S in Eq. (11)), which is simply the difference in the coefficients between forecasts (e.g., $F_t^P(R_{t+1})$ for professionals) and realized returns (R_{t+1}). In the regression of (forecasts) of returns on the predictor, the t-statistics that are presented below each estimated coefficient are based on Newey-West standard errors with four lags. For the second stage regression, t-statistics are based on standard errors derived from a stationary bootstrap. We also report the adjusted R^2 and the number of years in the sample.

	BM				DEF			
	R_{t+1}		$F_t^P(R_{t+1})$		R_{t+1}		$F_t^P(R_{t+1})$	
X	5.55		6.80		3.84		3.54	
	(1.74)		(6.48)		(1.60)		(3.05)	
$\widehat{X^{Shock}}$		11.77		10.26		9.19		9.24
		(1.29)		(2.69)		(1.82)		(2.23)
$\widehat{X^{Res.}}$		3.35		5.60		1.61		1.10
		(0.82)		(5.06)		(0.55)		(1.01)
Forecast Error on $\widehat{X^{Shock}}$				1.51				-0.05
				(0.18)				(-0.01)
Adj. R^2	0.03	0.03	0.37	0.40	0.03	0.04	0.14	0.29
# years	55	55	54	54	55	55	54	54

	DY				DP			
	R_{t+1}		$F_t^P(R_{t+1})$		R_{t+1}		$F_t^P(R_{t+1})$	
X	5.80		5.15		7.46		3.41	
	(1.25)		(4.35)		(2.74)		(2.96)	
$\widehat{X^{Shock}}$		9.18		12.68		5.84		13.55
		(0.73)		(3.02)		(0.96)		(2.24)
$\widehat{X^{Res.}}$		4.82		3.23		7.75		1.62
		(0.90)		(2.16)		(2.44)		(1.73)
Forecast Error on $\widehat{X^{Shock}}$				-3.50				-7.71
				(-0.27)				(-0.80)
Adj. R^2	0.03	0.03	0.25	0.30	0.09	0.09	0.25	0.44
# years	55	55	55	55	56	56	55	55

	F1				GDP			
	R_{t+1}		$F_t^P(R_{t+1})$		R_{t+1}		$F_t^P(R_{t+1})$	
X	2.73		3.50		3.23		4.25	
	(1.75)		(5.18)		(1.48)		(5.99)	
$\widehat{X^{Shock}}$		4.87		6.01		3.97		4.51
		(1.74)		(4.83)		(1.74)		(6.09)
$\widehat{X^{Res.}}$		1.62		2.00		-0.77		2.82
		(0.76)		(2.33)		(-0.12)		(1.73)
Forecast Error on $\widehat{X^{Shock}}$				-1.14				-0.54
				(-0.35)				(-0.23)
Adj. R^2	0.01	0.02	0.39	0.49	0.01	0.01	0.38	0.38
# years	56	56	55	55	55	55	54	54

	IK				IP			
	R_{t+1}		$F_t^P(R_{t+1})$		R_{t+1}		$F_t^P(R_{t+1})$	
X	2.27		1.62		4.90		7.46	
	(1.31)		(1.98)		(1.36)		(6.93)	
$\widehat{X^{Shock}}$		2.71		5.28		8.77		9.58
		(0.76)		(4.36)		(1.70)		(5.41)
$\widehat{X^{Res.}}$		1.99		-0.73		-1.52		4.02
		(0.87)		(-0.84)		(-0.25)		(1.79)
Forecast Error on $\widehat{X^{Shock}}$				2.57				0.81
				(0.69)				(0.15)
Adj. R^2	0.06	0.06	0.08	0.28	0.00	0.02	0.33	0.37
# years	55	55	55	55	55	55	54	54

Table VI: **(Forecasts of) Returns in Bad Shock Regimes**
This table presents average one-year ahead (from the end of quarter t to $t+1$, i.e., t plus one year) realized stock returns and forecasts of this quantity (made in quarter t) separated between bad shock quarters and normal times. The bad shock dummy is equal to one in quarter t when one or more of the macro shocks (suitably cumulated over a one-year period) are above their 90th percentile. We consider separately the case of Cycle and TFP shocks and the case with all three shocks in Panels A and B, respectively. We also present the difference with t-statistics based on Newey-West standard errors with four lags.

Panel A: Cycle + TFP				
	Professional	Non-Prof.	Objective	Realized
Bad shock	10.842	5.453	8.586	14.546
Normal	3.116	6.008	7.318	5.107
Difference	7.726	-0.555	1.268	9.438
t-statistic	(4.138)	(-1.123)	(4.932)	(2.219)
Panel B: Cycle + TFP + Oil				
	Professional	Non-Prof.	Objective	Realized
Bad shock	9.335	5.897	8.198	8.666
Normal	2.985	5.981	7.323	5.956
Difference	6.351	-0.084	0.875	2.710
t-statistic	(3.784)	(-0.186)	(3.541)	(0.639)

Internet Appendix

Appendix A. Definition of stock return predictors

We study a total of 22 predictors. Unless otherwise specified, the predictors span the full post-war sample studied in this paper. Fifteen of these predictors are downloaded from Amit Goyal's website and defined as in Goyal and Welch (2007):

- **BM (book-to-market):** the ratio of the book-to-market value of the Dow Jones Industrial Average.
- **CSP (cross-sectional premium):** the spread in returns between high and low beta stocks, as defined in Polk, Thompson, and Vuolteenaho (2006). This series is available until December 2002.
- **DE (dividend-to-earnings):** the spread between log dividends (defined as 12-month moving sums of dividends paid on the S&P 500 index) and log earnings (defined as 12-month moving sums of earnings on the S&P 500 index).
- **DEF (default spread):** the BAA-minus-AAA corporate bond yield spread.
- **DFR (default return spread):** the difference between long-term corporate bond and long-term government bond returns.
- **DY (dividend yield):** the difference between the log of dividends and the log of lagged prices.
- **EP (earnings-to-price ratio):** the difference between the log of earnings and the log of prices.
- **IK (investment-to-capital ratio):** the ratio of aggregate (private nonresidential fixed) investment to aggregate capital for the whole economy, as defined in Cochrane (1991).
- **INFL (inflation):** is the percentage change Consumer Price Index (All Urban Consumers) from the Bureau of Labor Statistics

- **LTR (long term rate of returns):** the long term bond returns provided by Ibbotson's *Stocks, Bonds, Bills and Inflation* . Yearbook

- **LTY (long term yield):** yield on long term government bonds from Ibbotson's *Stocks, Bonds, Bills and Inflation* .

- **NTIS (net equity expansion):** ratio of 12-month moving sums of net issues by NYSE listed firms divided by the total end-of-year NYSE market capitalization.

- **SVAR (stock variance):** the sum of squared daily returns on the S&P 500.

- **TBL (Tresury bill):** the yield on 3 month Treasury bills.

- **TERM (term spread):** the difference between the long term yield on government bonds and the Treasury bill.

The remaining seven predictors are defined as follows:

- **CAY:** the log consumption-wealth ratio of Lettau and Ludvigson (2001), obtained from Sydney Ludvigson's website. The sample period is March 1952 to September 2019.

- **DP (dividend-to-price ratio):** the repurchase-adjusted log dividend price ratio of the CRSP value-weighted index. Obtained from the replication package of Nagel and Xu (2022).

- **EXPD :** the long-run exponential average of past per-capita real aggregate dividend growth. Obtained from the replication package of Nagel and Xu (2022).

- **F1 (real factor):** the first principal component of a broad set of macroeconomic variables as used in Ludvigson and Ng (2009). Obtained from Sydney Ludvigson's website and the sample period is April 1960 to December 2020.

- **GDP (real GDP growth):** log yearly real GDP growth (based on the series GDPC1 from the FRED website).

- **IP (industrial production):** log yearly industrial production growth (based on the series INDPRO from the FRED website).

- **VIX2 (VIX squared):** is the square of the CBOE Volatility index, which starts in January 1990. For the earlier period, we follow Nagel and Xu (2023) and extend the series with the NVIX index constructed by Manela and Moreira (2017). We estimate the relation between VIX and NVIX in the overlapping sample and impute VIX pre-1990 with the fitted values from this regression.

Appendix B. Negative excess return forecasts

In this paper we focus on a component of professionals' stock return forecasts that reflects remarkably well the discount rate news contained in macro shocks. Our claim is not that professionals are always right and we leave for future work the analysis of other interesting components of their forecast. For instance, professionals forecast negative returns about 25% of the time. What drives these negative forecasts and how do they compare to positive forecasts?

To answer these questions, we first calculate the out-of-sample R^2 for all quarters that the professional forecast is positive versus negative. As shown in Table B.1 below, the professionals' forecast generates a large out-of-sample R^2 of 11.9%, almost twice the R^2 of the objective forecast (6.2%), in the 96 quarters in which their forecast is positive. When their forecast is negative, however, the fit is much worse than the objective forecast at an R^2 of -28.7% versus 14.1%. This finding is consistent with the idea that these negative forecasts are disconnected from the macro fundamentals that are key in our paper. Indeed, time-varying discount rates cannot explain why the professionals would sometimes forecast an excess return below -5%.

