The Stock Market's Two Truths: Subjective Beliefs and Objective Reality

Aizhan Anarkulova* Scott Cederburg† Yi Zhou‡

March 14, 2025§

Abstract

We demonstrate that *objective* stock market return decompositions inform about the representative investor's *subjective* beliefs. Our novel method for objective return decomposition uses model forecast combinations as state variables in the vector autoregression (VAR) and exhibits remarkable stability compared with the conventional VAR approach. In the context of our theoretical framework, empirical patterns in discount-rate and cash-flow news imply the representative investor underreacts to fundamental news and has noisy subjective cash-flow expectations. These effects intensify in pre-recession periods, as prices rise due to cash-flow optimism despite objectively deteriorating fundamentals, indicating the wedge between subjective beliefs and objective reality significantly contributes to market crashes.

JEL classifications: C32, C53, D84, G10, G12, G14, G17, G40

Key words: Return decomposition, objective and subjective beliefs, model forecast combination, vector autoregression

*Goizueta Business School, Emory University. Email: aizhan.anarkulova@emory.edu.

†Eller College of Management, University of Arizona. Email: cederburg@arizona.edu.

‡Eller College of Management, University of Arizona. Email: yizh@arizona.edu.

§We thank Michael O'Doherty, Yirong Wang, and seminar participants at the 2024 Financial Management Association for helpful comments and suggestions.

1 Introduction

What moves stock prices? To answer this fundamental question, seminal work by Campbell and Shiller (1988a, 1988b) and Campbell (1991) decomposes stock returns into changes in expectations about future cash flows and discount rates. Under the full-information rational expectations (FIRE) framework, these objective return components align perfectly with the representative investor's subjective beliefs about future returns and cash flows. Survey evidence, however, reveals systematic deviations between investor beliefs and rational forecasts [e.g., Vissing-Jorgensen (2003), Amromin and Sharpe (2013), and Greenwood and Shleifer (2014)]. This documented wedge breaks the equivalence between objective and subjective return decompositions and creates intricate feedback mechanisms between the two. Specifically, return realizations do not simply depend on the realized cash flows and fundamental shocks, but on how the representative investor perceives and reacts to them. This investor's subjective beliefs and risk preferences determine asset prices and returns, which in turn determine objective expected returns given the true fundamentals. The central challenge in studying these phenomena is that objective and subjective expectations are not directly observable and are difficult to precisely measure.

In this paper, we address the empirical challenge of estimating *objective* return and cash-flow expectations and propose a theoretical framework for using these estimates to characterize *subjective* beliefs. Our empirical methodology implements Campbell's (1991) vector autoregression (VAR) approach to estimate objective discount-rate and cash-flow news, but overcomes its sensitivity to the choice of state variables [e.g., Chen and Zhao (2009)]. Rather than relying on a few individual return predictors, we distill information from many predictors into three persistence-based model forecast combinations and include them as VAR state variables. Our theoretical framework builds on the fact that realized stock market returns must simultaneously reconcile with two truths: (i) the representative investor's subjective discount-rate and cash-flow news components and (ii) the objective discount-rate and cash-flow news counterparts. Thus, subjective and objective return decompositions are inherently linked. Since the representative investor's subjective beliefs determine prices and returns, objective discount-rate news implicitly reflects these subjective views given objective cash-flow news. Our results reveal that a substantial share of return volatility stems from systematic belief distortions: namely, underreaction to fundamental shocks and noisy subjective cash-flow beliefs. The subjective-objective wedge is pronounced in pre-recession periods, and subjective cash-flow optimism proves to be central to stock market crashes as recessions begin.

Our approach to studying subjective beliefs requires reliable estimates of objective news com-

ponents. Our novel VAR with forecast combinations delivers these, as opposed to the conventional VAR. We empirically demonstrate the magnitude of the conventional VAR's problem in Figure 1 by examining the time-series correlation of discount-rate news estimates across different model specifications. Using monthly data from January 1927 to December 2021, we estimate the 2,300 possible VARs that each include the log market excess return and three market return predictors from a set of 25 widely used predictors. The results are striking. Among all 2,643,850 model pairs, only 67% yield positively correlated discount-rate news estimates, a bare minimum requirement for methodological stability. Even more concerning, nearly half (49%) of correlations fall between -0.2 and 0.2. Two researchers making independent choices on state variables with a conventional VAR are unlikely to estimate discount-rate news series that are statistically or economically similar, underscoring the need for our robust methodology.

Our return decomposition method is simple to implement. We first estimate univariate predictive regressions for 25 return predictors. We then divide these variables into three groups based on persistence levels (low, medium, and high) and calculate the mean conditional return forecast within each group. Finally, we use these persistence-based forecast combinations as state variables in the VAR to estimate objective discount-rate and cash-flow news. This approach offers three key methodological advantages: (i) it incorporates information from many predictors, which account for diverse factors including macroeconomic risks, valuations, sentiment, and microstructure effects, while maintaining parsimony; (ii) it benefits from the stability of model forecast combinations; and (iii) it captures return expectations across various time horizons.

These methodological advantages translate into substantial improvements in empirical performance. First, the VAR achieves remarkable predictive power, with an R^2 of 4.93% for monthly excess market returns compared to an average of just 1.76% for standard three-variable VAR specifications. This enhanced predictability stems from our ability to incorporate information from 25 distinct return predictors while avoiding the dimensionality and overfitting problems that plague "kitchen sink" models. Second, our method generates news components that cleanly separate cash-flow and discount-rate effects — past discount-rate news strongly forecasts annual returns with an R^2 of 20.22% (t-statistic of 3.64) but not dividend growth, while past cash-flow news robustly predicts annual dividend growth with an R^2 of 18.13% (t-statistic of 5.01) but not returns. Third, our approach captures rich dynamics across time horizons by modeling expectations with different persistence levels, ranging from monthly to multi-year forecasts. Fourth, our estimates of return components demonstrate exceptional stability to modeling choices, remaining highly correlated even when varying the set of included predictors or persistence cutoffs.

We use our robust methodology and theoretical framework to draw insights about subjective beliefs from the empirical patterns in objective discount-rate and cash-flow news. Two initial findings are informative. First, objective discount-rate news exhibits higher volatility than cash-flow news with monthly standard deviations of 5.90% versus 2.84%, respectively. FIRE models explain this result with highly volatile risk premiums; within our framework, variation in subjective cash-flow optimism or pessimism can drive objective discount-rate volatility. Second, and more striking, the objective discount-rate and cash-flow news components exhibit a positive correlation of 0.46. This estimate is inconsistent with many FIRE models that predict a negative correlation, as poor cash-flow news coincides with increasing discount rates in the models. In our theoretical framework, the sign of this correlation is informative about subjective beliefs, and the positive correlation estimate implies that the representative investor underreacts to objective cash-flow news.

Our analysis of market behavior around recessions reinforces the disconnect between subjective beliefs and objective reality. Examining the 15 NBER-dated recessions in our sample, we find systematic patterns inconsistent with FIRE. In the two years preceding economic peaks, the market experiences positive cumulative excess returns averaging 16%, despite deteriorating objective cash-flow news that anticipates the recession (cumulative effect of -11% on realized returns). This pre-recession price run-up without support from fundamentals causes a substantial decline in objective expectations of future returns (cumulative effect of 20% on realized returns). Once recessions begin, the market experiences statistically significant declines [consistent with Gómez-Cram (2022) and Kroencke (2022)], with cumulative excess returns reaching an average low of -8% in month six of the recession. Objective cash-flow news remains relatively flat once the recession begins, so the price decline implies unexpected increases in objective expected returns (cumulative effect of -12%). These systematic patterns — increasing pre-recession prices despite deteriorating fundamentals and sharply decreasing recession prices when fundamental news is objectively neutral — provide additional evidence that investors sluggishly react to objective cash-flow news.

We also provide direct evidence of the subjective-objective belief gap by comparing the return components from objective and subjective models. For the subjective return decomposition, we use analyst forecasts as a proxy for the representative investor's subjective expectations following Chen, Da, and Zhao (2013). The correlation between subjective and objective cash-flow news is remarkably low at just 0.21. Subjective cash-flow news also displays much weaker predictability for future dividend growth compared with objective cash-flow news. These results further support our conclusion that subjective cash-flow beliefs are noisy and deviate from objective reality.

The objective-subjective gap is particularly pronounced in periods preceding economic down-

turns. While our prior analysis shows deteriorating objective cash-flow news accompanied by decreasing objective discount rates before recessions, analyst forecasts reveal the opposite pattern. The subjective decomposition implies that the price run-ups are driven by strongly positive subjective cash-flow news, with subjective discount-rate news having a slight negative contribution. The decidedly optimistic subjective viewpoint helps explain why investors continue to hold stocks even as objective expected returns turn negative. Only later, as evidence of poor fundamentals and overpriced stocks builds, do investors react.

Our paper contributes to three broad strands of literature: (i) subjective expectations and deviations from FIRE, (ii) return decompositions, and (iii) return predictability.

First, while the rational expectations hypothesis assumes investors form expectations that match objective probability distributions [e.g., Muth (1961), Lucas (1978), and Cochrane (2011)], survey evidence reveals that investors' beliefs deviate from rational expectations [e.g., Vissing-Jorgensen (2003), Amromin and Sharpe (2013), and Greenwood and Shleifer (2014)].[1] Direct examination of subjective beliefs requires survey data, and the need for more and better surveys has been recognized in the literature [e.g., Manski (2004)]. Two major issues arise when focusing on surveys to learn about subjective beliefs: (i) much of the survey data has a short forecast horizon, whereas long-run expectations are likely important for understanding why stock prices move, and (ii) surveys show that subjective expectations vary considerably across different investor types, leaving doubt about which investor group best proxies for the representative investor. Instead of relying solely on survey data, recent studies by Chen, Hansen, and Hansen (2020) and Ghosh and Roussellet (2023) attempt to identify subjective investor beliefs by simultaneously considering surveys and asset prices. We contribute to this literature by characterizing subjective beliefs through the lens of our framework that relates objective and subjective news components. Our primary approach directly examines returns with no reliance on survey data, such that it considers all relevant horizons and studies the representative investor who sets prices.

Second, our paper contributes to the literature on return decompositions. Researchers have used Campbell's (1991) return decomposition approach to study intertemporal risk, monetary policy,

[1] Recent work examines subjective cash flow expectations [e.g., Chen, Da, and Zhao (2013) and De la O and Myers (2021)] or return expectations across investor types such as retail investors [e.g., Greenwood and Shleifer (2014) and Giglio, Maggiori, Stroebel, and Utkus (2021)] or institutional investors [e.g., Andonov and Rauh (2022); Dahlquist and Ibert (2024); and Couts, Gonçalves, and Loudis (2024)]. Other papers study subjective beliefs for other asset classes [e.g., Piazzesi, Salomao, and Schneider (2015) and Buraschi, Piatti, and Whelan (2022) for bonds; Piazzesi and Schneider (2009) for real estate; and Nagel and Xu (2023) for asset classes including stocks, bonds, currencies, and commodities]. See Brunnermeier, Farhi, Koijen, Krishnamurthy, Ludvigson, Lustig, Nagel, and Piazzesi (2021) and Adam and Nagel (2023) for comprehensive reviews.

foreign exchange, and sources of market return predictability, among many other topics.[2] We develop a novel empirical method that substantially improves the estimation of cash-flow news and discount-rate news for stock market returns. Our simple approach addresses the primary criticism of conventional VAR models, which is sensitivity to the choice of state variables [e.g., Chen and Zhao (2009)], and delivers more robust and reliable estimates of the objective discount-rate and cash-flow news.

Third, our work relates to the literature on return predictability and time-varying expected returns.[3] Building on prior work on forecast combinations [e.g., Bates and Granger (1969); Timmermann (2006); and Rapach, Strauss, and Zhou (2010)], we combine forecasts from multiple predictive models to yield superior estimates of expected returns. Relative to existing papers, we innovate by grouping predictors based on their persistence levels, allowing our VAR to capture return dynamics across different time horizons. Our persistence-based combination approach achieves consistently strong predictive performance both in-sample and out-of-sample, which is a difficult feat to achieve [e.g., Goyal and Welch (2008) and Goyal, Welch, and Zafirov (2024)].

The remainder of the paper is organized as follows. Section 2 develops our theoretical framework for understanding subjective expectations through the lens of the objective return decomposition and introduces our stock return decomposition methods. Section 3 describes our data sources and provides descriptive statistics. Section 4 validates our estimation method, analyzes the implications of objective news components for subjective beliefs, and directly compares objective and subjective news components. Section 5 concludes.

2 Return decompositions and methods

We decompose unexpected stock market returns into cash-flow news and discount-rate news components, based on either objective expectations or subjective expectations. Section 2.1 develops these decompositions and explores the implied connections between objective and subjective news components. Section 2.2 introduces our novel method for estimating the objective news components. Section 2.3 describes our approach for estimating the subjective news components following Chen, Da, and Zhao (2013).

[2]See, e.g., Campbell (1996) and Gonçalves (2021) for intertemporal risk; Patelis (1997) and Bernanke and Kuttner (2005) for monetary policy; Bekaert and Hodrick (1992) and Atanasov and Nitschka (2015) for foreign exchange; and Lamont (1998) and Rapach, Ringgenberg, and Zhou (2016) for sources of market return predictability using the return decomposition framework.

[3]See Rapach and Zhou (2013) for a review of the literature on stock return forecasting.

2.1 Objective and subjective return decompositions

Campbell and Shiller (1988a, 1988b) develop a framework relating prices to future cash flows and returns. After loglinearization, the log price-dividend ratio follows

$$p_t - d_t = \frac{k}{1-\rho} + E_t \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} - r_{f,t+1+j} - r^e_{t+1+j}), \tag{1}$$

where Δd_{t+1} is log dividend growth, $r_{f,t+1}$ is the log risk-free rate, r^e_{t+1} is the log excess return, E_t is the objective expectation operator, and ρ and k are loglinearization constants.[4] Equation (1) shows that prices are high when expected future dividend growth (in excess of the risk-free rate) is high and/or when future expected excess returns are low. Under subjective expectations,

$$p_t - d_t = \frac{k}{1-\rho} + E^*_t \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} - r_{f,t+1+j} - r^e_{t+1+j}), \tag{2}$$

where E^*_t is the subjective expectation operator. With FIRE, the representative investor's subjective expectations equal objective expectations, making equations (1) and (2) equivalent.

Without FIRE, investor views may depart from objective reality. The representative investor forms beliefs about future cash flows and discount rates, and she sets a price that reflects her subjective beliefs [i.e., equation (2)]. This price level and the objective reality of future cash flows determine the objective expected excess returns [i.e., equation (1)].[5] Intuitively, objective expectations of future returns are low (high) when the price is too high (low) given the true cash-flow fundamentals.

As time passes, the representative investor observes signals about objective fundamentals that affect her subjective beliefs. Positive cash-flow surprises, for example, may lead her to become more optimistic about future growth. A change in subjective beliefs about cash flows and discount rates will cause the representative investor to set a new price, which drives the current-period market return. We can capture the effects of changes in beliefs within Campbell's (1991) return decomposition framework.

Unexpected excess stock returns can be decomposed into cash-flow news and discount-rate news

[4] We set the loglinearization constant $\rho = 1/(1 + \exp(\overline{d-p}))$ to 0.9972, which is calculated from the average log monthly dividend-price ratio ($\overline{d-p}$) in our sample. The decomposition assumes the terminal condition $\lim_{j\to\infty} \rho^j (p_{t+j} - d_{t+j}) = 0$.

[5] For simplicity, this intuition abstracts from the potential feedback mechanism in which asset prices can influence real investment decisions. See Bond, Edmans, and Goldstein (2012) and Goldstein and Yang (2017) for excellent reviews of the real effects of financial markets.

using either expectation operator. In objective expectations,

$$\begin{aligned} r^e_{t+1} - E_t r^e_{t+1} &= (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} - r_{f,t+1+j}) - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r^e_{t+1+j} \\ &= N_{CF,t+1} - N_{DR,t+1}. \end{aligned} \tag{3}$$

The news components $N_{CF,t+1}$ and $N_{DR,t+1}$ reflect changes in objective expectations about dividend growth and excess returns, respectively. In subjective expectations,

$$\begin{aligned} r^e_{t+1} - E^*_t r^e_{t+1} &= (E^*_{t+1} - E^*_t) \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} - r_{f,t+1+j}) - (E^*_{t+1} - E^*_t) \sum_{j=1}^{\infty} \rho^j r^e_{t+1+j} \\ &= N^*_{CF,t+1} - N^*_{DR,t+1}. \end{aligned} \tag{4}$$

Positive subjective cash-flow news ($N^*_{CF,t+1} > 0$) represents an upward revision in the representative investor's expectations of future cash flows. She drives the price higher, such that the return is positive. In contrast, positive subjective discount-rate news ($N^*_{DR,t+1} > 0$) reflects an increase in discount rates for future cash flows, which results in a price decline and a negative return.

To provide a framework for understanding the connections between changes in subjective and objective beliefs, we subtract equation (4) from equation (3) to get

$$N_{DR,t+1} = N^*_{DR,t+1} - (N^*_{CF,t+1} - N_{CF,t+1}) + (E_t r^e_{t+1} - E^*_t r^e_{t+1}). \tag{5}$$

Collectively, equations (1) to (5) illustrate that objective expected returns reflect subjective beliefs, both in levels and in changes over time. In levels, equations (1) and (2) show that objective expected returns are low when prices — determined by subjective beliefs — are high relative to objective fundamentals. In changes, equation (5) demonstrates that news about objective expected returns has three sources. First, when the representative investor raises her subjective discount rates ($N^*_{DR,t+1} > 0$), she drives down the price and mechanically increases objective expected returns. Second, when she becomes more optimistic about cash flows relative to objective reality ($N^*_{CF,t+1} > N_{CF,t+1}$), she drives up the price relative to fundamentals, thereby decreasing objective expected returns. Third, there is a (typically) small, mechanical effect of the initial gap between objective and subjective expected returns ($E_t r^e_{t+1} - E^*_t r^e_{t+1}$).

To demonstrate how the properties of objective discount-rate and cash-flow news reflects the behavior of subjective beliefs, we present three special cases that depend on the following conditions.

Conditions 1A and 1B relate to cash-flow news, and Conditions 2A and 2B relate to discount-rate news.

Condition 1A. *Subjective and objective cash-flow news are perfectly positively correlated but differ in magnitude (i.e., pure overreaction or underreaction).*

Condition 1B. *Subjective and objective cash-flow news are imperfectly correlated.*

Condition 2A. *The subjective expected excess return is constant.*

Condition 2B. *The change in subjective expected excess return is (directly or inversely) proportional to the change in subjective expected cash-flow growth at each horizon.*

While these conditions simplify reality, they shed light on key mechanisms for the relations between subjective beliefs and objective expected returns.

Case 1. *Conditions 1A and 2A — perfect cash-flow news correlation with constant subjective expected return.*

This case, which closely approximates the model of Bordalo, Gennaioli, La Porta, and Shleifer (2024), illustrates how cash-flow overreaction and underreaction affect objective expected returns. Under Condition 1A, subjective cash-flow news can be expressed as $N^*_{CF,t+1} = (1+\theta)N_{CF,t+1}$, where θ measures the degree of overreaction $[\theta > 0]$ or underreaction $[\theta \in (-1, 0)]$ to objective cash-flow news. Under Condition 2A, $N^*_{DR,t+1} = 0$ and $E^*_t r^e_{t+1} = E^* r^e_{t+1}$, as the subjective expected excess return is constant. Equation (5) therefore simplifies to

$$N_{DR,t+1} = -\theta N_{CF,t+1} + (E_t r^e_{t+1} - E^* r^e_{t+1}). \tag{6}$$

When $\theta > 0$, the representative investor overreacts to cash-flow news. A positive fundamental shock ($N_{CF,t+1} > 0$) increases her subjective optimism relative to reality, and the resulting price overshoot causes a decrease in objective expected returns ($N_{DR,t+1} < 0$). Conversely, when $\theta \in (-1, 0)$, the representative investor underreacts. Given a positive fundamental shock, underreaction in subjective beliefs causes a muted price increase relative to the change in the fundamental value, which increases objective expected returns ($N_{DR,t+1} > 0$). Thus, overreaction induces a negative relation between $N_{DR,t+1}$ and $N_{CF,t+1}$, whereas underreaction creates a positive one.

Panel A of Table I provides illustrative examples of these mechanisms.[6] Consider a positive 10% objective cash-flow shock. A FIRE investor would increase the price by 10%. An overreacting

[6]For simplicity, the examples in Table I assume $E_t r^e_{t+1} = E^*_t r^e_{t+1} = 0$ and no change to the FIRE risk premium.

representative investor with $\theta = 0.5$ perceives a 15% cash-flow shock and, with constant subjective expected returns, increases the price by 15%. The -5% objective discount-rate news reconciles equation (3) and implies a decrease in objective expected returns. In contrast, an underreacting investor with $\theta = -0.5$ increases the price by only 5%, with the 5% objective discount-rate news reflecting an increase in objective expected returns.

Finally, the variance of $N_{DR,t+1}$ that is inherited from representative investor misreactions to $N_{CF,t+1}$ is $\theta^2 Var(N_{CF,t+1})$, so objective discount-rate news is more volatile the larger is the degree of overreaction or underreaction.

Case 2. *Conditions 1A and 2B — perfect cash-flow news correlation with time-varying subjective expected return.*

This case extends the previous one to illustrate how changes in subjective expected returns affect objective expected returns. Under Condition 2B, changes in subjective discount rates are proportional to changes in subjective cash-flow expectations. We express this relation as $N^*_{DR,t+1} = \frac{\lambda}{1+\theta} N^*_{CF,t+1}$, where the sign of λ determines whether subjective expected returns move directly ($\lambda > 0$) or inversely ($\lambda < 0$) with subjective cash-flow expectations. Given this expression and Condition 1A, $N^*_{DR,t+1} = \lambda N_{CF,t+1}$ and subjective discount rates are either procyclical ($\lambda > 0$) or countercyclical ($\lambda < 0$). Equation (5) simplifies to

$$N_{DR,t+1} = (\lambda - \theta) N_{CF,t+1} + (E_t r^e_{t+1} - E^* r^e_{t+1}). \tag{7}$$

Survey evidence on subjective expected returns is mixed. If $\lambda > 0$ [e.g., Giglio, Maggiori, Stroebel, and Utkus's (2021) evidence on retail investors], positive cash-flow news increases subjective discount rates, dampening the price response. If $\lambda < 0$ [e.g., Dahlquist and Ibert's (2024) evidence on institutional investors], positive cash-flow news decreases subjective discount rates, amplifying the price response. Given the mixed survey evidence across investor types, there is uncertainty about the sign of λ for the representative investor.

Panel B of Table I illustrates the mechanisms. Consider first procyclical discount rates ($\lambda = 0.2$). If the investor overreacts ($\theta = 0.5$), the 15% perceived increase in cash flows is partially offset by higher subjective discount rates, resulting in an 13% price increase. The -3% objective discount-rate news reflects mispricing that will eventually be corrected. With underreaction ($\theta = -0.5$), procyclical discount rates compound the effect, leading to substantial undervaluation with only a 3% price increase. The increase in expected subsequent returns is reflected by the 7% objective discount-rate news.

When discount rates are countercyclical ($\lambda = -0.2$), the effects reverse. Overreaction ($\theta = 0.5$) is amplified as the 15% perceived increase in cash flows combines with a decrease in subjective discount rates to push the price up by 17%. This large overvaluation implies low objective return expectations, reflected in the -7% objective discount-rate news. With underreaction ($\theta = -0.5$), the decrease in discount rates partially offsets the sluggish response to fundamental news, leading to a 7% price increase. The 3% objective discount-rate news indicates modest undervaluation.

Equation (7) and Table I show that time-varying subjective expected returns affect the relation between objective discount-rate and cash-flow news. Procyclical subjective discount rates ($\lambda > 0$) attenuate the negative relation between $N_{DR,t+1}$ and $N_{CF,t+1}$ caused by overreaction but exacerbate the positive relation caused by underreaction. Countercyclical subjective discount rates ($\lambda < 0$) have the opposite effect, strengthening overreaction but weakening underreaction.

The objective discount-rate news variance depends on $(\lambda-\theta)^2 Var(N_{CF,t+1})$, so $N_{DR,t+1}$ is more volatile in the cases with amplification (i.e., $\lambda > 0$ and $\theta < 0$ or $\lambda < 0$ and $\theta > 0$).

Case 3. *Conditions 1B and 2A — imperfect cash-flow news correlation with constant subjective expected return.*

This case extends Case 1 to show the effects of noise in subjective cash-flow expectations beyond overreaction or underreaction. Under Condition 1B, subjective cash-flow news can be expressed as $N^*_{CF,t+1} = (1+\theta)N_{CF,t+1} + \epsilon_{t+1}$, where ϵ_{t+1} is orthogonal to $N_{CF,t+1}$. Equation (5) becomes

$$N_{DR,t+1} = -\theta N_{CF,t+1} - \epsilon_{t+1} + (E_t r^e_{t+1} - E^* r^e_{t+1}). \tag{8}$$

Subjective belief revisions that are unrelated to the fundamental shock may be too optimistic ($\epsilon_{t+1} > 0$) or pessimistic ($\epsilon_{t+1} < 0$) in any given period. Panel C of Table I shows that the price effects of such optimism (pessimism) drive down (up) objective expected returns. Because this variation is orthogonal to objective cash-flow news, higher volatility of ϵ_{t+1} will lead to greater attenuation of the correlation between objective discount-rate news and cash-flow news.

The variance of objective discount-rate news has a $\theta^2 Var(N_{CF,t+1})$ term, as in Case 1, and a $Var(\epsilon_{t+1})$ term, which reflects the variance of cash-flow expectation errors that are unrelated to fundamentals. If subjective cash-flow beliefs are quite noisy, $N_{DR,t+1}$ will be very volatile.

2.2 Measuring objective news components

The framework in Section 2.1 demonstrates that the objective cash-flow news and discount-rate news components are informative about the representative investor's subjective beliefs. Estimating the objective return components in equation (3) presents significant empirical challenges. The conventional approach of using a vector autoregression (VARs) with a small set of predictors to estimate objective cash-flow news and discount-rate news has proven inadequate, as demonstrated by low correlations across model specifications in Figure 1. Our novel methodology addresses shortcomings of this approach by combining information from multiple predictors while maintaining parsimony. The strong, robust performance of our method suggests it provides a more accurate view of objective reality from which to learn about subjective beliefs.

We propose a simple approach to estimate objective discount-rate and cash-flow news using model forecast combinations in conjunction with the VAR approach.[7] In short, we use the means of model return forecasts from subsets of low-, medium-, and high-persistence predictors as state variables in the VAR and otherwise maintain Campbell's (1991) approach. Our method has several inherent advantages relative to the conventional approach.

First, we specify a parsimonious VAR that retains information from a large set of predictors. The literature has identified a multitude of return predictors, any of which could be reasonable to include. Candidate predictors reflect both common and unique information about future returns [e.g., Granger (1989) and Rapach, Strauss, and Zhou (2010)], such that we would lose information by discarding variables. The model combination approach strengthens common signals while retaining unique information from predictors. Granger (1989) further notes that researchers should include as many forecasts as possible. Our method allows us to do so while avoiding the overfitting issues of "kitchen sink" models [e.g., Goyal and Welch (2008)].

Second, our approach gains the noted stability of the model combination approach. Individual predictors are affected differently by problems like structural breaks [e.g., Pesaran and Timmermann (2005)] and unknown misspecification biases [e.g., Clemen (1989) and Stock and Watson (2004)]. Timmermann (2006) describes the substantial diversification benefits of combining forecasts from different models, and evidence across decades demonstrates that model forecast combination improves performance [e.g., Clemen (1989), Timmermann (2006), and Petropoulos et al. (2022)]. Specific to our context, Rapach, Strauss, and Zhou (2010) demonstrate several desir-

[7] Model forecast combination has deep roots in the statistics literature [following a seminal study by Bates and Granger (1969)], and the approach we use is classified as an ensemble learning method within the field of machine learning.

able properties of stock market return predictions using a set of 15 predictors from Goyal and Welch (2008).

Third, we model expected market returns over the short, medium, and long terms. Acknowledging the importance of horizon is crucial for estimating the forward-looking news components in equation (3). Individual predictors differ in their persistence and forecast market returns at high or low frequencies. Restricting to a few of these predictors would hamper our ability to model expected market returns across the disparate horizons that determine discount-rate news.

Let each of the K predictor variables have a sample period length of $T+1$ months. We estimate univariate predictive regressions,

$$r^e_{t+1} = a_i + b_i x_{i,t} + \eta_{i,t+1}, \tag{9}$$

over the months $t = 1, \ldots, T$. We compute the time series of the fitted expected excess return, $\hat{r}^e_{i,t+1} = \hat{a}_i + \hat{b}_i x_{i,t}$, for variables $i = 1, \ldots, K$ and months $t = 1, \ldots, T+1$. To divide the predictors into subgroups based on persistence, we also estimate autoregressions,

$$x_{i,t+1} = \alpha_i + \phi_i x_{i,t} + \nu_{i,t+1}, \tag{10}$$

and retain the persistence estimate $\hat{\phi}_i$.

We group the K predictors into k subgroups according to $\hat{\phi}_i$. In our base case, $k = 3$ and we sort predictors into low-, medium-, and high-persistence groups, with persistence cutoffs based on model fit (using the Akaike information criterion) for estimating market returns at horizons from one month to five years. We label the predictors in increasing order of persistence. The K_1 predictors with estimated half-lives of less than three months (i.e., $\log(0.5)/\log(\hat{\phi}_i) < 3$, so the cutoff for $\hat{\phi}_i$ is about 0.794) form the low-persistence group. The K_2 variables with half-lives greater than or equal to three months but less than four years (i.e., $3 \leq \log(0.5)/\log(\hat{\phi}_i) < 48$, so the upper cutoff is about 0.986) are the medium-persistence group, and the remaining K_3 predictors are the high-persistence group.[8] The $k = 3$ case generates rich dynamics in return forecasts across horizons.[9]

[8] The persistence cutoff is associated with a half-life of 48 months for the case with $k = 2$.

[9] In the Internet Appendix, we decompose discount-rate news variance by horizon to study these dynamics. The discount-rate news variance for the $k = 1$ case is heavily weighted towards the short run, as 45% of variance is from discount-rate shocks at horizons of one year or less, 85% is within three years, and almost 100% is within ten years. For the $k = 3$ case, 26% of discount-rate news variance is from shocks at horizons of one year or less, 61% is within three years, and 93% is within ten years, such that the model gains meaningful information about discount rates at short and long horizons. See the Internet Appendix for additional details.

The model combination literature considers a variety of approaches to weight individual forecasts, and we choose a simple mean combination approach. Equal weights are commonly used in the combination literature and typically perform well relative to alternative weighting strategies [Timmermann (2006)].[10] Within each predictor subgroup, the mean combination approach simply averages the return forecasts, $\hat{r}^e_{i,t+1} = \hat{a}_i + \hat{b}_i x_{i,t}$, from each of the candidate models in the subgroup. For example, the low-persistence group produces

$$\tilde{r}^e_{1,t+1} = \frac{1}{K_1} \sum_{i=1}^{K_1} \hat{r}^e_{i,t+1}. \tag{11}$$

Each $\tilde{r}^e_{t+1}$ vector is $(T+1) \times 1$, with the first T elements being predictors that are available in months $t = 1, \ldots, T$ to forecast market returns in months $t = 2, \ldots, T+1$ and the last element $\tilde{r}^e_{T+2}$ being required for estimation of the VAR.

Our method is to use these combination return forecasts as state variables in the VAR. In the base case with $k = 3$, we estimate the VAR,

$$\begin{bmatrix} r^e_{t+1} \\ \tilde{r}^e_{1,t+2} \\ \tilde{r}^e_{2,t+2} \\ \tilde{r}^e_{3,t+2} \end{bmatrix} = a + A \begin{bmatrix} r^e_t \\ \tilde{r}^e_{1,t+1} \\ \tilde{r}^e_{2,t+1} \\ \tilde{r}^e_{3,t+1} \end{bmatrix} + \varepsilon_{t+1}. \tag{12}$$

Following Campbell (1991), the discount-rate news component implied by the VAR is

$$N_{DR,t+1} = e1' \rho A (I - \rho A)^{-1} \varepsilon_{t+1}, \tag{13}$$

where $e1$ is a $(k+1) \times 1$ vector with the first element equal to one and the remaining k elements equal to zero and I is a $(k+1) \times (k+1)$ identity matrix. The cash-flow news component is

$$N_{CF,t+1} = (e1' + e1' \rho A (I - \rho A)^{-1}) \varepsilon_{t+1}. \tag{14}$$

[10] Timmermann (2006) shows that applying equal weights to the forecasts of many models is optimal if the forecast errors from each model have equal variances and pairwise correlations. In the market return predictability literature, predictive R^2s are quite low in monthly return forecasts. Denoting $\widehat{E_t r^e_{t+1}}$ as a forecast from a predictive regression, a low predictive R^2 implies that $Var(\widehat{E_t r^e_{t+1}})$ is small relative to $Var(r^e_{t+1})$. Because of this fact, the forecast error $r^e_{t+1} - \widehat{E_t r^e_{t+1}}$ will have similar variance to, and will be highly correlated with, r^e_{t+1} for each model, and equal weights will be close to optimal. The mean combination approach also avoids sampling errors and model uncertainty in the estimation of optimal weights [Palm and Zellner (1992)].