To see this more directly, we now estimate the second stage of Eq. (12) (repeated here

Table B.1: **OOS R^2 Conditional on Sign of the Professionals' Forecast**
This table presents the OOS R^2 for the professionals' and the objective forecast in quarters when the professionals' forecast is positive versus negative. In the former case, this R^2 is calculated as $1 - \frac{\sum_{t \in F_t^P(R_{t+1})>0}(R_{t+1}-F_t(R_{t+1}))^2}{\sum_{t \in F_t^P(R_{t+1})>0}(R_{t+1}-\overline{R_{t+1}})^2}$. Here, $F_t(R_{t+1})$ is a forecast and $\overline{R_{t+1}}$ is the historical expanding window mean of realized returns R_{t+1}.

	Professional	Objective	# of quarters
$F_t^P(R_{t+1}) > 0$	0.119	0.062	96
$F_t^P(R_{t+1}) < 0$	-0.287	0.141	33

for convenience) using quantile regressions:

$$F_t^P(R_{t+1}) = \lambda_S^P \ \widehat{X_{t-\frac{1}{4}}^{Shock}} + \lambda_R^P \ \widehat{X_{t-\frac{1}{4}}^{Res.}} + \gamma^P C_{t-\frac{5}{4}-\kappa^\star} + v_{t+1}^P. \tag{18}$$

If the professionals' forecast adequately represents the discount rate news contained in macro shocks, we expect the sensitivity, λ_S^P, of their forecast to the instrumented discount rate proxy $\widehat{X_{t-\frac{1}{4}}^{Shock}}$ (i.e., the component of X due to macro shocks) to be increasing from the left tail of the distribution of forecasts (i.e., among the most negative return forecasts) to the right tail. Under this assumption, the professionals' forecast will incorporate the positive discount rate impact of bad macro shocks. In turn, their forecast will reflect lower discount rates in the absence of bad shocks or in the presence of good macro shocks. This discount rate story cannot help to explain variation among the most negative of forecasts, however.

Figure B.1 below presents the coefficient estimates (λ_S^P in Panel A and λ_R^P in Panel B) for three quantiles $\tau = 0.05, 0.50, 0.95$ for ten discount rate proxies, including the two uncertainty measures of Section *4.3.a* and the eight shock-sensitive predictors defined in Section *4.3.b*. To start, note that the estimate at the median ($\tau = 0.50$) is close to the OLS estimates reported in Tables III and VI.

Consistent with the idea that professionals adeptly incorporate discount rate news, we see in Panel A that the coefficient estimate λ_S^P is increasing from $\tau = 0.05$ to 0.95 for all ten proxies. For instance, at the 5th quantile, the impact of instrumented real uncertainty (Unc) is 4.8%, which is relative to 8.8% at the 95th quantile. On average across the ten proxies, the coefficient estimate is about 70% larger for $\tau = 0.95$ than for $\tau = 0.05$. Thus, macro shocks explain more variation in return forecasts in the right tail of the distribution than in the left tail. Consistent with the idea that the residual component captures something other than discount rate shocks, we see in Panel B that the coefficient estimate λ_R^P does not vary across the quantiles in a robust manner. Moreover, when we compare the coefficients on the instrumented and residual component for each proxy, we see that the difference increases from the 5th to the 95th quantile. On average across the ten proxies, the difference between λ_S^P and λ_R^P equals 5% for $\tau = 0.05$ and 10% for $\tau = 0.95$. In sum, the evidence provides more support for the discount rate channel hypothesis $\lambda_S^P > \lambda_R^P$ in the right tail of the distribution of the professionals' return forecast.

Intuitively, it may well be that the professional's model is ill-equipped to process a valuation ratio that is unreasonably high or low relative to macro fundamentals. We leave for future work the interesting question of how professionals may assess (the resolution of) such apparent mispricing and what other signals may feed into their most negative forecasts.

(A) Coefficient on macro-shock component

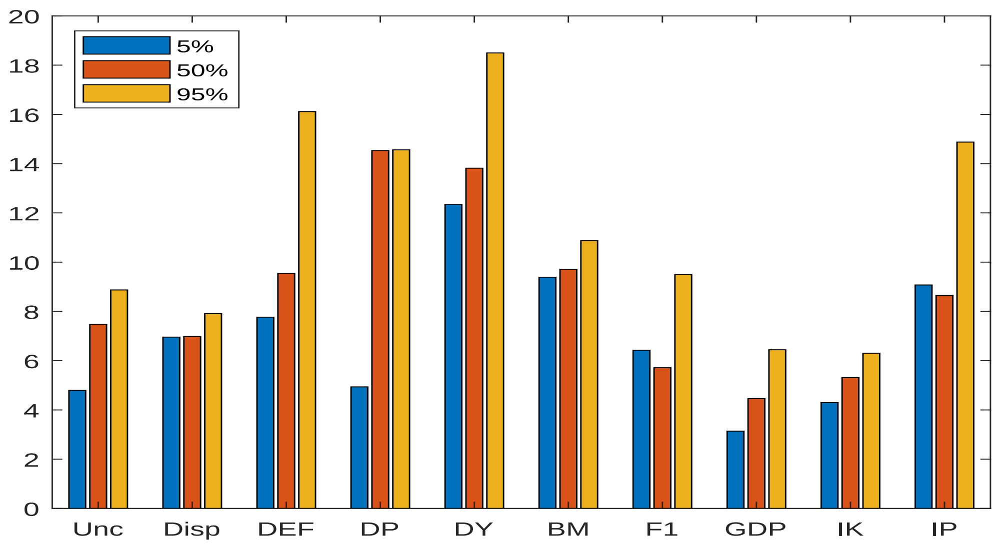

(B) Coefficient on residual component

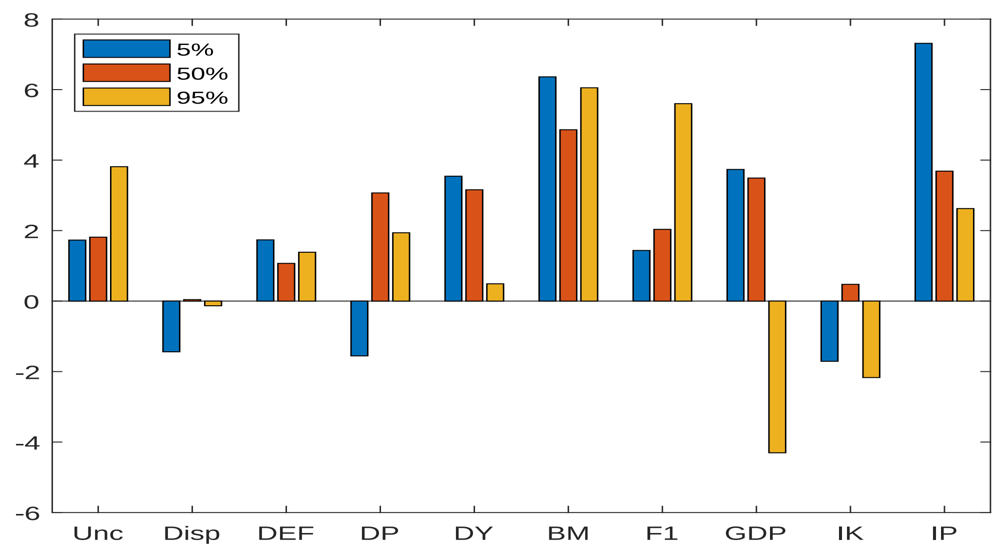

FIGURE B.1: **Second-Stage Coefficient Estimates from Quantile Regressions**
This figure reports point estimates for the coefficients λ_S^P (Panel A) and λ_R^P (Panel B) in Eq. 18. These coefficients are estimated using quantile regression for three quantiles $\tau = 0.05, 0.50, 0.95$.

Appendix C. Long-term return forecasts

The Survey of Professional Forecasters (SPF) reports professional forecasts of the annual-average rate of return to equities (S&P 500) over the next 10 years (including the survey quarter). This forecast is included in all first-quarter surveys from 1992 onwards. To make these forecasts comparable to the one-year excess return forecasts from the Livingston survey studied in our paper, we subtract the 10-year constant maturity Treasury yield and then take the average across forecasters.