2.3 Measuring subjective news components

Our subjective return decomposition follows Chen, Da, and Zhao (2013), such that we use analyst forecasts to measure subjective expectations of future cash flows. Each firm's implied cost of capital (ICC) is estimated by combining the earnings forecasts with current stock prices [see, e.g., Pástor, Sinha, and Swaminathan (2008)].[11] Cash-flow news captures price changes attributable to revisions in expected future cash flows holding ICC constant, while discount-rate news reflects price changes due to revisions in required returns holding expected cash flows constant. To construct market-level measures, we aggregate firm-level components using market capitalization weights, ensuring our decomposition captures broad market expectations.[12] Capital gain equals the sum of cash-flow news and discount-rate news in this framework, so we multiply the subjective discount-rate news by negative one for ease of comparison to our objective return decomposition.

3 Data and descriptive statistics

Our empirical analysis employs two approaches to decompose stock market returns: (i) an objective approach based on realized market returns and return predictors, and (ii) a subjective approach using valuations and analyst cash flow forecasts. This section describes our data for each approach.

3.1 Objective return decomposition

The sample is a balanced panel of monthly stock market excess returns and 25 predictor variables spanning January 1927 to December 2021. The market excess returns are from Kenneth French's website.[13] We use log excess returns multiplied by 100 to express them in percentage form. The 25 predictors are introduced below, and details on variable construction are in Appendix A.

Our predictor set begins with the expansive database of Goyal and Welch (2008), which provides 14 widely-studied variables that reflect varying economic aspects: valuation ratios (dividend-price ratio, dividend yield, earnings-price ratio, and book-to-market ratio), corporate decisions (dividend-earnings ratio and net stock issues), market conditions (stock variance), and macroeconomic indicators (Treasury bill yield, long-term yield, long-term return, term spread, default spread, default

[11] This calculation assumes steady-state growth rates in the terminal period, consistent with standard valuation models.

[12] Detailed procedures for sample selection and ICC calculation follow Chen, Da, and Zhao (2013).

[13] We thank Kenneth French for making these data available (https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html).

return, and inflation).[14]

We augment these predictors with three variables from John Campbell's work on return decompositions that do not have close analogues in Goyal and Welch (2008): the relative bill rate from Campbell (1991), the smoothed earnings-price ratio from Shiller (2000) and Campbell and Vuolteenaho (2004), and the small-stock value spread from Campbell and Vuolteenaho (2004).

We also include eight variables from recent studies in the top three finance journals. These predictors measure corporate decisions [net payout yield from Boudoukh, Michaely, Richardson, and Roberts (2007)], statistical features of returns [average correlation from Pollet and Wilson (2010); tail risk from Kelly and Jiang (2014); and average skewness from Jondeau, Zhang, and Zhu (2019)], market valuations [nearness to Dow historical high from Li and Yu (2012) and partial least squares (PLS) aggregated book-to-market ratio from Kelly and Pruitt (2013)], and market liquidity [Amihud (2002) illiquidity measure from Chen, Eaton, and Paye (2018) and zero return frequency from Lesmond, Ogden, and Trzcinka (1999) and Chen, Eaton, and Paye (2018)].

For ease of comparison across predictors, we standardize each predictor to have zero mean and unit standard deviation. Some variables are hypothesized to be positive return predictors and others negative.

3.1.1 Univariate and model combination predictability

We first examine each variable in isolation with univariate regressions of monthly excess returns on individual lagged predictors. For each variable, Table II reports the predictive regression coefficient estimate, its ordinary least squares (OLS) t-statistic, the predictive regression R^2. Statistical significance is based on OLS standard errors and two-tailed tests. We also report the autoregressive (AR) coefficient estimate from an AR(1) for each predictor.

The univariate results in Table II reveal considerable heterogeneity in predictive power across variables. Of our 25 predictors, only eight demonstrate statistically significant predictability, with varying levels of significance: payout yield and PLS aggregated book-to-market at the 1% level; book-to-market, net stock issues, Treasury bill yield, and smoothed earnings-price at the 5% level; and dividend yield and zero return frequency at the 10% level. The economic magnitude of predictability is modest. The strongest predictor, PLS aggregated book-to-market, achieves an R^2 of 2.23%, while others explain 0.70% or less of return variation.[15] The persistence properties of the

[14]We thank Amit Goyal for making these data available (https://sites.google.com/view/agoyal145).

[15]We extensively validate our constructed predictors: (i) for variables with partial time series available from the authors, we ensure high correlations in the overlapping periods; (ii) for several variables without available time series, we obtain return predictability results similar to the original studies if we match the sample period; and (iii)

predictors also vary substantially. For example, the valuation ratios show high persistence (i.e., AR(1) coefficients above 0.98), while some macroeconomic indicators (long-term return, default return, and inflation) and average skewness exhibit lower persistence (AR(1) coefficients below 0.50). This variation proves crucial to capture rich horizon effects in expected returns within our approach.

Figure 2 illustrates the model combination approach, displaying fitted expected excess returns in percent per month for different predictor groupings. Panel A shows the case in which all 25 predictors form a single group ($k = 1$). The estimates appear to capture economically meaningful variation in expected returns, in support of Rapach, Strauss, and Zhou's (2010) arguments for using combination forecasts. The fitted expected return rises sharply during NBER recessions (shaded areas) and shows a secular decline over the sample period, consistent with estimates by Fama and French (2002) and Avdis and Wachter (2017). Panels B and C group predictors based on their persistence levels, with $k = 2$ and $k = 3$, respectively. While the forecast combinations maintain countercyclical patterns, their distinct persistence levels generate notably different time-series dynamics.

Table III presents evidence on the predictive power of our model combinations. The single-group combination ($k = 1$) demonstrates strong predictability, with a t-statistic of 6.16 and an R^2 of 3.23%.[16] This performance is remarkable for market return predictability with a monthly horizon and far exceeds the performance of the best individual predictor, illustrating the benefits of model combination.

Our base case ($k = 3$) reveals distinct horizons of predictability. The medium-term predictor ($\tilde{r}^e_{2,t+1}$) has the strongest performance (t-statistic of 6.89 and an R^2 of 4.01%). The short-term predictor ($\tilde{r}^e_{1,t+1}$) and long-term predictor ($\tilde{r}^e_{3,t+1}$) are also statistically significant at the 1% level with t-statistics of 3.20 and 2.67, respectively. Grouping the individual predictors by persistence produces the desired variation in horizons, with AR(1) coefficients of 0.389 (half-life of about one month), 0.965 (half-life of about 19 months), and 0.993 (half-life of about 100 months). These distinct persistence patterns, ranging from monthly to multi-year horizons, suggest that return predictability operates through multiple economic channels.

for variables for which replication proves challenging, we closely match the independent replication attempts from Kothari and O'Doherty (2023) and Goyal, Welch, and Zafirov (2024).

[16] As additional validation of the model combination approach, Appendix B presents results on out-of-sample forecasting. The forecast combinations have statistically significant out-of-sample R^2s [following Goyal and Welch (2008)] in all designs, and the economic magnitudes of performance gains are large according to Campbell and Thompson (2008). Kelly, Malamud, and Zhou (2024) emphasize Sharpe ratios of out-of-sample trading strategies, and the forecast combination strategy produces statistically significant Sharpe ratio gains relative to a buy-and-hold market strategy.

3.2 Subjective return decomposition

The subjective return decomposition, following Chen, Da, and Zhao (2013), uses data from Thomson Reuters I/B/E/S Estimates Database from 1985 to 2021. I/B/E/S Summary Statistics contains consensus analyst forecasts for earnings per share (EPS) across multiple forecast horizons. Firm-level accounting data are from Compustat.

For each firm in our sample, we collect quarterly data (March, June, September, and December) on earnings forecasts for different time horizons: current fiscal year ($FY1$), next two fiscal years ($FY2$ and $FY3$), and long-term growth (LTG). To ensure data quality, we require firms to have: (i) non-missing forecasts for $FY1$, $FY2$, and LTG; (ii) coverage by at least three analysts; and (iii) a minimum of 16 months of continuous coverage. We obtain share prices and shares outstanding from I/B/E/S, with all earnings forecasts and prices adjusted for stock splits.

4 Results

We study how subjective beliefs move stock prices using return decompositions. We do so in two complementary ways. First, we examine the time-series properties of objective cash-flow and discount-rate news, which are informative about the representative investor's subjective expectations according to the theoretical framework in Section 2.1. Sections 4.1 to 4.3 take this approach. Section 4.1 presents estimates of objective return components from our novel model combination VAR approach, and Section 4.2 demonstrates the stability and validity of the estimates. Section 4.3 draws inferences about the representative investor's subjective beliefs based on the dynamics of objective cash-flow and discount-rate news in the full sample and around recessions. Second, we use the subjective return components estimated from analyst forecasts to directly compare objective and subjective beliefs. Section 4.4 considers the joint behavior of objective and subjective return components in the historical time series and around recessions.

4.1 VAR estimates

We begin by presenting estimates for the VAR described in Section 2.2. As previously described, our approach uses a model combination approach to specify a parsimonious VAR that retains information from a large number of predictors. The predictors are grouped by persistence to capture horizon effects in expected market returns. Our base case has $k = 3$ groups. Table IV shows the A matrix of VAR coefficients, the R^2 from each regression, and the Σ matrix for the VAR residuals. OLS t-statistics for the VAR coefficients are in parentheses.

The top row of Table IV contains estimates for the excess market return predictability regression in the VAR. The lagged excess market return is a statistically insignificant predictor. The $\tilde{r}^e_{1,t+1}$ and $\tilde{r}^e_{2,t+1}$ model forecast combinations are significant predictors (t-statistics of 2.28 and 5.89, respectively), whereas the high-persistence variable $\tilde{r}^e_{3,t+1}$ is a positive but insignificant predictor (t-statistic of 0.49) after controlling for the other three variables. The predictive regression R^2 is an impressive 4.93%.

The VAR coefficient estimates in Table IV also show that the combination forecasts are the dominant predictors of themselves with the expected persistence levels. The Σ estimate implies short-term continuation (due to a positive covariance between the shocks to r^e_{t+1} and $\tilde{r}^e_{1,t+2}$) and mean reversion in the medium and long terms (due to negative covariances between shocks to r^e_{t+1} and both of $\tilde{r}^e_{2,t+2}$ and $\tilde{r}^e_{3,t+2}$). The persistence-based state variables allow the VAR to capture rich horizon-based patterns in expected returns.

4.2 Validation and stability of our method

The objective return decomposition in equation (3) splits excess returns into three components: expected excess return, discount-rate news, and cash-flow news. We proceed to show that our method successfully models each of the three components: (i) the estimated expected excess return forecasts excess returns in the next month, (ii) the estimated discount-rate news forecasts future excess returns, and (iii) the estimated cash-flow news forecasts future cash flows.

As previously discussed, Table IV shows that our model for the expected excess return is impressive with a predictive regression R^2 of 4.93%. In comparison, a standard VAR approach directly using three predictor variables as state variables tends to produce much less evidence of predictability. In untabulated results, we find that of the 2,300 such models that each use three of the 25 predictors, just six (0.26%) produce an R^2 above 4.93%. The average R^2 is 1.76%, and 86% of models have R^2s below 2.00%. Our method is highly successful in distilling information from the 25 predictors to model the objective expectation of the excess market return.

Next, we validate our estimates of objective discount-rate news and cash-flow news. A properly specified model should produce news components that align with their theoretical properties according to equation (3): discount-rate news should predict excess returns but not cash flows, and cash-flow news should predict dividend growth in excess of the risk-free rate but not excess returns. Our results show that our method is highly effective at achieving these goals.

The news components apply to all future horizons, such that past news at different lags are infor-

mative about the current return and dividend growth. To aggregate past news, we follow Campbell, Giglio, and Polk (2013) and use a trailing exponentially weighted moving average approach. They place a weight of 0.08 on the current quarterly estimate and a weight of 0.92 on the past moving average. We use a monthly analogue with weights of $[1 - 0.92^{(1/3)}]$ and $0.92^{(1/3)}$, so the moving average for cash-flow news, for example, is

$$MA_t(N_{CF}) = 0.0274 \times N_{CF,t} + 0.9726 \times MA_{t-1}(N_{CF}). \tag{15}$$

Figure 3 plots $MA_t(N_{CF})$ and $-MA_t(N_{DR})$, where the negative sign reflects the negative relation between discount-rate news and unexpected returns in equation (3). Our model implies that the historic price increases leading up to the market crashes in 1929 and the early 2000s were accompanied by sharp declines in objective expected returns. We also estimate that objective cash-flow news contributed strongly to price drops at the onset of the Great Depression, upon the bursting of the Internet bubble, and in late 2008 when the Great Recession deepened.

To validate our estimates, Table V reports results from forecasts of annual returns and dividend growth.[17] The monthly regressions use overlapping annual return and dividend growth observations, so we use Newey-West (1987) standard errors with 11 lags. Panel A forecasts the annual log excess market return,

$$\log(R^e_{t\to t+12}) = a_r + b_r MA_t(N_{DR}) + c_r MA_t(N_{CF}) + \epsilon_t, \tag{16}$$

with c_r either restricted to be zero or left unrestricted. Panel B forecasts the annual log dividend growth in excess of the log risk-free rate,

$$\log(\Delta D_{t\to t+12}) - \log(R^f_{t\to t+12}) = a_d + b_d MA_t(N_{DR}) + c_d MA_t(N_{CF}) + \eta_t, \tag{17}$$

with b_d either restricted or unrestricted.

The predictive regression results in Table V are consistent with our hypotheses. In Panel A, discount-rate news successfully predicts stock market returns with a coefficient of 10.64 and a t-statistic of 3.64. The predictive regression R^2 for annual returns is very high at 20.22%. In Panel B, cash-flow news strongly forecasts dividend growth with a coefficient of 13.78 (t-statistic of 5.01) and an R^2 of 18.13%. Importantly, cash-flow news does not forecast returns (coefficient of 0.29 with a

[17] We calculate the cumulative annual logs of the excess market return and the risk-free rate. We follow Koijen and Van Nieuwerburgh (2011) to construct monthly dividends on the CRSP value-weighted stock market portfolio with no reinvestment and sum the monthly dividends over each 12-month period to avoid dividend seasonalities. We then calculate a monthly time series of annual log dividend growth by dividing the annual sum of dividends by its 12-month lag and taking the log.

t-statistic of 0.06) and discount-rate news does not forecast dividends (coefficient of -0.31 with a t-statistic of -0.27). The R^2s are virtually unchanged when moving from restricted to unrestricted regressions. Our approach is successful in isolating the discount-rate and cash-flow components.

Our model achieves these outcomes by specifying a parsimonious VAR that retains information from a broad set of predictors through our model combination approach. A limitation of the conventional VAR approach is that researchers must restrict themselves to just a few predictors to maintain parsimony. Figure 4 compares our approach with the set of 2,300 VAR models that directly use three of the 25 predictor variables as state variables. Panels A, C, and E summarize the results for return forecasts, and Panels B, D, and F show results for dividend growth forecasts.

Figure 4 demonstrates that our model combination approach dominates the conventional approach. When evaluating statistical significance, the t-statistics for the discount-rate news coefficient and the cash-flow news coefficient each exceed those for about 99% of the VARs. For return forecasts, our predictive regression achieves a higher R^2 compared with any of the VARs (20.22% R^2 versus a maximum of 14.67%). For dividend growth forecasts, our approach's R^2 of 18.13% exceeds the R^2s for 94% of the conventional VARs. These results demonstrate the importance of using model forecast combination to model expected returns in this setting.

Researchers implementing our approach could reasonably make different modeling choices from ours. In Appendix B, we demonstrate the stability of our approach to the choices of predictors and persistence cutoffs. Researchers could choose different predictors and a different number of predictors K. We show that our method produces news estimates that are very highly correlated with our base case as long as researchers use enough variables to include at least four or five predictors in each persistence subgroup. Frequently updated data on Amit Goyal's website allow researchers to readily access 15 or more predictors to easily implement our method with large K. For the persistence cutoff choices, our estimates are stable while varying the shorter cutoff half-life between one and 36 months or the longer cutoff half-life between 12 months and 120 months.

4.3 Objective news component implications for subjective beliefs

Our theoretical framework in Section 2.1 demonstrates that properties of objective cash-flow and discount-rate news are informative about the representative investor's subjective expectations. Section 4.3.1 studies the full time series. Section 4.3.2 concentrates on the periods around recessions.

4.3.1 Full sample inferences

Table VI shows time-series statistics for objective discount-rate news and cash-flow news. Two important empirical patterns emerge. The monthly standard deviation of N_{DR} is 5.90%, which is twice as high as the N_{CF} standard deviation of 2.84%. The correlation between N_{DR} and N_{CF} is positive and equal to 0.46, which is near Campbell, Giglio, and Polk's (2013) restricted VAR estimate of 0.58. These empirical patterns in objective return components provide valuable insights into the representative investor's subjective beliefs, as they suggest systematic deviations from rational expectations.

The larger volatility of discount-rate news relative to cash-flow news is consistent with Campbell's (1991) seminal finding that stock return variance is largely driven by fluctuating discount rates. While this volatility pattern can be accommodated within FIRE frameworks, the positive correlation between discount-rate news and cash-flow news presents a direct challenge to FIRE predictions. Under FIRE, objective discount-rate news reflects time-varying risk premiums driven by fundamental economic mechanisms. For example, Campbell and Cochrane (1999) show how habit formation generates countercyclical risk aversion and expected returns, and Bansal and Yaron (2004) demonstrate that long-run risks combined with time-varying economic uncertainty can produce volatile risk premiums. Importantly, FIRE frameworks such as Campbell and Cochrane (1999) and Gabaix (2012) also predict that discount-rate news and cash-flow news should be reinforcing. Poor cash-flow news in these models is accompanied, on average, by an increase in the risk premium which decreases the price. This reinforcing behavior manifests as a model-implied negative correlation between discount-rate and cash-flow news, which is inconsistent with our positive estimate of 0.46.

Our framework suggests alternative mechanisms to explain these two empirical patterns. Starting with the volatility results, Case 1 demonstrates that either overreaction or underreaction to cash-flow news generates variation in objective discount-rate news. The effect on variance is proportional to θ^2, where θ is the parameter capturing the degree of overreaction or underreaction. Case 2 shows that time-varying subjective discount rates can either amplify or attenuate these patterns. With underreaction, procyclical subjective discount rates ($\lambda > 0$) increase discount-rate volatility. Countercyclical rates ($\lambda < 0$) have the opposite effect. Case 3 provides a distinct perspective on the high volatility of objective discount-rate news. Whereas FIRE models attribute this volatility to time-varying risk premiums, Case 3 demonstrates that noisy subjective cash-flow expectations can generate substantial variation in objective discount-rate news even with constant risk premiums.

An important piece of the puzzle is the positive correlation between discount-rate news and cash-flow news. Under Case 1, overreaction ($\theta > 0$) produces a negative correlation as price movements overshoot fundamentals, while underreaction ($\theta < 0$) generates a positive correlation as prices adjust sluggishly. Under Case 2, with underreaction, procyclical subjective discount rates ($\lambda > 0$) strengthen the positive correlation between objective discount-rate and cash-flow news. Under Case 3, noisy subjective cash-flow beliefs attenuate the correlation.

Taken together, the positive correlation between news components points toward underreaction to cash-flow news as the primary feature of the representative investor's subjective beliefs. Case 3, which incorporates both underreaction and noisy subjective cash-flow expectations, appears most consistent with the observed high objective discount-rate news volatility and the observed positive (but moderate in magnitude) correlation between objective discount-rate and cash-flow news. These mechanisms generate substantial variation in objective discount rates without requiring highly volatile risk premiums.

4.3.2 Recession period inferences

Given the importance of understanding market downturns and their implications for investors, we study the dynamics of objective expected returns, discount-rate news, and cash-flow news around recessions. The onsets of recessions are coincident with many of the large stock market declines in our sample, and Figure 3 suggests important discount-rate and cash-flow effects around these periods. We consider the two years leading up to a recession and the first two years of recession and recovery to examine systematic patterns in the return components.

Figure 5 plots the cumulative market returns, expected return contributions, discount-rate news contributions, and cash-flow news contributions averaged across the 15 NBER-dated recessions in our sample. Panel A tracks the average cumulative excess market return over the 24 months preceding a recession with month zero representing the month of the economic peak, and the red dotted lines provide a 90% bootstrap confidence interval. Panel B plots the cumulative excess market return during the first 24 months of the recession and recovery period, with month one representing the first month of the NBER recession. Panels C and D show the cumulative effects of objective expected excess returns. Panels E and F provide the cumulative contributions of objective discount-rate news (i.e., $-N_{DR}$, such that the plotted discount-rate effects contribute positively to returns), and Panels G and H are analogous for objective cash-flow effects (i.e., N_{CF}). The cumulative effects for the pre-recession (recession/recovery) period in Panels C, E, and G (Panels D, F, and H) sum to the cumulative excess returns in Panel A (Panel B).

Figure 5 demonstrates dynamics around recessions that are difficult to reconcile with FIRE. Panel A shows that pre-recession periods are marked, on average, by market increases that reach nearly 20% before leveling off around six months before the recession onset and averaging a cumulative excess return of 16% during the pre-recession period. Expected excess returns (Panel C) initially positively contribute to realized returns, but the slope flattens and turns slightly negative as the peak approaches. The negative slope of this contribution implies that the average monthly expected excess return is negative in the months before a recession. Our estimates suggest very large, positive cumulative discount-rate news (Panel E) leading to a recession. About 20% of the pre-recession price run-up is accounted for through declining future expected returns. Finally, Panel G shows that the declining discount rates are accompanied by poor cash-flow news, with an average cumulative effect of about -11%. The partially offsetting effects of objective discount-rate news ($-N_{DR}$ as plotted) and cash-flow news (N_{CF}) are consistent with the positive correlation between N_{DR} and N_{CF} reported in Table VI. On the whole, the pre-recession periods have rising prices in the face of poor objective cash-flow news, such that objective expected returns fall precipitously.

The recession/recovery periods in Figure 5 initially have large, statistically significant market declines (Panel B), a point recently emphasized by Gómez-Cram (2022). The cumulative excess return reaches an average low of -8% in month six of the recession. Panel D implies that a slight negative expected return contributes to this decline, but Panel F shows that the poor average excess return realization is accompanied by a large -12% discount-rate contribution from unexpected increases in objective expected returns. Once the recession starts, objective cash-flow news (Panel H) is roughly flat throughout the recession and recovery.

In untabulated results, we find these patterns to be particularly stark for a small set of six recessions with significantly negative expected returns according to a bootstrap analysis. The accompanying peaks include the 1929 and 2001 highs along with the economic peaks in 1937, 1969, 1973, and 1981. These episodes had, on average, price run-ups that occur despite poor objective cash-flow news, implying large declines in objective expected returns. These recessions also drive nearly all of the Figure 5 pattern of negative average excess returns early in recessions due to large, negative expected returns (averaging nearly -1% per month for the first six months) and discount-rate news contributions (averaging a cumulative -17% over the first six months).

In summary, our recession analysis reveals systematic patterns that seem fundamentally inconsistent with FIRE. It is unsurprising that objective cash-flow news is negative on average leading up to the onset of a recession. What is surprising is that, as objective cash-flow growth prospects

worsen, the stock market enjoys substantial gains. The pre-recession advances in the stock market drive objective expected returns into negative territory on average, a pattern that is difficult to reconcile within a FIRE framework. The patterns are, however, consistent with a representative investor who underreacts and has noisy cash-flow expectations, as in Case 3. Rather than properly perceiving the worsening growth prospects leading up to the recession, investors drive prices ever higher, which causes objective expected returns to decline. Investors only sluggishly react to the objectively poor cash-flow news as the signals become more overt with the onset of a recession. Stocks suffer as the recession begins, as anticipated by the negative objective discount-rate news in the pre-recession period.

4.4 Direct evidence of the objective-subjective belief gap

Our previous analyses indicate a substantial wedge between objective and subjective expectations, including around economic downturns. These inferences are based on our theoretical framework and empirical estimates of objective cash-flow and discount-rate news. In the last set of results, we turn to direct measures of subjective beliefs.

We use analyst forecasts to proxy for the representative investor's subjective expectations of future cash flows. Direct measures of investor beliefs are not readily available, so analyst information is often used as a proxy for investor beliefs [Adam and Nagel (2023)]. As detailed in Section 2.3, we follow Chen, Da, and Zhao (2013) for the subjective return decomposition. Because we use analyst data, the sample period for this analysis begins in 1985.

Table VII reports results from predictive regressions for returns and dividend growth, analogous to Table V. We use the $MA(N_{DR})$ and $MA(N_{CF})$ variables from our approach over this period and the one-year discount-rate news and cash-flow news estimates from Chen, Da, and Zhao (2013) as predictors.[18] The sample period is April 1985 to December 2021, and the subjective regressions have a quarterly frequency.

Panel A of Table VII reports that the objective $MA(N_{DR})$ variable positively forecasts returns with a t-statistic of 2.87 and an R^2 of 10.63% over this sample period. The subjective measure achieves lower predictive ability with a t-statistic of 1.73 and an R^2 of 3.72%. In Panel B, the objective $MA(N_{CF})$ displays impressive dividend growth forecastability with a t-statistic of 7.43 and an R^2 of 46.69%. The subjective measure positively forecasts dividend growth but with much weaker predictability: a t-statistic of 2.50 and an R^2 of 7.18%. The lower predictability of subjective

[18] We experimented with various smoothing methods for the Chen, Da, and Zhao (2013) estimates to aggregate information across periods. None was as effective as using the raw estimates.

return components for future returns and dividend growth indicates that subjective beliefs deviate from objective reality.

Panel A of Table VIII reports time-series properties for the objective and subjective return components, where the subjective news components are quarterly estimates following Chen, Da, and Zhao (2013). Discount-rate news accounts for more return variance than does cash-flow news under both objective and subjective expectations. Whereas the objective effects of discount-rate news and cash-flow news from our approach have a 0.31 correlation over this sample period, the subjective effects are even more strongly positively correlated at 0.69. Regardless of whether we examine objective or subjective expectations, the positive correlation is in line with Campbell, Giglio, and Polk's (2013) conclusion that market downturns reflect declining cash-flow expectations or increasing discount rates, but usually not both.

The direct comparison of objective and subjective return components in Panel B of Table VIII demonstrates a large disconnect between objective and subjective beliefs. The correlation between the discount-rate effects from the two models is moderate at 0.46. More striking is the correlation between cash-flow news measures, which is just 0.21. This finding and the weak forecasting ability of subjective cash-flow news for future dividend growth in Table VII are indicative of noisy subjective cash-flow expectations, which contributes to the large volatility of objective discount-rate news.

Importantly, these low correlations across objective and subjective expectations manifest around recessions. Figure 6 is analogous to Figure 5, but plots the cumulative capital gain, discount-rate news contribution, and cash-flow news contribution from the subjective return decomposition. The effects in Panels C and E (Panels D and F) sum to the cumulative capital gain in Panel A (Panel B). The sample period in Figure 6 covers only four recessions, but we verified that the important patterns in objective expectations about discount-rate news and cash-flow news from Figure 5 carry over to this sample.

Figure 6 demonstrates that subjective expectations paint a much different picture surrounding recessions relative to the objective decomposition. Prices run up while approaching the economic peak, but the attribution of these gains reveals a fundamental disconnect between objective reality and subjective beliefs.

Whereas the objective decomposition shows deteriorating cash-flow news accompanied by decreasing discount rates before recessions (Figure 5), analyst forecasts reveal the opposite pattern — increases in market valuations are accompanied by strongly positive subjective cash-flow news, with subjective discount-rate news having a slight negative contribution. This reversal in the perceived drivers of market performance suggests investors fundamentally misinterpret and underreact

to negative objective economic signals preceding downturns. The four recessions covered by the sample period did not, on average, have significantly negative returns or capital gains in the initial months of the recession, so we can glean less information about the recession and recovery periods from this analysis.

Taken together, Figures 5 and 6 are consistent with a large recession-period gap between objective and subjective beliefs. Whereas objective cash-flow expectations deteriorate in advance of recessions, investors (proxied by analysts) become more bullish about future cash-flow growth in these periods. Investors perceive that expected returns remain robustly positive as prices continue to rise. In the objective reality, the steady price increases combined with the worsening cash-flow outlook lead to lower and lower expected returns.

These patterns align with the underreaction mechanism from our theoretical framework in Section 2.1, where subjective beliefs adjust sluggishly to changing fundamentals. Furthermore, the substantial divergence between objective and subjective cash-flow expectations (correlation of just 0.21) provides strong empirical support for the noisy subjective cash-flow beliefs described in Case 3. The wedge between objective and subjective cash-flow expectations may explain why investors were readily holding stocks during the periods with negative objective expectations of future returns. They underreact to negative cash-flow signals as the recession approaches, leading to overly optimistic expectations that only gradually adjust as the recession unfolds.

Our findings have important implications for asset pricing theory. Rather than attributing market movements primarily to rational time-varying risk premiums, our results suggest that a substantial portion of stock price dynamics stems from systematic belief distortions — namely, underreaction to fundamental information combined with noisy subjective cash-flow expectations. Our estimates of volatility and correlation for the discount-rate and cash-flow news components provide new empirical moments for asset pricing models to match.

5 Conclusion

Growing survey evidence indicates that investor beliefs depart from FIRE. Given this context, addressing the fundamental question of why stock prices move is confounded by the inherent unobservability of objective and subjective beliefs. We develop a theoretical framework to demonstrate that the behavior of the *objective* discount-rate news and cash-flow news components are informative about the representative investor's *subjective* beliefs. Our novel return decomposition method improves on the conventional approach and provides robust estimates of these objective compo-

nents. Our analysis characterizes subjective beliefs using stock prices and state variables and, as such, avoids the potential pitfalls of relying solely on survey data.

Our empirical results indicate substantial deviations from FIRE. The properties of objective discount-rate and cash-flow news imply that the representative investor systematically underreacts to fundamental shocks and has noisy cash-flow expectations. Objective discount-rate variation is primarily driven by a dynamic subjective-objective gap in cash-flow expectations, contrasting with the FIRE framework's attribution of such variation to time-varying risk premiums. These effects are particularly pronounced in the periods leading up to recessions, and overly optimistic subjective cash-flow expectations help explain why investors happily invest in stocks despite negative objective expected returns. Our study contributes to a growing literature that documents that departures from rationality affect stock prices, and we demonstrate how objective return decompositions reflect subjective beliefs. Successful asset pricing models should jointly explain objective and subjective beliefs about cash flows and discount rates, and we provide important new empirical moments to match.

References

Adam, Klaus, and Stefan Nagel, 2023, Expectations Data in Asset Pricing, in Rüdiger Bachmann, Giorgio Topa, and Wilbert van der Klaauw, eds.: *Handbook of Economic Expectations* (Academic Press), Cambridge, MA.

Amihud, Yakov, 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.

Amromin, Gene, and Steven A. Sharpe, 2013, From the horse's mouth: Economic conditions and investor expectations of risk and return, *Management Science* 60, 845–866.

Andonov, Aleksandar, and Joshua D. Rauh, 2022, The return expectations of public pension funds, *Review of Financial Studies* 35, 3777–3822.

Atanasov, Victoria, and Thomas Nitschka, 2015, Foreign currency returns and systematic risks, *Journal of Financial and Quantitative Analysis* 50, 231–250.

Avdis, Efstathios, and Jessica A. Wachter, 2017, Maximum likelihood estimation of the equity premium, *Journal of Financial Economics* 125, 589–609.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Bates, John M., and Clive W. J. Granger, 1969, The combination of forecasts, *Operational Research Quarterly* 20, 451–468.

Bekaert, Geert, and Robert J. Hodrick, 1992, Characterizing predictable components in excess returns on equity and foreign exchange markets, *Journal of Finance* 47, 467–509.