To see how these long-term return forecasts respond to shocks, we run the following regression:

$$F^P_{t+1/4}(R_{t:t+10}) = \alpha + \beta N_t + \gamma F^P_{t-\frac{3}{4}}(R_{t-1:t+9}) + \epsilon, \tag{19}$$

where $F^P_{t+1/4}(R_{t+1:t+10})$ is the long-term excess return forecast made in the first quarter of each year $t+1$ (for the years $t+1$ to $t+10$ and $F^P_{t-\frac{3}{4}}(R_{t:t+9})$ is its lag) and N_t is the sum of four quarterly shocks that hit the economy throughout year t. As a benchmark, we run a similar regression for the one-year excess return forecasts studied elsewhere in our paper:

$$F^P_t(R_{t+1}) = \alpha + \beta N_t + \gamma F^P_{t-1}(R_t) + \epsilon, \tag{20}$$

where $F^P_t(R_{t+1})$ is the forecast made in the last quarter of each year t for the return in the next year $t+1$ (i.e., the first year covered by the professionals' 10-year return forecast). Since the response to the Oil shock in the data is delayed, we lag the Oil shock by an additional year for this exercise. The annual shocks N_t are normalized to have unit standard deviation.

We report coefficient estimates in Table C.1 below. We find that the evidence for the long-term forecast in the SPF is surprisingly consistent with the short-term forecast in the Livingston survey, especially noting the many caveats. These caveats include: (i) the relatively short annual sample of at most 29 observations from 1992 to 2020; (ii) we are comparing professionals' forecasts from two very different survey questions (e.g., there is more uncertainty in the long-term, such that one would expect those forecasts to rely more on shrinkage); and, (iii) the long-term forecasts are in excess of a 10-year bond yield that compensates for more risks than the 1-year yield used to calculate the excess 1-year return forecast. To be precise, both the professional 1- and 10-year return forecasts respond positively to each type of shock. For both Cycle and TFP shocks, the impact on the 10-year forecast is significant at 37 and 32 basis points, respectively. For the Oil shock, the impact is smaller and insignificant at 10 basis points. The impacts on the 1-year forecast are larger ranging from 1.69% (Oil) to 4.64% (Cycle).

A simple back-of-the-envelop calculation based on the persistence in the professionals' 1-year return forecasts shows that these estimates are not only qualitatively, but also quantitatively quite consistent. The annual autocorrelation of the professionals' consensus 1-year return forecast from the Livingston survey is about 0.3. Thus, we can approximate the average of the impact of each shock on the next 10 years of returns as: $\frac{1}{10}\sum_{k=1}^{10}$ short-term impact $\times\, 0.3^{k-1} \approx$ short-term impact/7. In other words, the persistence of the professional 1-year excess return forecast in the Livingston survey implies that the increase in the 10-year average return forecast should be a factor of seven smaller than the increase in the 1-year return forecast. This number is close to what we find for the SPF survey that asks professionals the difficult question to forecast these long-term returns, especially for the Cycle and TFP shocks. The slightly worse fit for the Oil shock is perhaps unsurprising noting that the delayed impact of Oil shocks complicates comparison of the coefficients from the two regressions. In fact, we find in Section 4 that Cycle and TFP shocks are relatively stronger instruments for expected returns.

In conclusion, professionals' long- and short-term forecasts both reflect that expected returns are persistently high after bad macro-shocks. We study the persistence of expected return (forecasts) in more detail in Section 5, where we impose the present-value restriction.

TABLE C.1: **Professionals' Long- versus Short-Term Forecasts**

This table presents coefficient estimates from the regressions in Equations (19) and (20). Standard errors are Newey-West with two annual lags.

	β	t-stat	γ	t-stat
Ten-year forecast (Equation (19))				
Cycle	-0.37	(-2.70)	0.15	(0.70)
TFP	-0.32	(-2.48)	0.18	(0.96)
Oil	-0.10	(-0.75)	0.19	(0.74)
One-year forecast (Equation (20))				
Cycle	-4.64	(-2.80)	0.15	(0.52)
TFP	-2.69	(-1.55)	0.48	(2.94)
Oil	-1.69	(-1.27)	0.61	(5.83)

Appendix D. Professional forecasts of GDP

Figure 2 plots the IRF for realized GDP and forecasts of GDP from the SPF (quarterly) and Livingston (bi-annual) surveys to all three shocks. These responses are estimated through the following equations that follow the timing convention used in Bianchi, Ludvigson, and Ma (2022):

$$Y_{t+K+1} = \alpha_K + \beta_K N_t + \sum_l \gamma_l C_{t-l} + \epsilon_{t+K+1}, \tag{21}$$

where Y_{t+K+1} is either the forecast of yearly GDP growth made at $t+K$ ($F_{t+K}(GDP_{t+K+1})$) or the four quarter growth in real per capita GDP from $t+K$ to $t+K+1$ (GDP_{t+K+1}). The vector C_{t-l} collects the controls and when $Y_{t+K} = F_{t+K}(GDP_{t+K+1})$ contains one year worth of lags of the forecast (e.g., $F_{t-1/2}(GDP_{t+1/2})$ and $F_{t-1}(GDP_t)$ for the Livingston case), realized GDP growth (from $t-\frac{5}{4}$ to $t-\frac{1}{4}$) and the shock. When $Y_{t+K} = GDP_{t+K}$, z_{t-l} contains one year worth of lags of realized GDP growth (from $t-\frac{5}{4}$ to $t-\frac{1}{4}$) and the shock. Dictated by the frequency of the forecasts, we use a quarterly frequency for SPF forecasts and realized GDP (i.e., $K = 0, \frac{1}{4}, \frac{1}{2}, ..., 4$ and $l = \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1$) and a biannual frequency for Livingston forecasts ($K = 0, \frac{1}{2}, 1, ..., 4$ and $l = \frac{1}{2}, 1$). To make the IRFs comparable between the quarterly and biannual cases, we multiply the coefficients β_K in both cases by the standard deviation of the biannual shock.

When plotting the response of realized GDP, we calculate the impulse response function also for $K = -1, -3/4, -1/2, -1/4$. As a consequence, the plot for realized GDP has four additional points on the x-axis relative to the impulse responses we report for forecasts of GDP (as well as forecasts of returns in Figure 3). These points show the response of realized GDP growth over annual periods that include the quarter the shock hits, but for which forecasters were not aware of the shock yet.

As highlighted in Bianchi, Ludvigson, and Ma (2022), the responses of professional GDP forecasts are best described as showing strong initial underreaction, with also some evidence of delayed overreaction (see further Angeletos, Huo, and Sastry, 2021). Both the Cycle and TFP shock have a large and immediate negative impact on realized GDP. For the Cycle shock, forecasts hardly respond at all. For the TFP shock, the trough response of forecasts is only about a third of realized GDP. Realized GDP responds with a delay of about 1 to 2 years to the Oil shock. At these horizons, the trough response of forecasts is only about a fourth of realized GDP. For Cycle and TFP shocks, we further see evidence of delayed overreaction, because the response of GDP forecasts turns positive and more so than realized GDP one (Cycle) and two (TFP) years out.

Appendix E. State-space systems

To estimate the parameters of the present-value system in Eqs. (13)-(17), we set up a Kalman filter with two state equations (defining $\widehat{g_t} = g_t - \gamma_0$ and $\widehat{\mu_t} = \mu_t - \delta_0$):

$$X_{t+1} = FX_t + \Gamma\epsilon_{t+1} \text{ where} \tag{22}$$

$$X_t = [\widehat{g_{t-1}}, \widehat{\mu_{t-1}}, \epsilon_t]', \; F = \begin{pmatrix} \gamma_1 & 0 & 0 & 1 & 0 \\ 0 & \delta_1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \text{ and } \Gamma = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{23}$$

$$\tag{24}$$

We also have two observation equations:

$$Y_t = M_0 + M_2 X_t \text{ where} \tag{25}$$

$$Y_t = [\Delta d_t, pd_{t-1}]', M_0 = [\gamma_0, A]' \text{ and } M_2 = \begin{pmatrix} 1 & 0 & 1 & 0 & 0 \\ B_2(\gamma_1 - \delta_1) & -B_1 & 0 & 0 & 0 \end{pmatrix}. \tag{26}$$

There are nine parameters that we estimate using the annual (end-of-year) repurchase-adjusted dividend growth and dividend-to-price series from Nagel and Xu (2022) for the sample period from 1952 to 2020:

Parameter	δ_0	δ_1	γ_0	γ_1	σ_d	σ_g	σ_μ	$\rho_{d\mu}$	$\rho_{\mu g}$
Estimate	0.100	0.755	0.064	0.121	0.024	0.131	0.048	0.748	0.456

As noted also in Koijen and Van Nieuwerburgh (2011), the parameter estimates resulting from a system using repurchase-adjusted series are somewhat different to those estimated in Van Binsbergen and Koijen (2010). That said, our main interest is in the responses of filtered expected returns and dividend growth to macro shocks, and those do not change qualitatively when we use the unadjusted series.

To filter the professional's forecast of dividend growth, we adapt the formulation of the state equations as follows:

$$X_{t+1}^P = F^P X_t^P + \Gamma^P \epsilon_{t+1}^P \tag{27}$$

$$X_t^P = [\widehat{g_t^P}, \widehat{\mu_t^P}]', \; \epsilon_{t+1}^P = [\epsilon_{g,t+1}^P, \epsilon_{\mu,t+1}^P], \; F^P = \begin{pmatrix} \gamma_1^P & 0 \\ 0 & \delta_1^P \end{pmatrix} \text{ and } \Gamma^P = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \tag{28}$$

Recall from footnote 15 that we do not observe the professional's forecast of cum dividend returns $(P_{t+1} + D_{t+1})/P_t$, because the forecast objective in the survey is the capital gain, P_{t+1}/P_t. To account for this difference, the observation equations become:

$$Y_t^P = M_0^P + M_2^P X_t^P, \text{ where} \tag{29}$$

$$Y_t^P = [\rho E_t^P(ln(P_{t+1}/P_t)) + \kappa - (1-\rho)pd_t, pd_t]', \tag{30}$$

$$M_0^P = [(\delta_0^P - (1-\rho)\gamma_0^P), A^P]' \text{ and } M_2^P = \begin{pmatrix} -(1-\rho) & 1 \\ B_2^P & -B_1^P \end{pmatrix}. \tag{31}$$

Here, $E_t^P(ln(P_{t+1}/P_t))$ denotes the mean forecast of the log capital gain.[29] Because ρ is close to one, the adjustment to the capital gain forecast in Eq. (30) turns out to be small in the data.

Ultimately, this system is easier to estimate than the one in Van Binsbergen and Koijen (2010), because there are two parameters less and there is no disturbance in either of the

[29]To arrive at this definition, start from the log-linear approximation of cum-dividend returns, take conditional expectations, use that $\mu_t = E_t(r_{t+1})$ and $g_t = E_t(ln(D_{t+1}/D_t))$, and rearrange to have observables

observation equations. Using semi-annual observations of the capital gain forecasts and the repurchase-adjusted dividend-to-price series from Nagel and Xu (2022) for the sample from June 1952 to December 2019, we estimate the following parameters:

Parameter	δ_0	δ_1	γ_0	γ_1	σ_g	σ_μ	$\rho_{\mu g}$
Estimate	0.08	0.71	0.040	0.87	0.023	0.04	0.56

One important difference with the results above is that the professional forecast of returns is much less persistent than the objective forecast (0.71^2 versus 0.755). Since the two sets of return and growth forecasts need to align with the observed price-to-dividend ratio through the present value relation, it comes as no surprise that the professional's implied dividend growth forecast is much more persistent than the objective forecast (0.87^2 versus 0.121). Combining the parameter estimates, we therefore also find that it is subjectively expected dividend growth that contributes most of the unconditional variation in the price-dividend ratio, even though it is subjectively expected returns that respond most strongly to shocks. These unconditional variance contributions are equal to 82% and 18%, respectively. These results can coexist because a large share of the variation in price-dividend ratios is endogenous or predictable and thus subtracted out through the VAR. Our impulse responses instead focus on the impact of shocks on the unpredictable components of expected returns and dividend growth.

on the left-hand side:

$$\begin{aligned}
r_{t+1} &= \kappa + \rho p_{t+1} + (1-\rho)d_{t+1} - p_t \\
E_t(r_{t+1}) &= \kappa + \rho E_t(ln(P_{t+1}/P_t) + p_t) + (1-\rho)E_t(ln(D_{t+1}/D_t) + d_t) - p_t \\
\mu_t &= \kappa + \rho E_t(ln(P_{t+1}/P_t) + p_t) + (1-\rho)(g_t + d_t) - p_t \\
\rho E_t(ln(P_{t+1}/P_t)) + \kappa + (1-\rho)(d_t - p_t) &= \mu_t - (1-\rho)g_t.
\end{aligned} \tag{32}$$

The same reasoning applies when taking the professional's expectations.

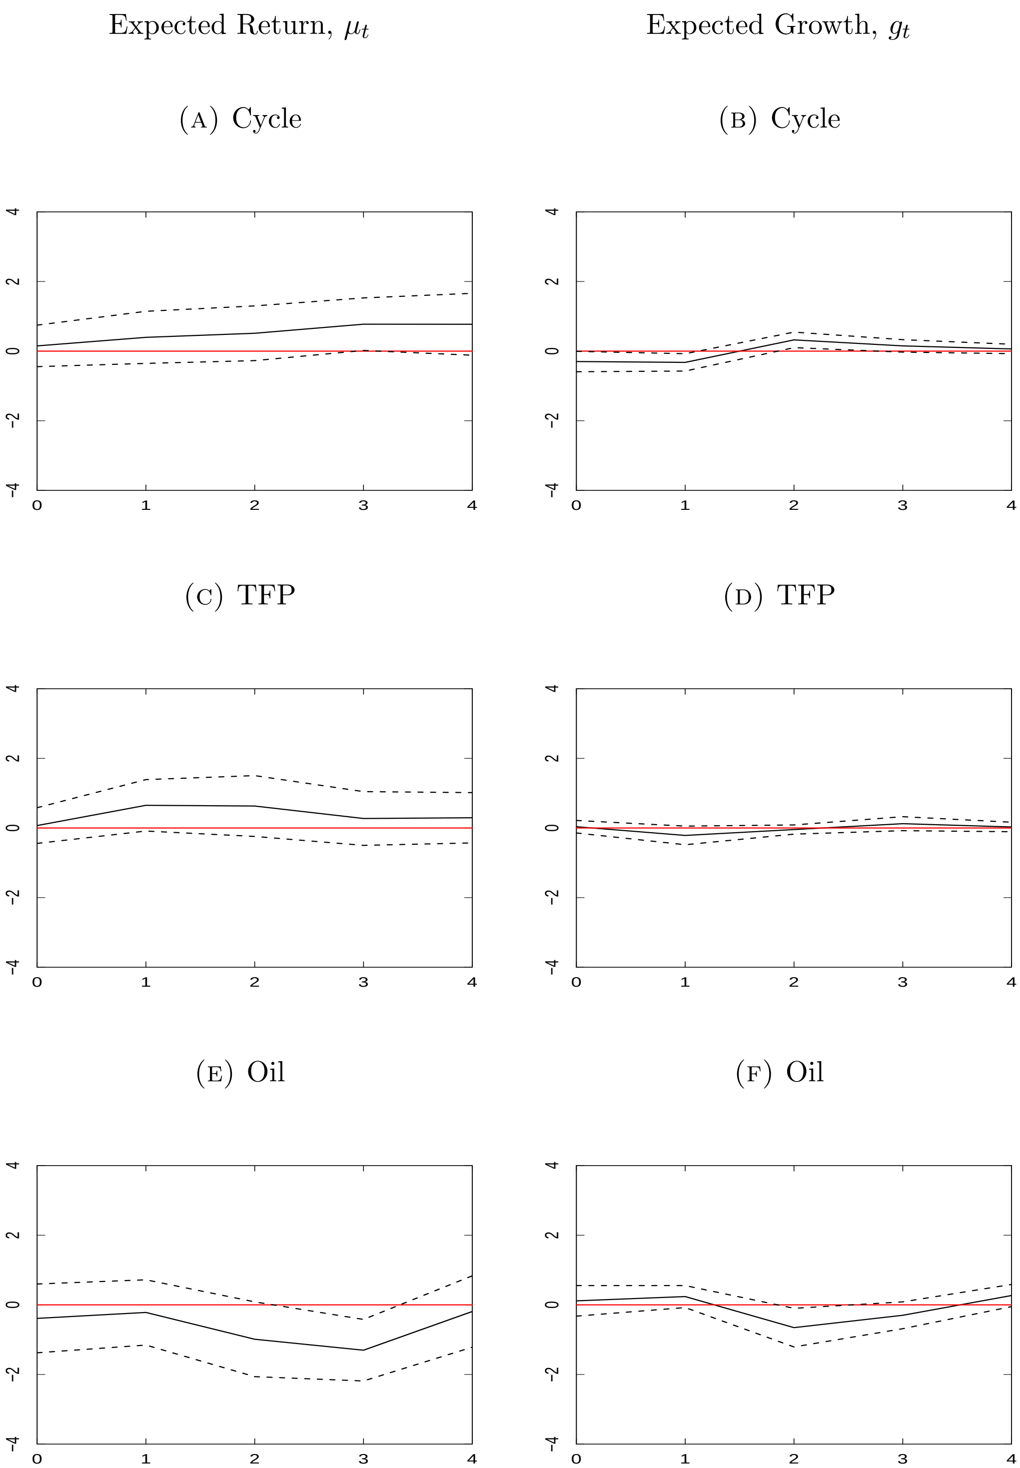

Figure E.1: **Responses of Objective Expectations Derived From the Present-Value Identity**
This figure presents the impulse response function (with 90%-confidence interval constructed using Newey-West standard errors with lag length set to one year) of filtered stock returns and dividend growth forecasts to three different macro shocks (Cycle, TFP and Oil). The filtered series are obtained as in Van Binsbergen and Koijen (2010), see Section 5.1. The responses are estimated using the local linear projections of Eq. (4) with annual data. The projections for stock returns (dividend growth) control for one year worth of lags of stock returns (dividend growth), the forecast and the shock.