Bernanke, Ben S., and Kenneth N. Kuttner, 2005, What explains the stock market's reaction to Federal Reserve policy?, *Journal of Finance* 60, 1221–1257.

Bond, Philip, Alex Edmans, and Itay Goldstein, 2012, The real effects of financial markets, *Annual Review of Financial Economics* 4, 339–360.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2024, Belief overreaction and stock market puzzles, *Journal of Political Economy* 132, 1450–1484.

Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael R. Roberts, 2007, On the importance of measuring payout yield: Implications for empirical asset pricing, *Journal of Finance* 62, 877–915.

Brunnermeier, Markus, Emmanuel Farhi, Ralph S. J. Koijen, Arvind Krishnamurthy, Sydney C. Ludvigson, Hanno Lustig, Stefan Nagel, and Monika Piazzesi, 2021, Perspectives on the future of asset pricing, *Review of Financial Studies* 34, 2126–2160.

Buraschi, Andrea, Ilaria Piatti, and Paul Whelan, 2022, Subjective bond returns and belief aggregation, *Review of Financial Studies* 35, 3710–3741.

Campbell, John Y., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.

——— , 1996, Understanding risk and return, *Journal of Political Economy* 104, 298–345.

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Campbell, John Y., Stefano Giglio, and Christopher Polk, 2013, Hard times, *Review of Asset Pricing Studies* 3, 95–132.

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

——— , 1988b, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661–676.

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

Campbell, John Y., and Tuomo Vuolteenaho, 2004, Bad beta, good beta, *American Economic Review* 94, 1249–1275.

Chen, Long, Zhi Da, and Xinlei Zhao, 2013, What drives stock price movements?, *Review of Financial Studies* 26, 841–876.

Chen, Long, and Xinlei Zhao, 2009, Return decomposition, *Review of Financial Studies* 22, 5213–5249.

Chen, Xiaohong, Lars Peter Hansen, and Peter G. Hansen, 2020, Robust identification of investor beliefs, *Proceedings of the National Academy of Science* 117, 33130–33140.

Chen, Yong, Gregory W. Eaton, and Bradley S. Paye, 2018, Micro(structure) before macro? The predictive power of aggregate illiquidity for stock returns and economic activity, *Journal of Financial Economics* 130, 48–73.

Clark, Todd E., and Kenneth D. West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291–311.

Clemen, Robert T., 1989, Combining forecasts: A review and annotated bibliography, *International Journal of Forecasting* 5, 559–583.

Cochrane, John H., 2011, Discount rates, *Journal of Finance* 66, 1047–1108.

Couts, Spencer J., Andrei S. Gonçalves, and Johnathan Loudis, 2024, The subjective risk and return expectations of institutional investors, Working paper, University of Southern California.

Dahlquist, Magnus, and Markus Ibert, 2024, Equity return expectations and portfolios: Evidence from large asset managers, *Review of Financial Studies* 37, 1887–1928.

De la O, Ricardo, and Sean Myers, 2021, Subjective cash flow and discount rate expectations, *Journal of Finance* 76, 1339–1387.

Fama, Eugene F., and Kenneth R. French, 2002, The equity premium, *Journal of Finance* 57, 637–659.

Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.

Ghosh, Anisha, and Guillaume Roussellet, 2023, Identifying beliefs from asset prices, Working paper, McGill University.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2021, Five facts about beliefs and portfolios, *American Economic Review* 111, 1481–1522.

Goldstein, Itay, and Liyan Yang, 2017, Information disclosure in financial markets, *Annual Review of Financial Economics* 9, 101–125.

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

Gonçalves, Andrei S., 2021, Reinvestment risk and the equity term structure, *Journal of Finance* 76, 2153–2197.

Goyal, Amit, and Ivo Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.

Goyal, Amit, Ivo Welch, and Athanasse Zafirov, 2024, A comprehensive 2022 look at the empirical performance of equity premium prediction, *Review of Financial Studies* 37, 3490–3557.

Granger, Clive W. J., 1989, Combining forecasts—Twenty years later, *Journal of Forecasting* 8, 167–173.

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

Hill, Bruce M., 1975, A simple general approach to inference about the tail of a distribution, *Annals of Statistics* 3, 1163–1174.

Jobson, J. D., and Bob M. Korkie, 1981, Performance hypothesis testing with the Sharpe and Treynor measures, *Journal of Finance* 36, 889–908.

Jondeau, Eric, Qunzi Zhang, and Xiaoneng Zhu, 2019, Average skewness matters, *Journal of Financial Economics* 134, 29–47.

Kelly, Bryan, and Hao Jiang, 2014, Tail risk and asset prices, *Review of Financial Studies* 27, 2841–2871.

Kelly, Bryan, Semyon Malamud, and Kangying Zhou, 2024, The virtue of complexity in return prediction, *Journal of Finance* 79, 459–503.

Kelly, Bryan, and Seth Pruitt, 2013, Market expectations in the cross-section of present values, *Journal of Finance* 68, 1721–1756.

Koijen, Ralph S. J., and Stijn Van Nieuwerburgh, 2011, Predictability of Returns and Cash Flows, *Annual Review of Financial Economics* 3, 467–491.

Kothari, Pratik, and Michael S. O'Doherty, 2023, Job postings and aggregate stock returns, *Journal of Financial Markets* 64, 100804.

Kroencke, Tim A., 2022, Recessions and the stock market, *Journal of Monetary Economics* 131, 61–77.

Lamont, Owen A., 1998, Earnings and expected returns, *Journal of Finance* 53, 1563–1587.

Lesmond, David A., Joseph P. Ogden, and Charles A. Trzcinka, 1999, A new estimate of transaction costs, *Review of Financial Studies* 12, 1113–1141.

Li, Jun, and Jianfeng Yu, 2012, Investor attention, psychological anchors, and stock return predictability, *Journal of Financial Economics* 104, 401–419.

Lucas, Robert E. Jr., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

Manski, Charles F., 2004, Measuring expectations, *Econometrica* 72, 1329–1376.

Memmel, Christoph, 2003, Performance hypothesis testing with the Sharpe ratio, *Finance Letters* 1, 21–23.

Muth, John F., 1961, Rational expectations and the theory of price movements, *Econometrica* 29, 315–335.

Nagel, Stefan, and Zhengyang Xu, 2023, Dynamics of subjective risk premia, *Journal of Financial Economics* 150, 103713.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Palm, Franz C., and Arnold Zellner, 1992, To combine or not to combine? Issues of combining forecasts, *Journal of Forecasting* 11, 687–701.

Pástor, Ľuboš, Meenakshi Sinha, and Bhaskaran Swaminathan, 2008, Estimating the intertemporal risk-return tradeoff using the implied cost of capital, *Journal of Finance* 63, 2859–2897.

Patelis, Alex D., 1997, Stock return predictability and the role of monetary policy, *Journal of Finance* 52, 1951–1972.

Pesaran, M. Hashem, and Allan Timmermann, 2005, Small sample properties of forecasts from autoregressive models under structural breaks, *Journal of Econometrics* 129, 183–217.

Petropoulos, Fotios, Daniele Apiletti, Vassilios Assimakopoulos, Mohamed Zied Babai, Devon K. Barrow, Souhaib Ben Taieb, Christoph Bergmeir, Ricardo J. Bessa, Jakub Bijak, John E. Boylan, Jethro Browell, Claudio Carnevale, Jennifer L. Castle, Pasquale Cirillo, Michael P. Clements, Clara Cordeiro, Fernando Luiz Cyrino Oliveira, Shari De Baets, Alexander Dokumentov, Joanne Ellison, Piotr Fiszeder, Philip Hans Franses, David T. Frazier, Michael Gilliland, M. Sinan Gönül, Paul Goodwin, Luigi Grossi, Yael Grushka-Cockayne, Mariangela Guidolin, Massimo Guidolin, Ulrich Gunter, Xiaojia Guo, Renato Guseo, Nigel Harvey, David F. Hendry, Ross Hollyman, Tim Januschowski, Jooyoung Jeon, Victor Richmond R. Jose, Yanfei Kang, Anne B. Koehler, Stephan Kolassa, Nikolaos Kourentzes, Sonia Leva, Feng Li, Konstantia Litsiou, Spyros Makridakis, Gael M. Martin, Andrew B. Martinez, Sheik Meeran, Theodore Modis, Konstantinos Nikolopoulos, Dilek Önkal, Alessia Paccagnini, Anastasios Panagiotelis, Ioannis Panapakidis, Jose M. Pavía, Manuela Pedio, Diego J. Pedregal, Pierre Pinson, Patrícia Ramos, David E. Rapach, J. James Reade, Bahman Rostami-Tabar, Michał Rubaszek, Georgios Sermpinis, Han Lin Shang, Evangelos Spiliotis, Aris A. Syntetos, Priyanga Dilini Talagala, Thiyanga S. Talagala, Len Tashman, Dimitrios Thomakos, Thordis Thorarinsdottir, Ezio Todini, Juan Ramón Trapero Arenas, Xiaoqian Wang, Robert L. Winkler, Alisa Yusupova, and Florian Ziel, 2022, Forecasting: Theory and practice, *International Journal of Forecasting* 38, 705–871.

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

Piazzesi, Monika, and Martin Schneider, 2009, Momentum traders in the housing market: Survey evidence and a search model, *American Economic Review: Papers and Proceedings* 99, 406–411.

Pollet, Joshua M., and Mungo Wilson, 2010, Average correlation and stock market returns, *Journal of Financial Economics* 96, 364–380.

Rapach, David, and Guofu Zhou, 2013, Forecasting Stock Returns, in Graham Elliott, and Allan Timmermann, eds.: *Handbook of Economic Forecasting* (Elsevier), pp. 328–383.

Rapach, David E., Matthew C. Ringgenberg, and Guofu Zhou, 2016, Short interest and aggregate stock returns, *Journal of Financial Economics* 121, 46–65.

Rapach, David E., Jack K. Strauss, and Guofu Zhou, 2010, Out-of-sample equity premium prediction: Combination forecasts and links to the real economy, *Review of Financial Studies* 23, 821–862.

Shiller, Robert J., 2000, *Irrational Exuberance*. Princeton University Press, Princeton, NJ.

Stock, James H., and Mark W. Watson, 2004, Combination forecasts of output growth in a seven-country data set, *Journal of Forecasting* 23, 405–430.

Timmermann, Allan, 2006, Forecast Combinations, in Graham Elliott, Clive W. J. Granger, and Allan Timmerman, eds.: *Handbook of Economic Forecasting* (North-Holland).

Vissing-Jorgensen, Annette, 2003, Perspectives on behavioral finance: Does "irrationality" disappear with wealth? Evidence from expectations and actions, *NBER Macroeconomics Annual* 18, 139–194.

Table I
Illustrative examples for the relation between subjective beliefs and objective reality.

The table provides illustrative examples for Cases 1, 2, and 3 of the relation between subjective beliefs and objective reality, introduced in Section 2.1. The parameters govern subjective cash-flow overreaction ($\theta > 0$) or underreaction ($\theta < 0$), subjective discount rate procyclicality ($\lambda > 0$) or countercyclicality ($\lambda < 0$), and noise in subjective cash-flow news that is too optimistic ($\epsilon_{t+1} > 0$) or pessimistic ($\epsilon_{t+1} < 0$). The examples consider the effects of a 10% objective cash-flow shock.

			FIRE	Subjective					Objective				
θ	λ	ϵ_{t+1}	$r^e_{t+1} - E_t r^e_{t+1}$	$r^e_{t+1} - E^*_t r^e_{t+1}$	=	$N^*_{CF,t+1}$	−	$N^*_{DR,t+1}$	$r^e_{t+1} - E_t r^e_{t+1}$	=	$N_{CF,t+1}$	−	$N_{DR,t+1}$
Panel A: Case 1 parameter scenarios													
0.5	—	—	10%	15%	=	15%	−	0%	15%	=	10%	−	−5%
−0.5	—	—	10%	5%	=	5%	−	0%	5%	=	10%	−	5%
Panel B: Case 2 parameter scenarios													
0.5	0.2	—	10%	13%	=	15%	−	2%	13%	=	10%	−	−3%
−0.5	0.2	—	10%	3%	=	5%	−	2%	3%	=	10%	−	7%
0.5	−0.2	—	10%	17%	=	15%	−	−2%	17%	=	10%	−	−7%
−0.5	−0.2	—	10%	7%	=	5%	−	−2%	7%	=	10%	−	3%
Panel C: Case 3 parameter scenarios													
0.5	—	3%	10%	18%	=	18%	−	0%	18%	=	10%	−	−8%
0.5	—	−3%	10%	12%	=	12%	−	0%	12%	=	10%	−	−2%
−0.5	—	3%	10%	8%	=	8%	−	0%	8%	=	10%	−	2%
−0.5	—	−3%	10%	2%	=	2%	−	0%	2%	=	10%	−	8%

Table II
Predictor variable univariate regressions.

The table summarizes results from univariate predictive regressions and autoregressions for 25 market return predictor variables that are commonly used in the finance literature. All of the predictor variables are standardized. For each variable, we report the predictive regression coefficient, its OLS t-statistic, and the R^2 (in percent) from a regression of the log excess market return on the lagged predictor value. We also report the autoregressive coefficient from a regression of the predictor variable on its own lag. The sample period is January 1927 to December 2021. An * (**) [***] on the predictive coefficient estimate indicates statistical significance at the 10% (5%) [1%] level in a two-tailed test.

Variable	Coefficient	t-stat	R^2(%)	AR(1)
Dividend-price ratio	0.18	1.11	0.11	0.995
Dividend yield	0.29*	1.84	0.30	0.994
Earnings-price ratio	0.25	1.59	0.22	0.988
Book-to-market ratio	0.32**	1.99	0.35	0.988
Dividend-earnings ratio	−0.07	−0.46	0.02	0.992
Stock variance	−0.06	−0.40	0.01	0.577
Net stock issues	−0.41**	−2.56	0.58	0.981
Treasury bill yield	−0.32**	−2.00	0.35	0.994
Long-term yield	−0.24	−1.54	0.21	0.997
Long-term return	0.26	1.64	0.24	0.049
Term spread	0.22	1.39	0.17	0.961
Default spread	0.11	0.71	0.04	0.975
Default return	0.25	1.57	0.22	−0.103
Inflation	−0.21	−1.30	0.15	0.482
Relative bill rate	−0.15	−0.94	0.08	0.741
Smoothed earnings-price ratio	0.35**	2.19	0.42	0.992
Small-stock value spread	−0.07	−0.42	0.02	0.984
Payout yield	0.45***	2.84	0.70	0.983
Average correlation	0.24	1.51	0.20	0.882
Nearness to Dow historical high	−0.08	−0.52	0.02	0.990
PLS aggregated book-to-market ratio	0.80***	5.10	2.23	0.968
Tail risk	0.12	0.73	0.05	0.777
Amihud illiquidity measure	0.16	1.04	0.09	0.946
Zero return frequency	0.27*	1.68	0.25	0.879
Average skewness	0.16	0.99	0.09	0.106

Table III
Model forecast combination univariate regressions.

The table summarizes results from univariate predictive regressions and autoregressions for the model forecast combination predictor variables for alternative values of k (i.e., the number of persistence-based subgroups) ranging from one to three. Each predictor variable is the fitted log excess market return forecast from the mean combination approach, and this predictor variable is standardized. We report the predictive regression coefficient, its OLS t-statistic (in parentheses), and the R^2 (in percent) from a regression of the log excess market return on the lagged predictor value. We also report the autoregressive coefficient from a regression of the predictor variable on its own lag. The sample period is January 1927 to December 2021. An * (**) [***] on the predictive coefficient estimate indicates statistical significance at the 10% (5%) [1%] level in a two-tailed test.

	$\tilde{r}^e_{1,t+1}$			$\tilde{r}^e_{2,t+1}$			$\tilde{r}^e_{3,t+1}$		
Specification	Coeff.	R^2	AR(1)	Coeff.	R^2	AR(1)	Coeff.	R^2	AR(1)
$k = 1$	0.96***	3.23	0.957						
	(6.16)								
$k = 2$	1.17***	4.76	0.880	0.42***	0.62	0.993			
	(7.54)			(2.67)					
$k = 3$	0.51***	0.90	0.389	1.07***	4.01	0.965	0.42***	0.62	0.993
	(3.20)			(6.89)			(2.67)		

Table IV
Vector autoregression for baseline model.