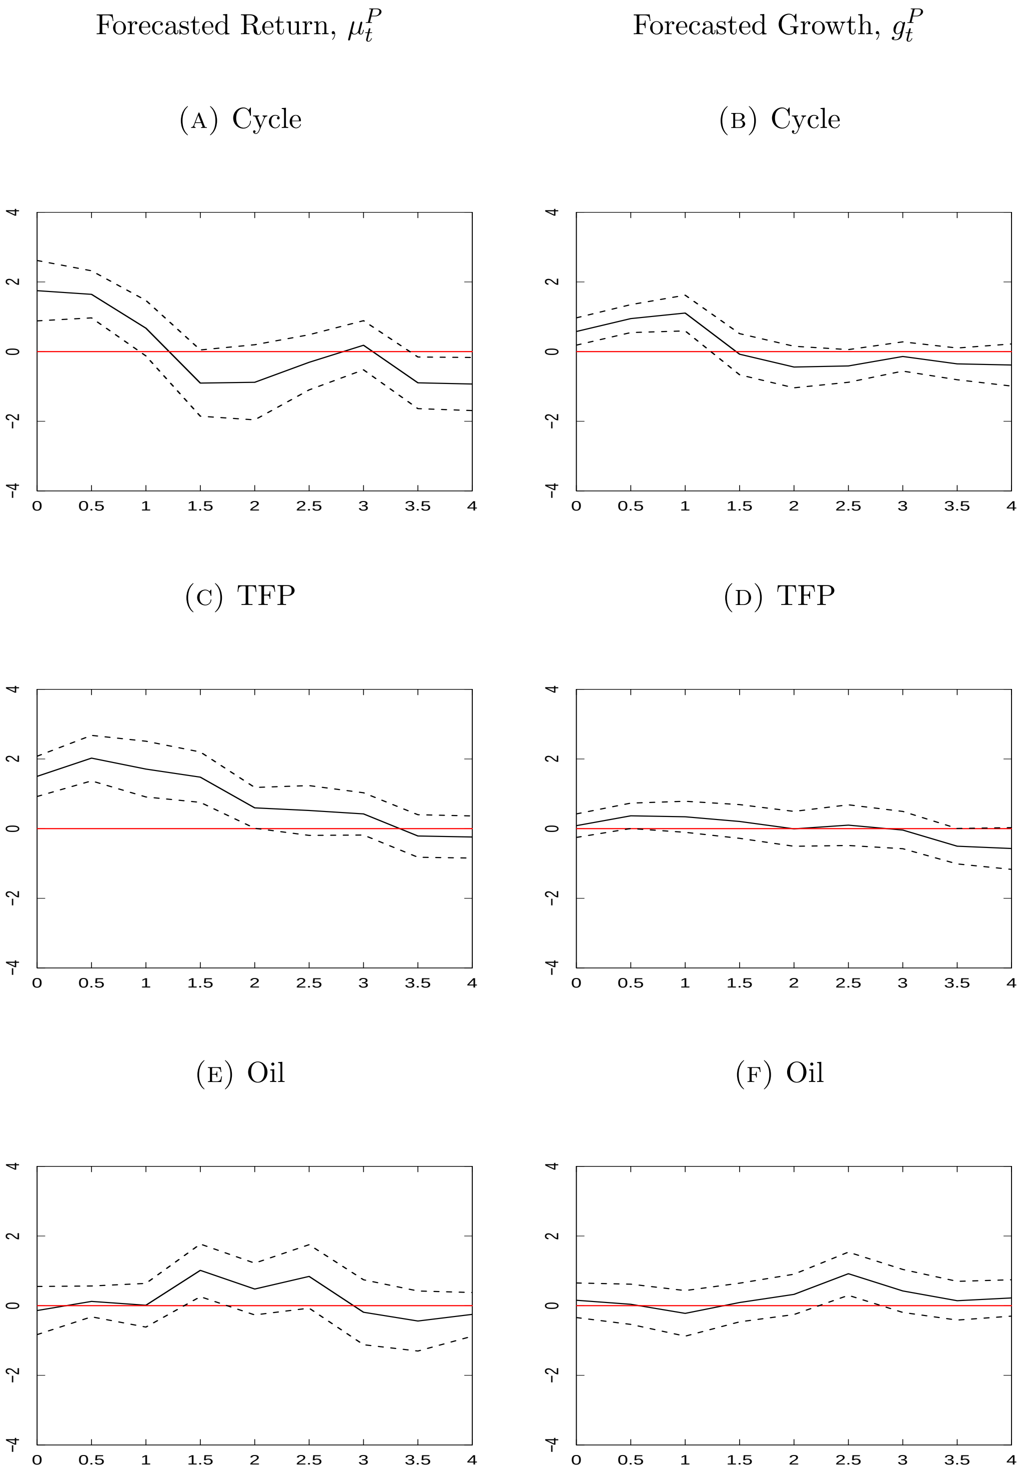

Figure E.2: **Responses of Subjective Expectations Derived From the Present-Value Identity**

This figure presents the impulse response function (with 90%-confidence interval constructed using Newey-West standard errors with lag length set to one year) of subjective expected stock returns and dividend growth to three different macro shocks (Cycle, TFP and Oil). The subjective forecasts are filtered through the present value model described in Section 5.2. The responses are estimated using the local linear projections of Eq. (4) with biannual data. The projections for stock returns (dividend growth) control for one year worth of lags of stock returns (dividend growth), the forecast and the shock.

Appendix F. Additional figures and tables

(A) Cycle (1955-2017)

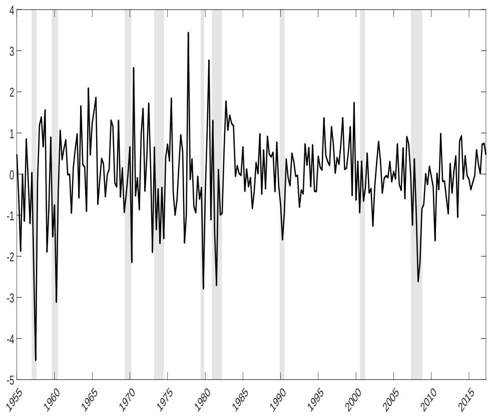

(B) TFP (1962-2019)

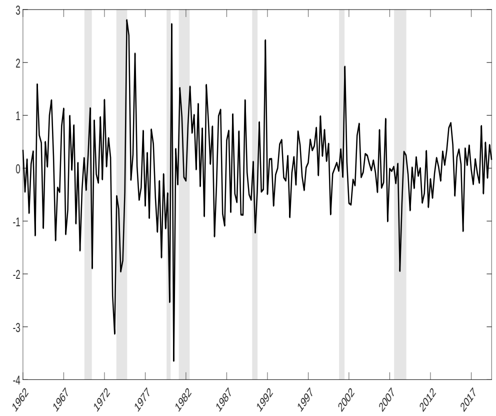

(C) Oil (1975-2019)

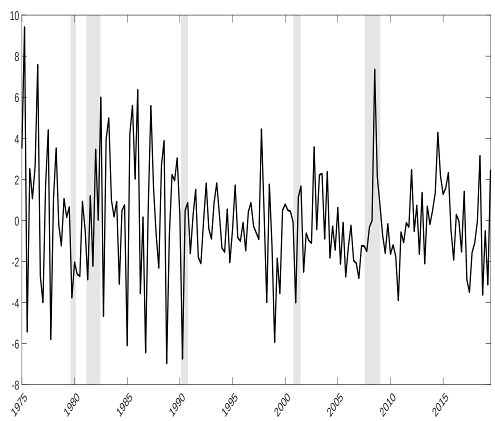

FIGURE F.1: **Macro Shocks**
This figure presents the time series of the three macro shocks that are the main focus in our paper: the business cycle shock of Angeletos, Collard, and Dellas (2020) in Panel A, the TFP news shock constructed using the identification method of Kurmann and Sims (2021) in Panel B and the oil supply shocks from Baumeister and Hamilton (2019) in Panel C. Shaded areas are NBER recessions.