The table summarizes results from a vector autoregression with the model forecast combinations with $k = 3$ predictor subgroups based on persistence. The predictors are sorted into subgroups based on their autoregressive half-lives, and we create model forecast combinations for low-persistence ($\tilde{r}^e_{1,t+1}$, half-lives less than three months), medium-persistence ($\tilde{r}^e_{2,t+1}$, three months to four years), and high-persistence ($\tilde{r}^e_{3,t+1}$, more than four years) predictors. The state variables in the VAR are the log excess market return and the fitted log excess market return forecasts from the low-, medium-, and high-persistence subgroups using the mean combination approach, and the fitted expected returns are standardized. We report the VAR coefficient matrix A, its OLS t-statistics (in parentheses), the R^2 (in percent) from each regression, and the covariance matrix Σ. The sample period is January 1927 to December 2021. An * (**) [***] on a VAR coefficient estimate indicates statistical significance at the 10% (5%) [1%] level in a two-tailed test.

	A					Σ			
	r^e_t	$\tilde{r}^e_{1,t+1}$	$\tilde{r}^e_{2,t+1}$	$\tilde{r}^e_{3,t+1}$	R^2	r^e_{t+1}	$\tilde{r}^e_{1,t+2}$	$\tilde{r}^e_{2,t+2}$	$\tilde{r}^e_{3,t+2}$
r^e_{t+1}	0.04	0.38**	0.99***	0.08	4.93	27.27	1.86	−0.24	−0.55
	(1.36)	(2.28)	(5.89)	(0.49)					
$\tilde{r}^e_{1,t+2}$	−0.03***	0.45***	0.07**	−0.02	17.63	1.86	0.82	−0.02	−0.03
	(−5.54)	(15.36)	(2.53)	(−0.73)					
$\tilde{r}^e_{2,t+2}$	−0.01***	−0.01	0.97***	0.02***	93.76	−0.24	−0.02	0.06	0.01
	(−9.16)	(−0.97)	(119.79)	(2.69)					
$\tilde{r}^e_{3,t+2}$	−0.00	−0.00	−0.02***	1.00***	98.57	−0.55	−0.03	0.01	0.01
	(−1.32)	(−0.81)	(−4.45)	(260.79)					

Table V
Return and dividend growth forecasts from discount-rate and cash-flow news estimates.

The table summarizes results from regressions that forecast annual log market return (Panel A) and log aggregate dividend growth (Panel B) using objective discount-rate news and cash-flow news estimates from the baseline VAR. The predictors are smoothed versions of discount-rate news ($MA_t(N_{DR})$) and cash-flow news ($MA_t(N_{CF})$) estimates that use a trailing exponentially weighted moving average following Campbell, Giglio, and Polk (2013). The sample period is January 1927 to December 2021. Due to overlapping data in annual return and dividend growth computations, the t-statistics in parentheses are based on Newey-West (1987) standard errors with a lag length equal to 11. An * (**) [***] on a coefficient estimate indicates statistical significance at the 10% (5%) [1%] level in a two-tailed test.

Model	$MA_t(N_{DR})$	$MA_t(N_{CF})$	R^2
Panel A: Annual return forecasts			
Baseline VAR	10.64^{***}		20.22
	(3.64)		
	10.58^{***}	0.29	20.22
	(3.20)	(0.06)	
Panel B: Annual dividend growth forecasts			
Baseline VAR		13.78^{***}	18.13
		(5.01)	
	−0.31	14.18^{***}	18.17
	(−0.27)	(4.75)	

Table VI
Objective discount-rate and cash-flow news variability.

The table reports the standard deviations of objective discount-rate news and objective cash-flow news and the correlation between them. The standard deviations are reported in percentage per month. The sample period is January 1927 to December 2021.

Statistic	Estimate
$\sigma(N_{DR})$	5.90
$\sigma(N_{CF})$	2.84
$\rho(N_{DR}, N_{CF})$	0.46

Table VII
Return and dividend growth forecasts from discount-rate and cash-flow news estimates, 1985-2021.

The table summarizes results from regressions that forecast annual log market return (Panel A) and log aggregate dividend growth (Panel B) using discount-rate news and cash-flow news estimates from the baseline VAR and the Chen, Da, and Zhao (2013) method. The predictors for the baseline VAR are smoothed versions of discount-rate news ($MA_t(N_{DR})$) and cash-flow news ($MA_t(N_{CF})$) estimates that use a trailing exponentially weighted moving average following Campbell, Giglio, and Polk (2013). The predictors for the analyst forecast revisions model of Chen, Da, and Zhao (2013) are the one-year discount-rate and cash-flow news series. Because these series are quarterly, the analyst revision predictive regressions are quarterly. The sample period is October 1985 to December 2021. Due to overlapping data in annual return and dividend growth computations, the t-statistics in parentheses are based on Newey-West (1987) standard errors with a lag length equal to 11 for the VAR regressions and three for the analyst revision model. An * (**) [***] on a coefficient estimate indicates statistical significance at the 10% (5%) [1%] level in a two-tailed test.

Model	$MA_t(N_{DR})$	$MA_t(N_{CF})$	R^2
Panel A: Annual return forecasts			
Baseline VAR	10.43^{***}		10.63
	(2.87)		
	10.96^{***}	-7.01	12.57
	(2.79)	(-1.53)	
Analyst revisions	0.16^{*}		3.72
	(1.73)		
	0.03	-0.16	5.79
	(0.20)	(-1.10)	
Panel B: Annual dividend growth forecasts			
Baseline VAR		15.89^{***}	46.69
		(7.43)	
	2.25^{*}	14.75^{***}	49.39
	(1.94)	(6.74)	
Analyst revisions		0.10^{**}	7.18
		(2.50)	
	0.09	0.16^{*}	9.60
	(0.89)	(1.63)	

Table VIII
Objective and subjective discount-rate and cash-flow news variability.

The table reports the variability statistics for both objective and subjective decompositions. Panel A reports the standard deviations of discount-rate news and cash-flow news and the correlation between them for each decomposition. The standard deviations are reported in percentage per month. Panel B reports the correlations between subjective and objective news components. The subjective series are quarterly, so the statistics are based on quarterly observations and the quarterly standard deviations are reported in percentage per month by multiplying them by $1/\sqrt{3}$. The sample period is April 1985 to December 2021.

Panel A: News component moments		
Statistic	Objective	Subjective
$\sigma(N_{DR})$	4.72	6.15
$\sigma(N_{CF})$	3.33	4.75
$\rho(N_{DR}, N_{CF})$	0.31	0.69
Panel B: Objective-subjective correlations		
Statistic	N_{DR}	N_{CF}
ρ	0.46	0.21

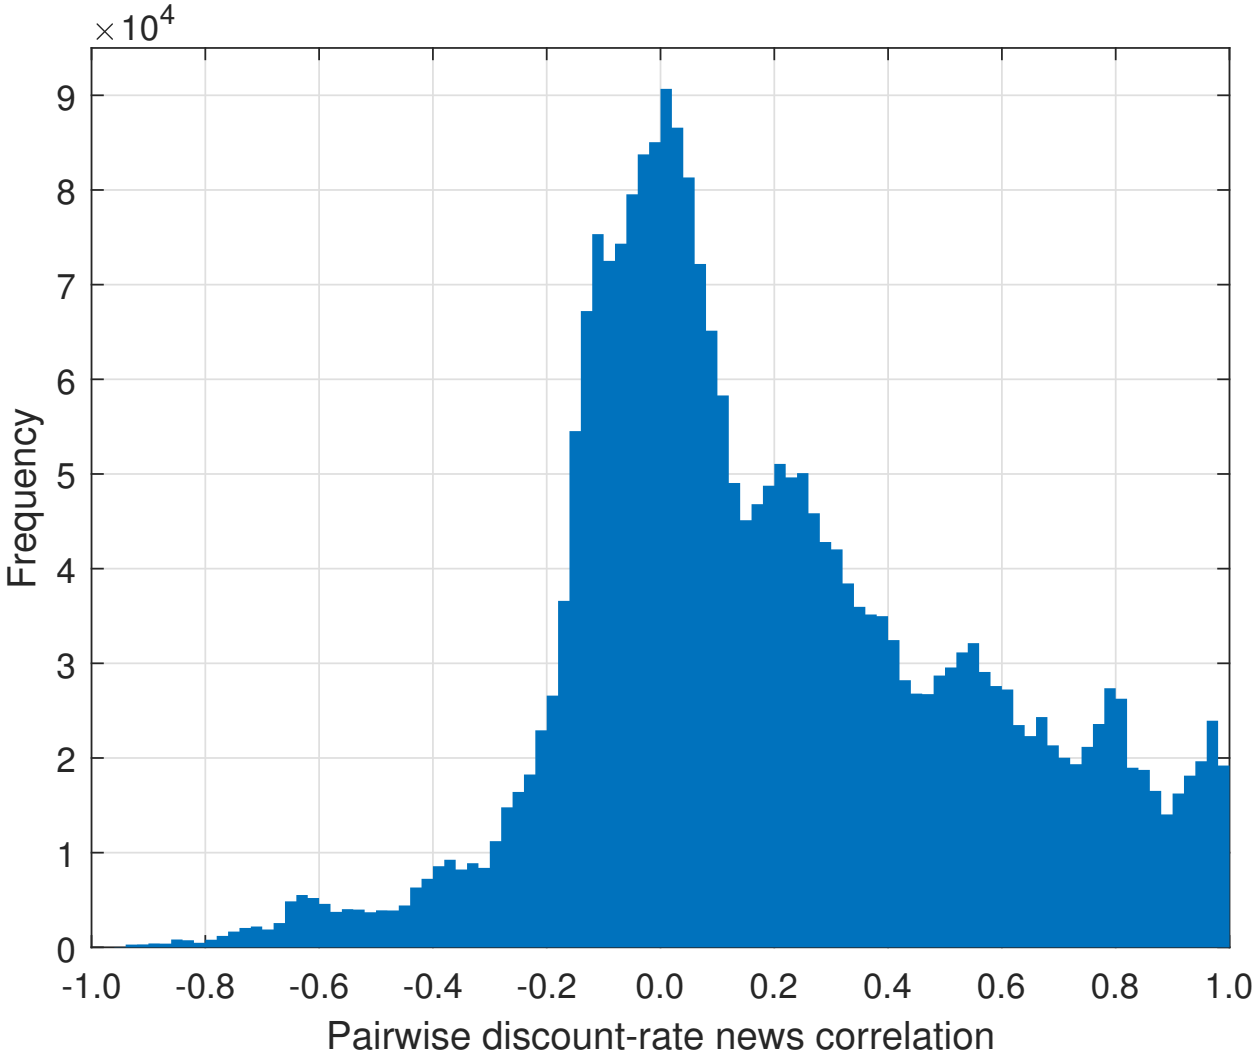


Figure 1. Correlation of discount-rate news across model pairs. The figure shows a histogram of the time-series correlation between the discount-rate news estimates across pairs of VAR models. The VAR models under consideration are all 2,300 possible combinations of the market excess return and three predictor variables from a set of 25 predictors. The sample period is January 1927 to December 2021.

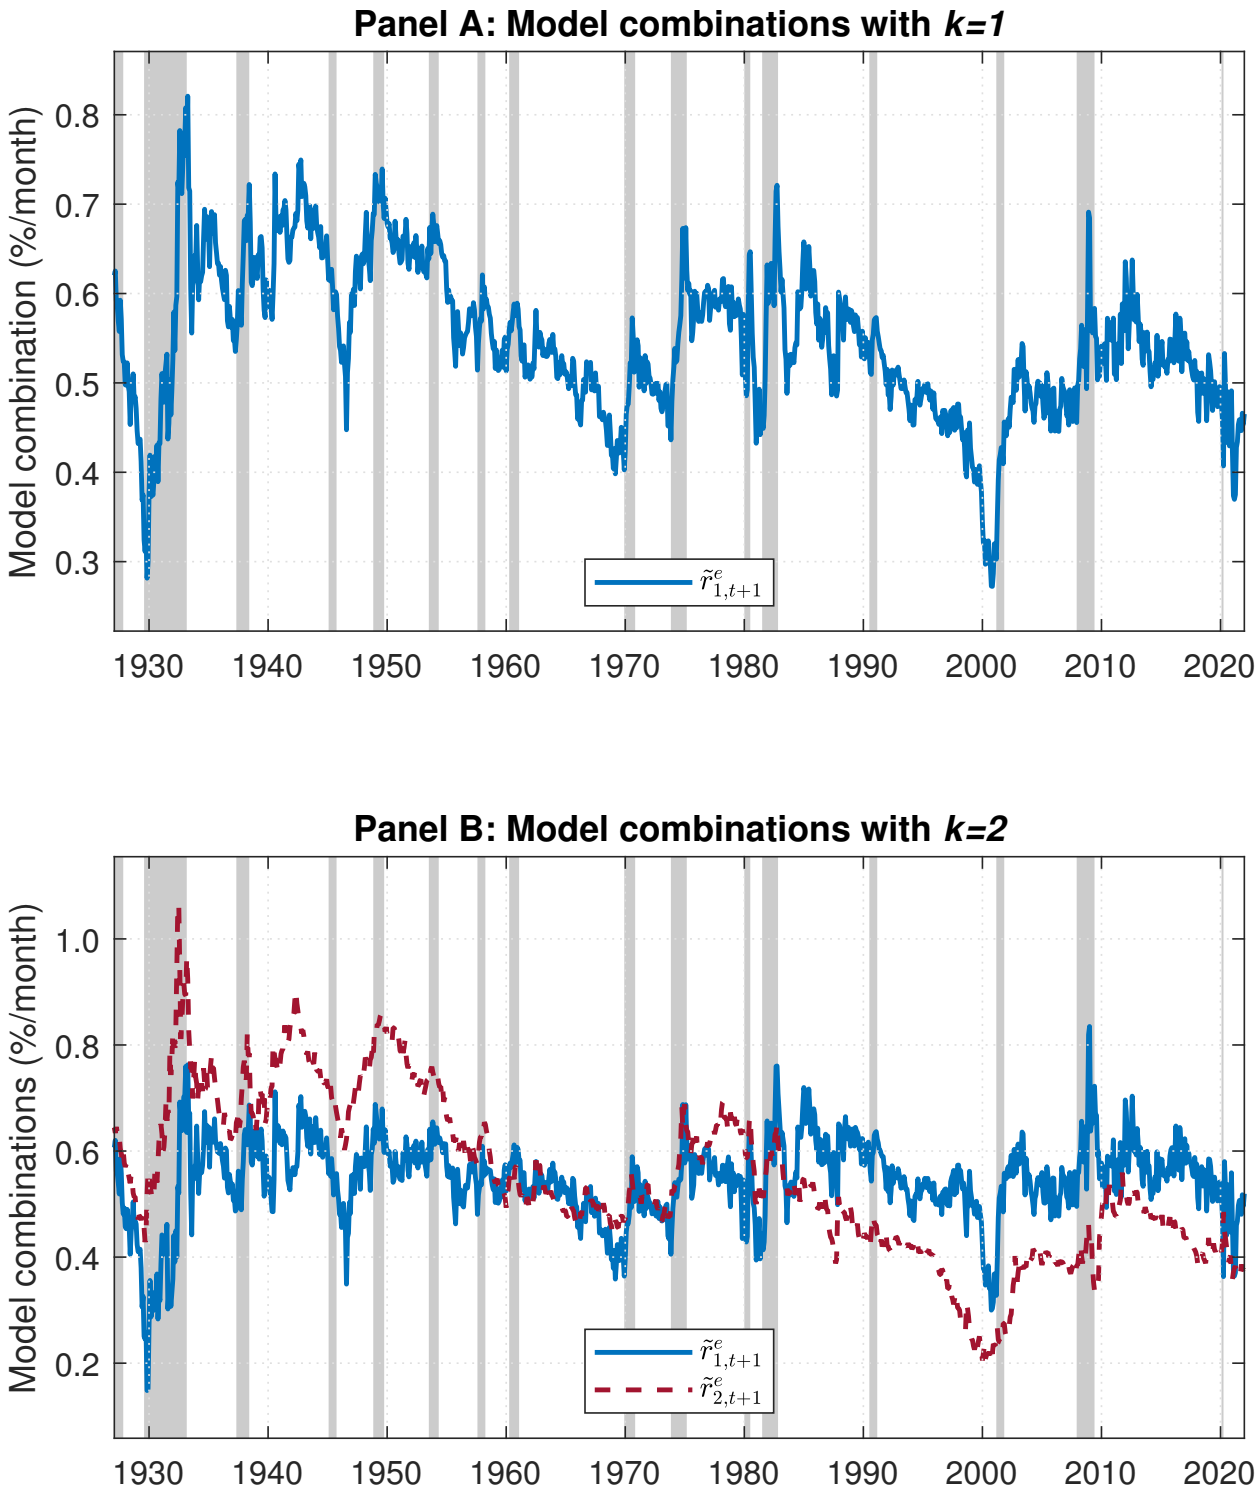
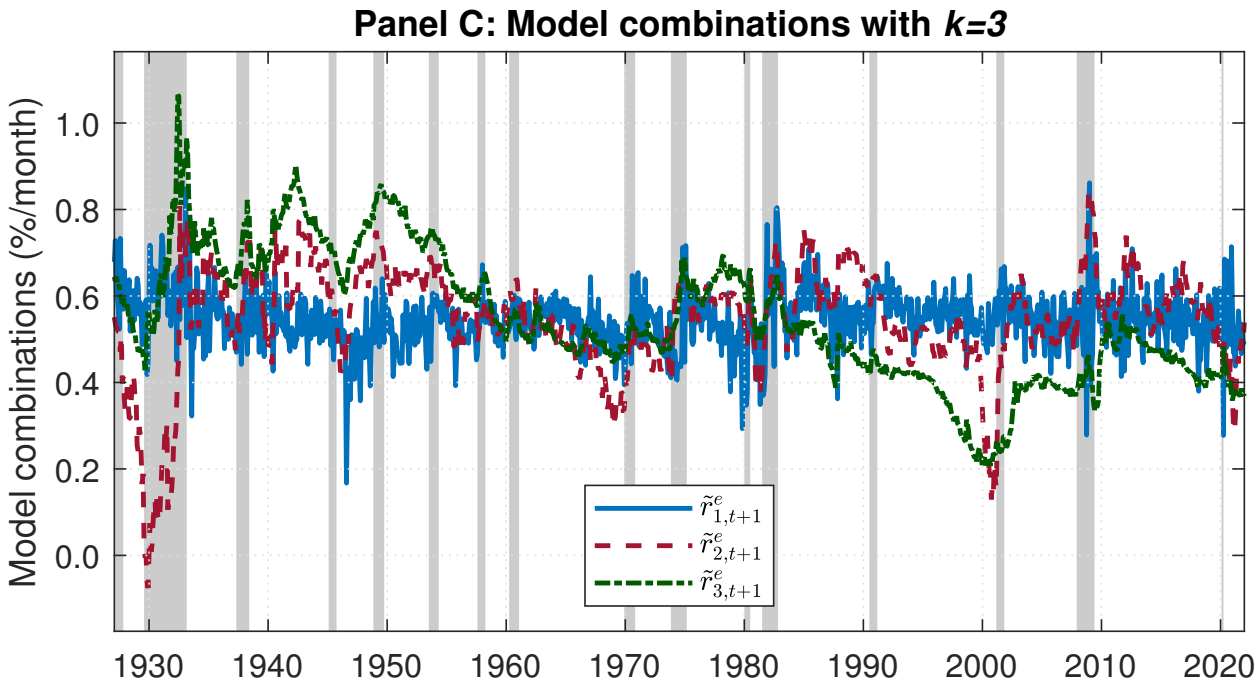


Figure 2. Model forecast combination predictor variables. The figure shows the time series of model forecast combination predictor variables for alternative values of k (i.e., the number of persistence-based subgroups) ranging from one (Panel A) to three (Panel C). Each predictor variable is the fitted log excess market return forecast from the mean combination approach. NBER recessions are shaded. The sample period is January 1927 to December 2021.


Panel C: Model combinations with k=3
Model combinations (%/month)
1.0
0.8
0.6
0.4
0.2
0.0
1930
1940
1950
1960
1970
1980
1990
2000
2010
2020
$\tilde{r}^e_{1,t+1}$
$\tilde{r}^e_{2,t+1}$
$\tilde{r}^e_{3,t+1}$

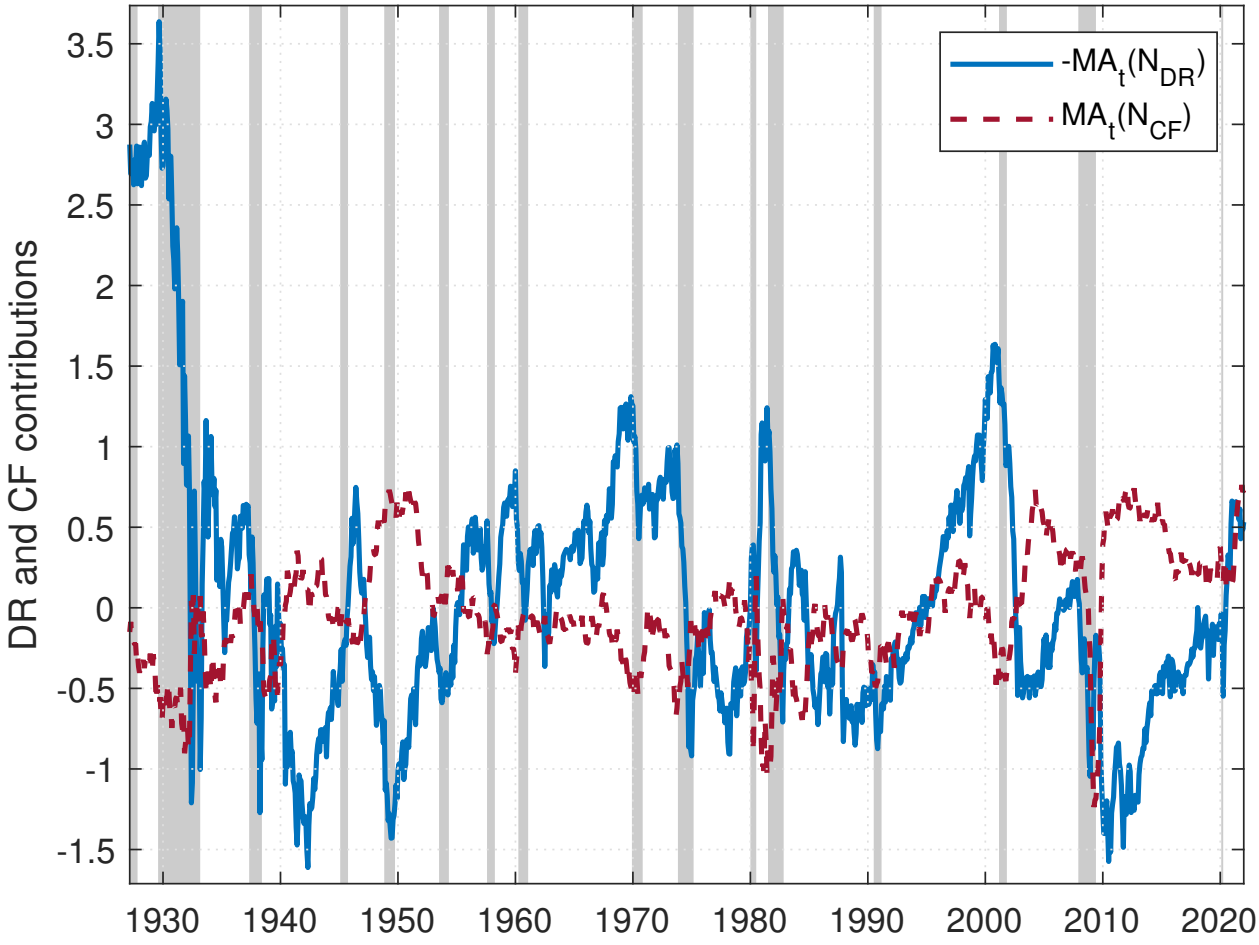


Figure 3. Smoothed discount-rate news and cash-flow news components. The figure shows the smoothed estimates from the baseline VAR of the contributions from discount-rate news and cash-flow news to market excess returns. The smoothed series use a trailing exponentially weighted moving average approach following Campbell, Giglio, and Polk (2013). A positive value of $-MA_t(N_{DR})$ (solid blue line) or $MA_t(N_{CF})$ (dotted red line) is indicative of a positive contribution to recent returns from discount-rate news or cash-flow news, respectively. NBER recessions are shaded. The sample period is January 1927 to December 2021.

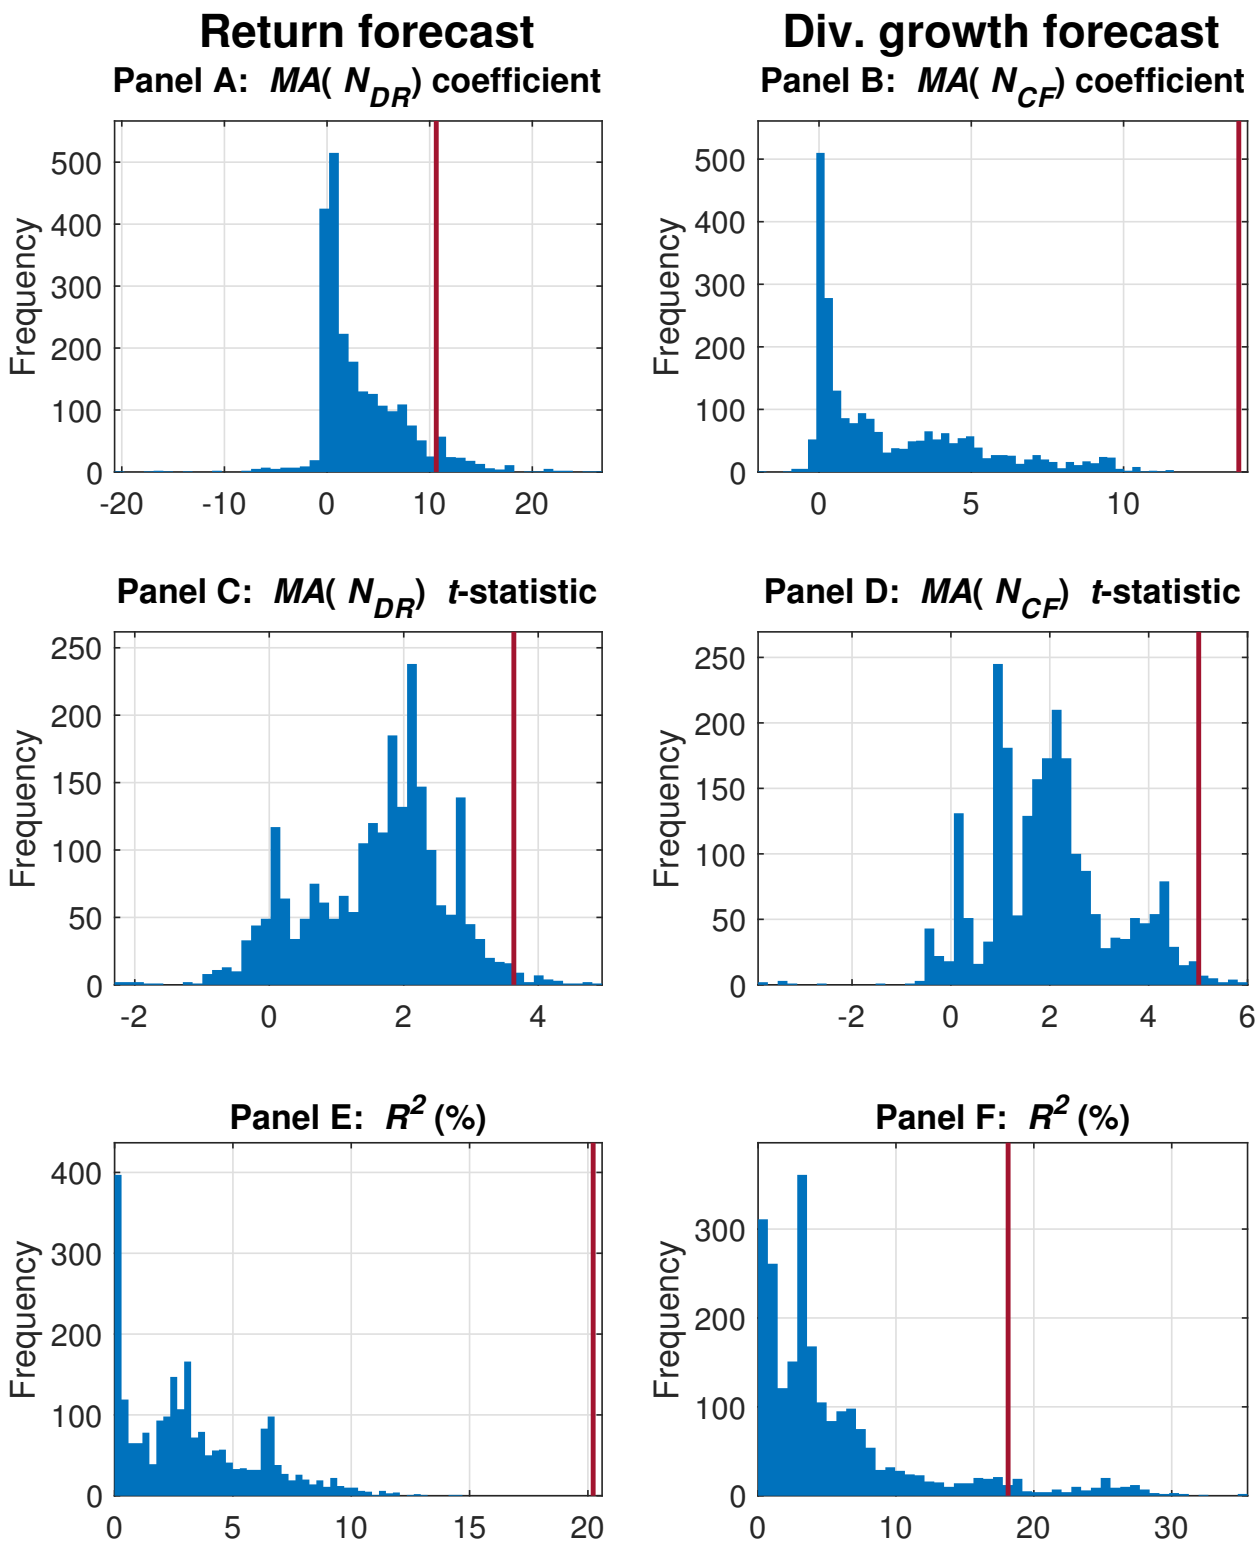


Figure 4. Return and dividend growth forecast regressions across models. The figure shows the coefficient estimates, t-statistics, and R^2s from predictive regressions for annual log excess market returns (Panels A, C, and E) and annual log dividend growth (Panels B, D, and F). The blue histograms show results for VARs with all 2,300 possible combinations of the market excess return and three predictor variables from a set of 25 predictors, and the red line shows results for our baseline VAR. The sample period is January 1927 to December 2021.