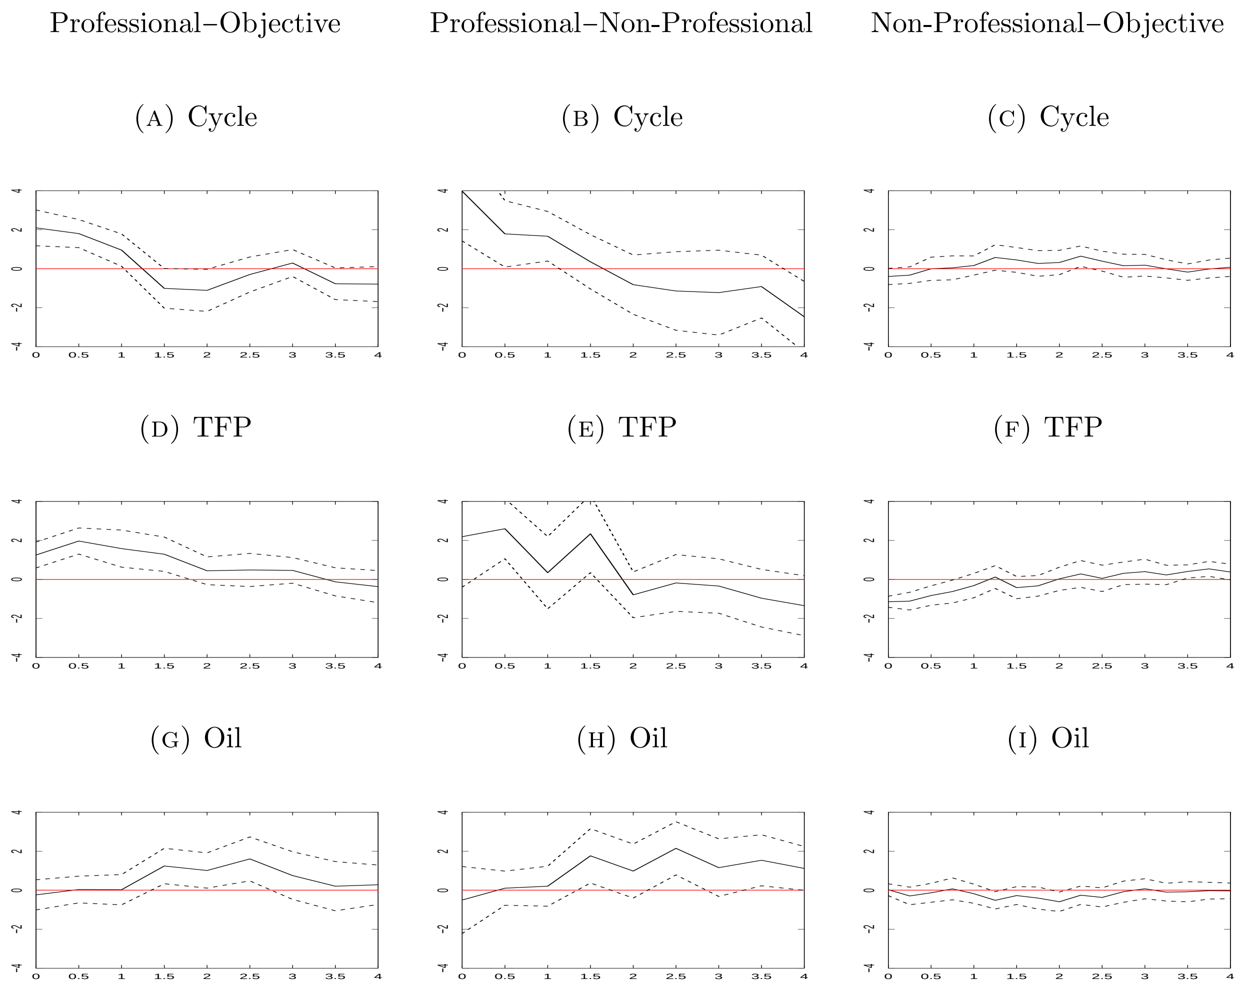

Figure F.2: **Impulse Responses of Differences in Forecasts to Macro Shocks.** This figure is similar to Figure 3 in the paper, but now we present the impulse response function for the difference between two forecasts.

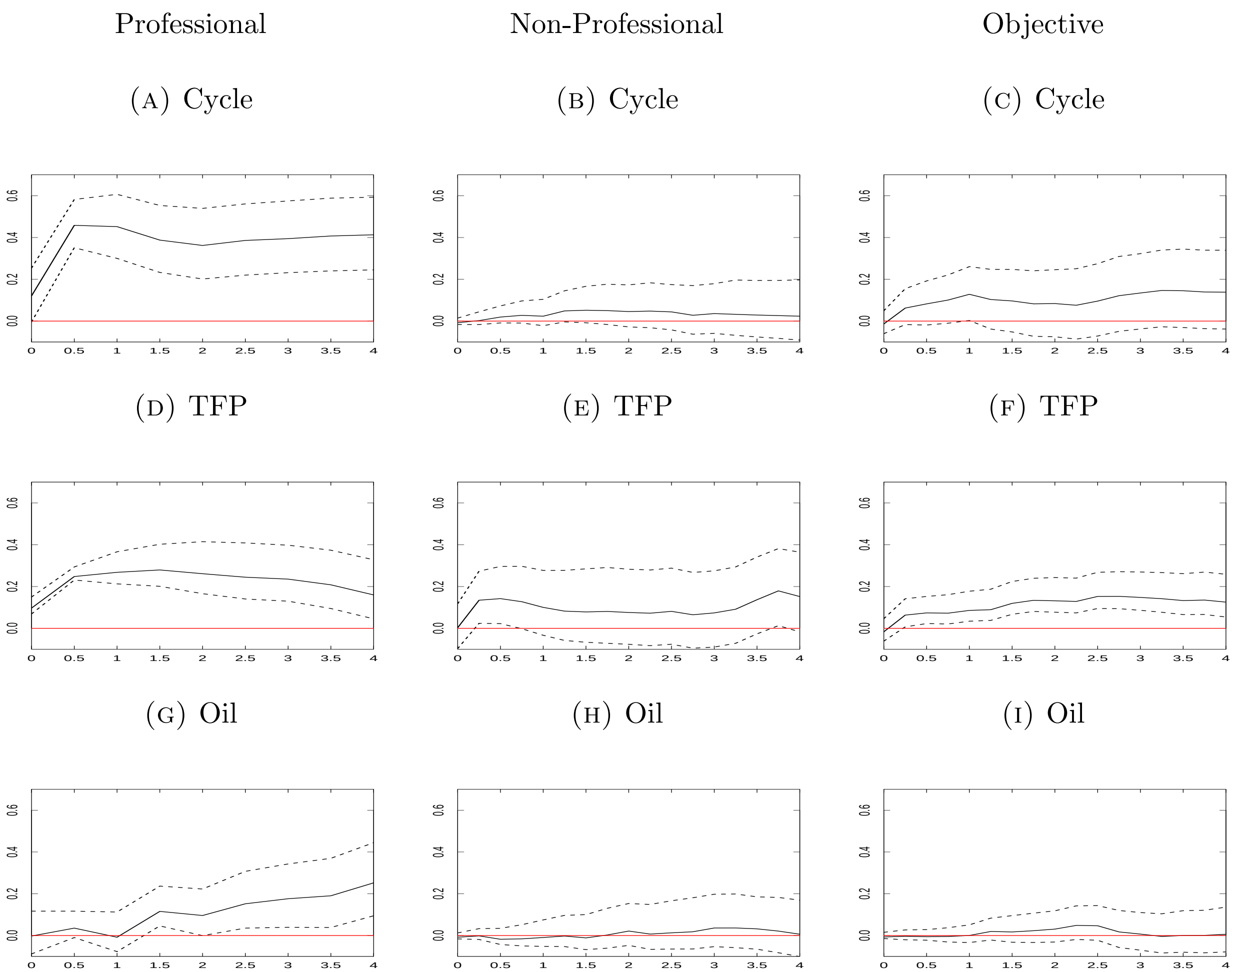

FIGURE F.3: **Forecast Error Variance Decomposition for Forecasts.** This figure is similar to Figure 4 in the paper, but now we report the bias-adjusted FEVD for each forecast time series separately.

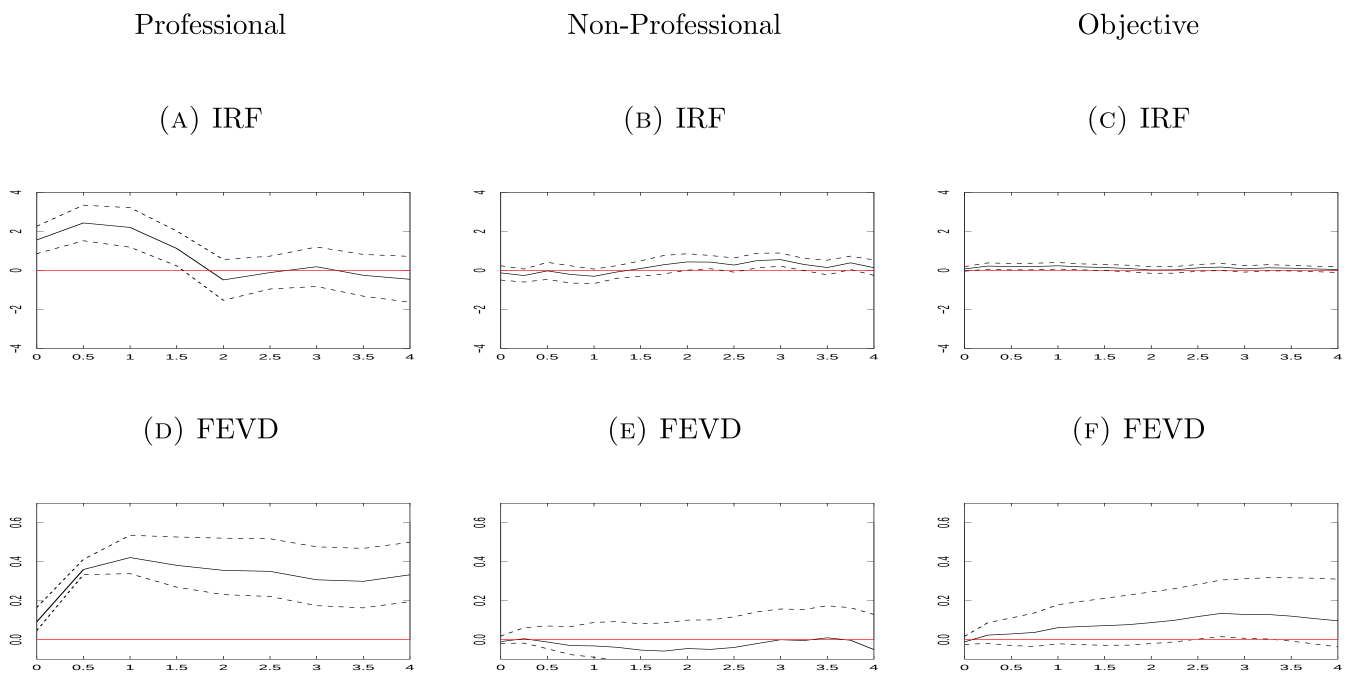

Figure F.4: **IRF and FEVD of Forecasts for IST Shocks.**
This figure presents evidence for an alternative technology shock, that is, the IST news shocks of Ben Zeev and Khan (2015). In the first row, we report the impulse response function for all three forecasts (as in Figure 3). In the second row, we report the forecast error variance decomposition (as in Figure F.3). The sample for IST shocks runs from 1952Q1 to 2012Q4.

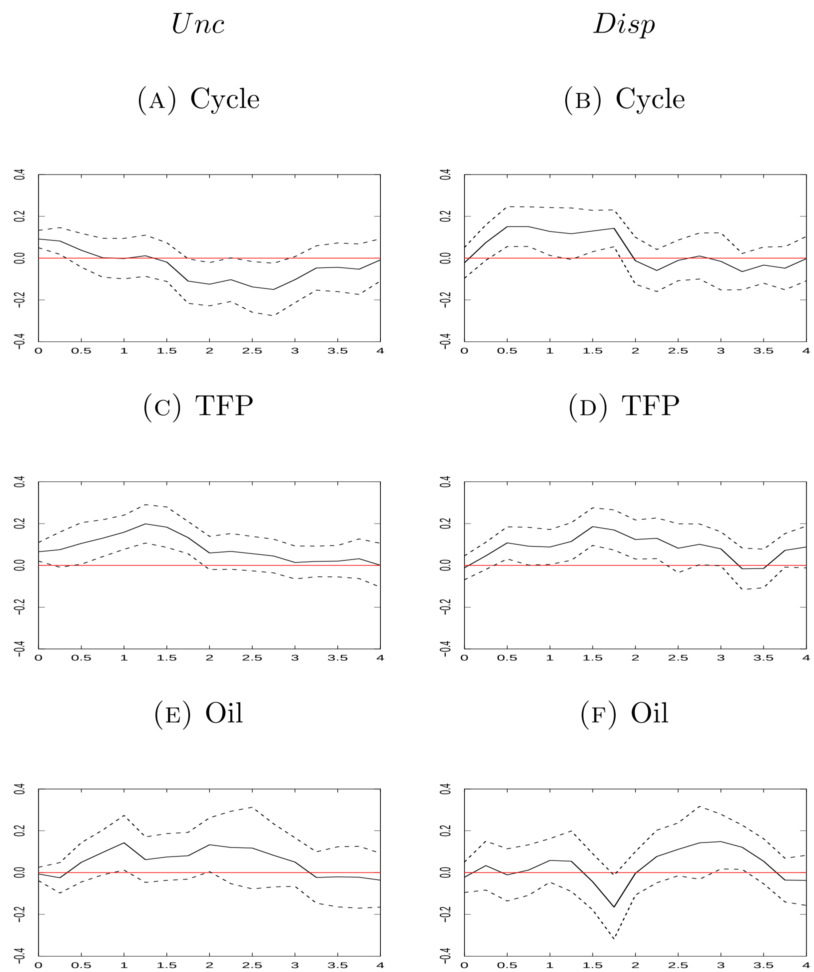

FIGURE F.5: **Impulse Responses of Real Uncertainty to Macro Shocks.** This figure is analogous to Figure 3 in the paper and reports the impulse response function for two measures of real uncertainty: the real uncertainty index of Ludvigson, Ma, and Ng (2021, *Unc*) and the dispersion in Livingston GDP forecasts (*Disp*).

Martin (2017) | Professional Forecast

(A) Cycle | (B) Cycle

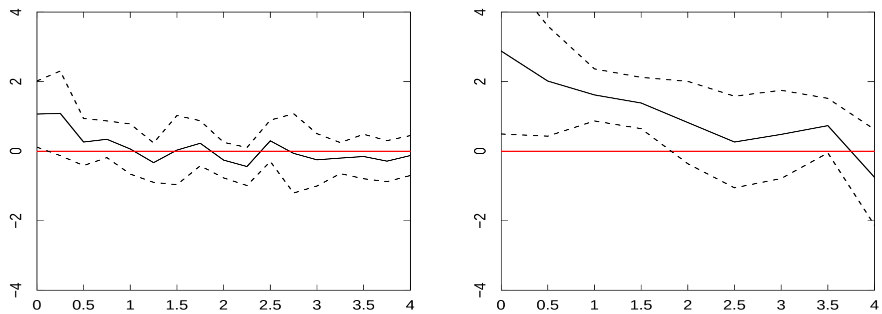

(C) TFP | (D) TFP

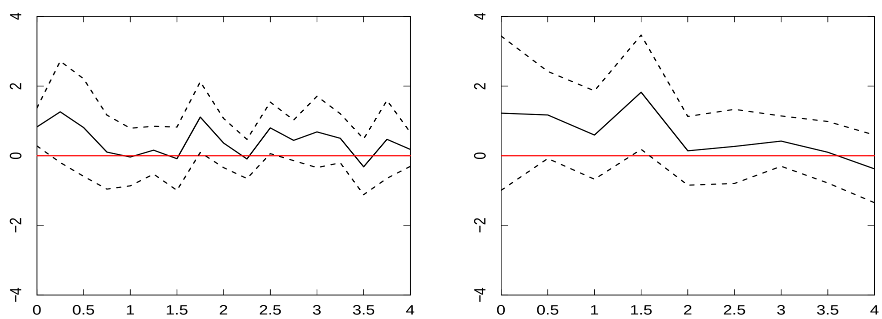

(E) Oil | (F) Oil

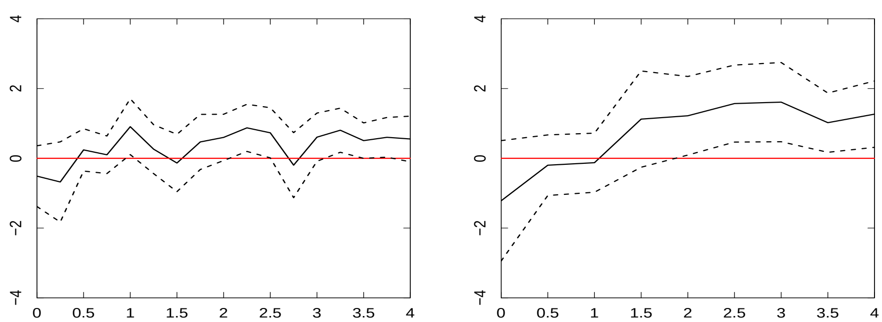

Figure F.6: **Impulse Responses for Martin's 2017 Equity Premium Lower Bound**
This figure is analogous to Figure 3 in the paper and presents the impulse response function of the one-year ahead equity premium lower bound from (Martin, 2017, left) and one-year ahead professional forecast (right) to three macro shocks (Cycle, TFP and Oil).

TABLE F.1: **(Forecasts of) Returns on Real Uncertainty Instrumented Using All Shocks.**
This table is similar to Table III of the paper. The only difference is that we now instrument real uncertainty using all three shock (Cycle, TFP and Oil). Since the Oil shock is only available from 1975 onwards, we pad this shock with zeroes. We check robustness to this approach in Table F.3 below.

Panel A: *Unc*

	R_{t+1}		$F_t^P(R_{t+1})$		$F_t^{NP}(R_{t+1})$		$F_t^O(R_{t+1})$	
Unc	0.71		3.60		-0.03		0.34	
	(0.32)		(2.92)		(-0.09)		(3.46)	
$\widehat{Unc}^{Shock}$		7.74		6.89		-0.54		0.62
		(1.46)		(2.62)		(-0.52)		(2.43)
$\widehat{Unc}^{Res.}$		-2.54		2.02		0.17		0.15
		(-0.97)		(1.85)		(0.34)		(1.19)
Forecast Error on $\widehat{Unc}^{Shock}$				0.85		8.28		7.12
				(0.33)		(2.86)		(2.73)
Adj. R^2	-0.01	0.03	0.23	0.30	0.02	0.02	0.27	0.30
# years	58	58	58	58	33	33	58	58

Panel B: *Disp*

	R_{t+1}		$F_t^P(R_{t+1})$		$F_t^{NP}(R_{t+1})$		$F_t^O(R_{t+1})$	
$Disp$	2.09		0.76		-0.02		0.36	
	(1.29)		(1.38)		(-0.07)		(3.29)	
$\widehat{Disp}_t^{Shock}$		6.62		5.65		0.43		0.76
		(1.29)		(1.96)		(0.67)		(2.95)
$\widehat{Disp}_t^{Res.}$		1.22		-0.14		-0.08		0.28
		(0.76)		(0.08)		(-0.32)		(2.85)
Forecast Error on $\widehat{Disp}_t^{Shock}$				1.31		5.16		5.03
				(0.66)		(2.73)		(2.80)
R^2	0.00	0.01	0.04	0.14	-0.02	-0.03	0.18	0.21
# years	58	58	58	58	33	33	58	58

TABLE F.2: **Controlling for Lagged Professional Forecast.**
Analogous to Tables II and III, we presents results from our 2SLS estimation when we control for lags of the professional forecast (in addition to lagged real uncertainty and lagged realized returns) in the first-stage regression of two alternative measures of real uncertainty on Cycle and TFP shocks (see Eq. (8), Panel A) as well as in the reduced-form regression of professional forecasts on real uncertainty (see Eq. (9), first and third columns in Panel B) and the second-stage regression of professional forecasts on instrumented real uncertainty (see Eq. (11), second and fourth column in Panel B). The lag of the professional forecast is chosen so that this forecast is made one quarter before any of the shocks included in the test hit the economy.

Panel A: First Stage

	Unc	*Disp*
Cycle	0.19	0.24
	(2.25)	(2.54)
TFP	0.32	0.25
	(4.35)	(2.66)
Adj.R^2	0.42	0.15
# years	55	55
κ^*	3/4	5/4
Eff. F-stat.	18.24***	10.605***

Panel B: Second Stage

	Unc		*Disp*	
X	3.85		0.66	
	(3.04)		(1.14)	
$\widehat{X^{Shock}}$		6.88		6.09
		(2.75)		(2.32)
$\widehat{X^{Res.}}$		2.32		-0.43
		(1.93)		(-0,68)
R^2	0.26	0.33	0.04	0.18
# years	54	54	55	55

Table F.3: **First Stage Coefficients Restricting the Sample to Oil Shocks.** In Panel A of this table, we report the first stage regression of Eq. (8) of the real uncertainty index of Ludvigson, Ma, and Ng (2021) on the three shocks. We now restrict the sample to the period after 1975, which is when we the oil supply shock of Baumeister and Hamilton (2019) is available. In Panel B, we re-estimate the second stage regression of Eq. (11) (analogous to what is reported in Panel A of Table III) using these alternative first-stage coefficient estimates. To be precise, we combine these coefficient estimates with the three shocks (as they become available over the period from 1963 to 2017) to construct the instrumented component of real uncertainty.

Panel A: First Stage

	Cycle+TFP+Oil
Cycle	0.19
	(2.13)
TFP	0.32
	(4.16)
Oil	0.05
	(0.68)
Adj.R^2	0.39
# years	42
κ^*	3/4

Panel B: Second Stage

	R_{t+1}		$F_t^P(R_{t+1})$		$F_t^{NP}(R_{t+1})$		$F_t^O(R_{t+1})$	
Unc	0.71		3.60		-0.03		0.34	
	(0.32)		(2.92)		(-0.09)		(3.46)	
$\widehat{Unc}^{Shock}$		6.86		6.67		-0.49		0.61
		(1.55)		(3.13)		(-0.43)		(2.84)
$\widehat{Unc}^{Res.}$		-2.41		1.99		0.06		0.21
		-0.86		(1.72)		(0.11)		(1.79)
Forecast Error on $\widehat{Unc}^{Shock}$				0.19		7.35		6.25
				(0.07)		(2.54)		(2.39)
Adj. R^2	-0.01	0.02	0.24	0.31	0.01	0.01	0.30	0.32
# years	55	55	54	54	32	32	55	55

TABLE F.4: **(Forecasts of) Returns on Macro Shocks (Reduced Form).**
This table presents results from the reduced-form regression of (forecasts of) returns on each of the three macro shocks (Cycle, TFP, and Oil). We control for lagged real uncertainty and lagged realized returns to make these results comparable to Panel A of Table III. We report the coefficient on the shocks, the t-statistics based on Newey-West standard errors with four lags, the number of years in the sample, the R^2, and the optimal lag choice κ, which is used to select the four consecutive quarters over which the shocks are cumulated (see the discussion in Section 4.1).

Panel A: Cycle

	R_{t+1}	$F_t^P(R_{t+1})$	$F_t^{NP}(R_{t+1})$	$F_t^O(R_{t+1})$
Coeff	3.04	2.99	-0.26	0.23
	(1.71)	(5.09)	(-0.78)	(2.72)
Adj.R^2	0.01	0.33	0.02	0.23
# years	58	58	58	58
κ^*	0/4	0/4	0/4	0/4

Panel B: TFP

	R_{t+1}	$F_t^P(R_{t+1})$	$F_t^{NP}(R_{t+1})$	$F_t^O(R_{t+1})$
Coeff	3.14	2.05	-0.17	0.22
	(1.95)	(3.10)	-(0.49)	(2.91)
Adj.R^2	0.02	0.15	0.02	0.27
# years	56	56	56	56
κ^*		3/4	3/4	3/4

Panel C: Oil

	R_{t+1}	$F_t^P(R_{t+1})$	$F_t^{NP}(R_{t+1})$	$F_t^O(R_{t+1})$
Coeff	-1.02	0.60	-0.25	-0.04
	(-0.50)	(1.08)	(-1.38)	(-0.51)
Adj.R^2	-0.01	0.08	0.04	0.3
# years	43	43	43	43
κ^*	3/4	3/4	3/4	3/4

Table F.5: **(Forecasts of) Returns on Shock-Insensitive Predictors.**
This table presents results from regressing (forecasts of) returns on shock-insensitive predictors (along the lines of the first column in each block of two columns in Table V). These coefficient estimates are plotted in Panel A of Figure 7. Shock-insensitive predictors are those for which at least one of the following two criteria is not met. First, the predictor must respond to the Cycle and TFP shock with the same sign. Second, the first stage must be strong (the effective F-statistic in Table IV must be significant at the 10%-level).

	CAY		CSP		DE		DFR	
	R_{t+1}	$F_t^P(R_{t+1})$	R_{t+1}	$F_t^P(R_{t+1})$	R_{t+1}	$F_t^P(R_{t+1})$	R_{t+1}	$F_t^P(R_{t+1})$
X	4.92	-1.06	10.89	1.75	1.21	0.94	0.64	1.10
	(2.47)	-(1.04)	(3.60)	(1.41)	(0.94)	(1.27)	(0.66)	(1.32)
Adj. R^2	0.04	0.06	0.09	0.09	-0.01	0.10	-0.01	0.09
# years	55	55	40	39	55	55	57	57

	EP		EXPD		INFL	
	R_{t+1}	$F_t^P(R_{t+1})$	R_{t+1}	$F_t^P(R_{t+1})$	R_{t+1}	$F_t^P(R_{t+1})$
X	2.27	-0.22	2.70	1.00	2.14	-2.10
	(1.15)	(0.15)	(1.44)	(1.89)	(0.84)	(-1.53)
Adj. R^2	0.03	0.07	0.12	0.15	-0.00	0.07
# years	56	56	56	55	57	56

	LTR		LTY		NTIS		SVAR	
	R_{t+1}	$F_t^P(R_{t+1})$	R_{t+1}	$F_t^P(R_{t+1})$	R_{t+1}	$F_t^P(R_{t+1})$	R_{t+1}	$F_t^P(R_{t+1})$
X	0.63	0.25	2.10	0.64	1.53	1.23	2.12	5.73
	(0.49)	(0.49)	(0.59)	(0.42)	(0.54)	(1.23)	(0.81)	(6.48)
Adj. R^2	-0.01	0.02	0.03	0.10	-0.00	0.03	-0.01	0.28
# years	57	56	55	54	56	56	57	57

	TBL		TERM		VIX2	
	R_{t+1}	$F_t^P(R_{t+1})$	R_{t+1}	$F_t^P(R_{t+1})$	R_{t+1}	$F_t^P(R_{t+1})$
X	0.83	-1.08	1.43	0.77	0.77	1.57
	(0.25)	(-0.55)	(0.71)	(0.82)	(0.63)	(2.47)
Adj. R^2	-0.00	0.15	0.01	0.14	-0.01	0.23
# years	56	55	56	55	56	55