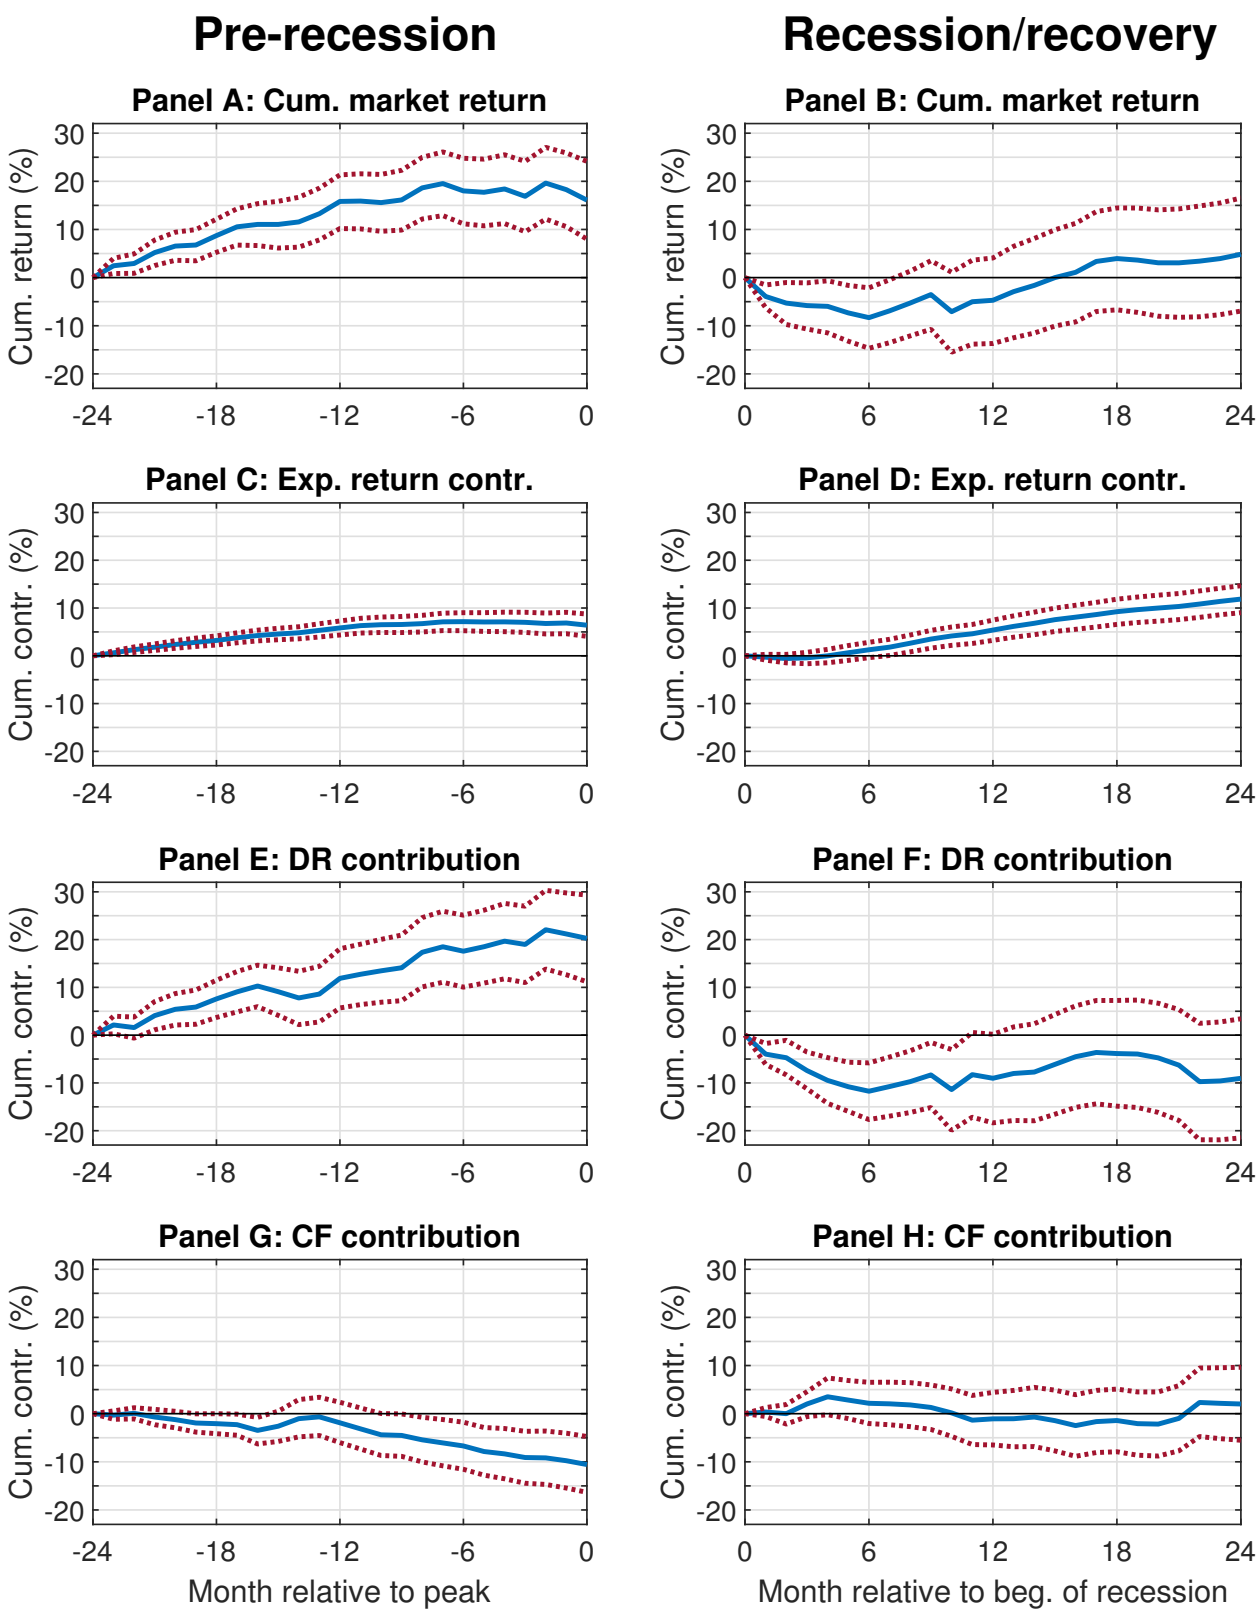


Figure 5. Baseline VAR objective news about returns, discount rates, and cash flows around recessions. The figure shows the cumulative market returns (Panels A and B), expected returns (Panels C and D), negative of discount-rate news (Panels E and F), and cash-flow news (Panels G and H) in the periods leading up to economic peaks (Panels A, C, E, and G) and following NBER recession starts (Panels B, D, F, and H). The 90% confidence intervals from a bootstrap simulation are shown as red dotted lines. The sample period is January 1927 to December 2021.

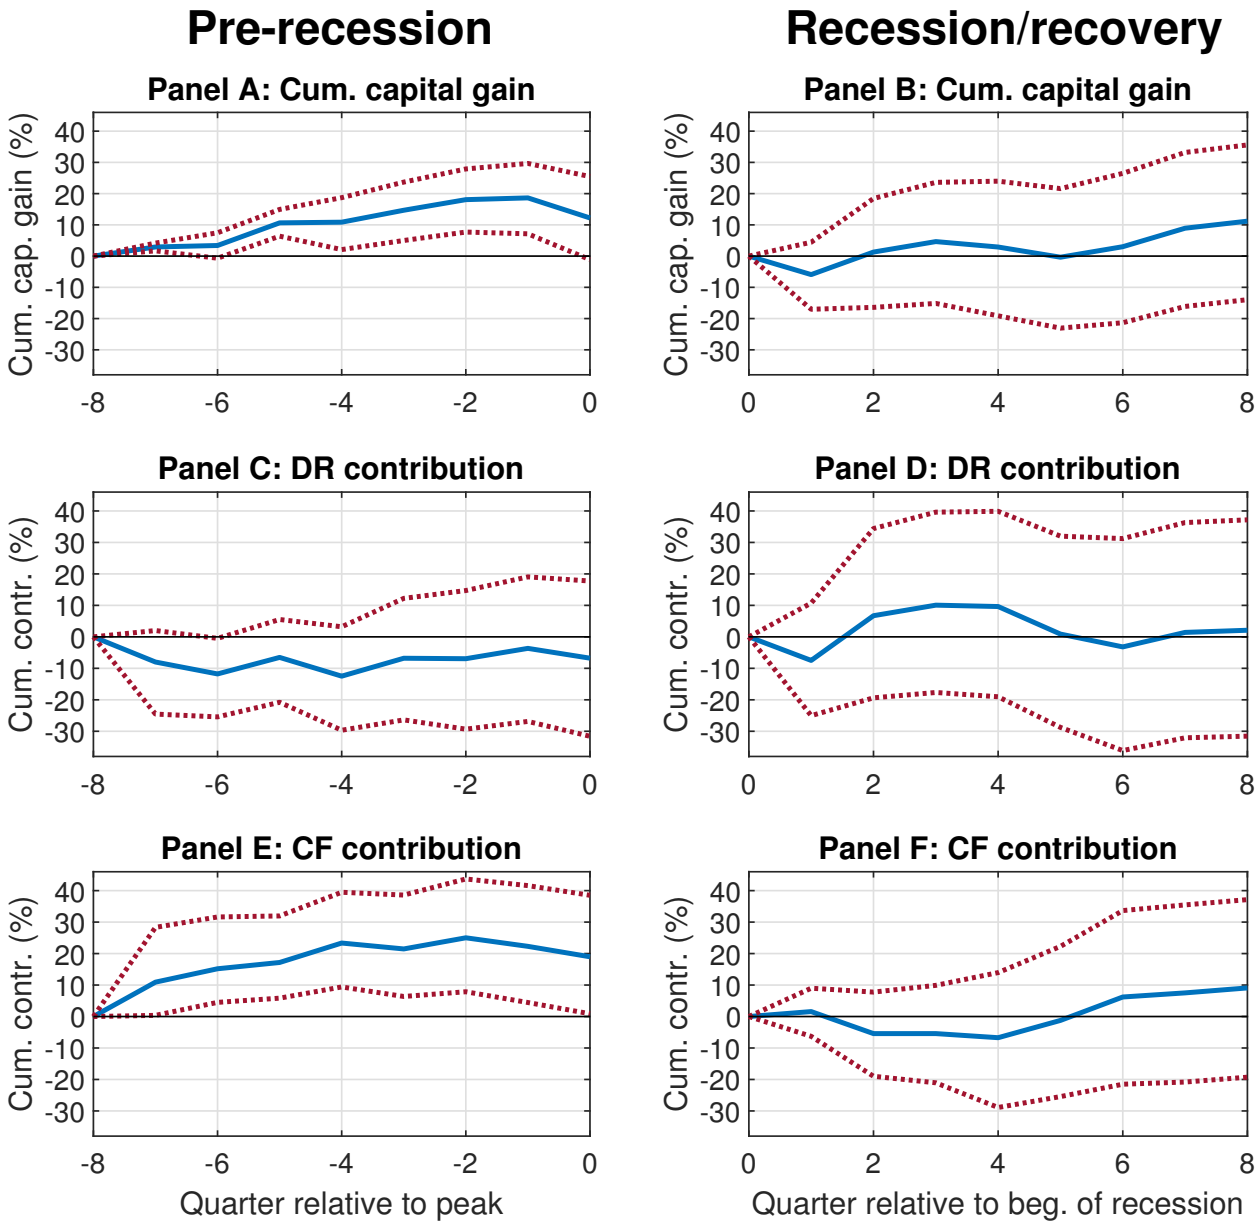


Figure 6. Subjective news about capital gains, discount rates, and cash flows around recessions. The figure shows the cumulative market capital gains (Panels A and B), implied costs of capital (Panels C and D), discount-rate news (Panels E and F), and cash-flow news (Panels G and H) from the Chen, Da, and Zhao (2013) method in the periods leading up to economic peaks (Panels A, C, E, and G) and following NBER recession starts (Panels B, D, F, and H). The 90% confidence intervals from a bootstrap simulation are shown as red dotted lines. The sample period is April 1985 to December 2021.

Internet Appendix

"The Stock Market's Two Truths: Subjective Beliefs and Objective Reality"

A Variable descriptions

This appendix describes the return predictor variables used in the vector autoregressions. We require predictors to have historical data by January 1927. We estimate VARs using monthly data, so we also require that we can calculate each variable on a monthly basis.

We consider a total of 25 predictors. The first 14 predictors are from Goyal and Welch (2008).

1. **Dividend-price ratio**: The log of the ratio of dividends paid on the S&P 500 over the past 12 months and the index level of the S&P 500.
2. **Dividend yield**: The log of the ratio of dividends paid on the S&P 500 over the past 12 months and the 12-month lag of the index level of the S&P 500.
3. **Earnings-price ratio**: The log of the ratio of earnings on the S&P 500 over the past 12 months and the index level of the S&P 500.
4. **Book-to-market ratio**: The ratio of book value and market value of the Dow Jones Industrial Average.
5. **Dividend-earnings ratio**: The log of the ratio of dividends on the S&P 500 over the past 12 months and earnings on the S&P 500 over the past 12 months.
6. **Stock variance**: The sum of squared daily returns on the S&P 500 over the past month.
7. **Net stock issues**: The ratio of the sum of net issues by NYSE stocks over the past 12 months and the total market capitalization of NYSE stocks.
8. **Treasury bill yield**: The yield on three-month Treasury bills.
9. **Long-term yield**: The yield on long-term Treasury bonds.
10. **Long-term return**: The return on long-term Treasury bonds over the past month.
11. **Term spread**: The difference between the yield on long-term Treasury bonds and the yield on three-month Treasury bills.
12. **Default spread**: The difference between the yield on BAA-rated corporate bonds and the yield on AAA-rated corporate bonds.
13. **Default return**: The difference between the return on long-term corporate bonds and the return on long-term Treasury bonds over the past month.
14. **Inflation**: The one-month percentage change in the Consumer Price Index (All Urban Consumers) lagged by one month to account for reporting delays.

The next three predictors are variables used in VARs by John Campbell that are not analogous to any of the variables listed above.

15. **Relative bill rate**: Following Campbell (1991), we compute the difference between the yield on one-month Treasury bills and the average yield on one-month Treasury bills over the prior 12 months. We thank Kenneth French for providing data for one-month Treasury bill yields.

16. **Smoothed earnings-price ratio**: Following Campbell and Vuolteenaho (2004), the smoothed earnings-price ratio is the log of the ratio of the average annual earnings on the S&P 500 over the past ten years and the index level of the S&P 500. We thank Amit Goyal for providing these data.

17. **Small-stock value spread**: Following Campbell and Vuolteenaho (2004), the small-stock value spread is the log of the ratio of the book-to-market ratio of the small value portfolio and the book-to-market ratio of the small growth portfolio, where the small value and small growth portfolios are from a two-by-three sort on size and book-to-market.[19] The June market-to-book ratio for each portfolio is the value-weighted average of the book value from fiscal year $t-1$ divided by the market value from December of year $t-1$. The log difference in the book-to-market ratios in subsequent months is adjusted by adding the log of the cumulative value-weighted return on the small growth portfolio since the past June and subtracting the log of the cumulative value-weighted return on the small value portfolio since the past June.

The remaining eight predictors are from recent studies published in the *Journal of Finance*, *Journal of Financial Economics*, or *Review of Financial Studies*.

18. **Net payout yield**: Following Boudoukh, Michaely, Richardson, and Roberts (2007), we compute the net payout yield as the log of 0.1 plus the ratio of the aggregate net payout by nonfinancial firms over the past 12 months and the aggregate market capitalization of nonfinancial firms. We use data from CRSP for NYSE-, AMEX-, and Nasdaq-listed firms with share codes of 10 and 11, and we exclude firms with Standard Industrial Classification (SIC) codes between 6000 and 6999. The aggregate dividend for each month is calculated from firm-level dividends that imputed from the returns and ex-dividend returns on nonfinancial firms. The aggregate net issuance for each month is calculated from firm-level net issuance, which is the product of the split-adjusted price per share and the monthly change in the split-adjusted number of shares outstanding, in which the split-adjusted price per share is the month-beginning and month-ending average. Aggregate net payout is aggregate dividends minus aggregate net issuances. We calculate $PY_t = \log(0.1 + NP_t)$, where NP_t is the ratio of the aggregate net payout by nonfinancial firms over the past 12 months and the aggregate market capitalization of nonfinancial firms.

19. **Average correlation**: Following Pollet and Wilson (2010), we calculate the value-weighted average cross-sectional correlation of daily returns for the 500 largest stocks by market capitalization and the end of each month. We use daily returns from CRSP for NYSE-, AMEX-, and Nasdaq-listed firms with share codes of 10 and 11. For each month t, we compute the standard deviation for each stock i, $\hat{\sigma}_{i,t}$, and the covariance between each pair of stocks i and j, $\hat{\sigma}_{ij,t}$, over the period from month $t-2$ to month t using non-missing daily returns. We require that each firm exists for all three months by ensuring it has a price observation at the end of month $t-3$. We then compute the correlation between each pair of stocks i and j as

$$\hat{\rho} = \frac{\hat{\sigma}_{ij,t}}{\hat{\sigma}_{i,t}\hat{\sigma}_{j,t}}. \tag{A1}$$

[19] We thank Kenneth French for making these data available (https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html).

Finally, we calculate the average correlation measure as

$$CORR_t = \sum_{i=1}^{500} \sum_{j \neq i} w_{i,t} w_{j,t} \hat{\rho}_{ij,t}, \tag{A2}$$

where $w_{i,t}$ and $w_{j,t}$ are the value weights of stocks i and j at the end of month t in a portfolio of the largest 500 stocks at the end of month t.

20. **Nearness to Dow historical high**: Following Li and Yu (2012), we calculate the ratio of the Dow Jones Industrial Index at the end of each month t and the historical high of daily closing index values of the Dow Jones Industrial Index. Data are from Bloomberg.

21. **Partial least squares aggregated book-to-market ratio**: Following Kelly and Pruitt (2013), we aggregate information from the book-to-market ratios of portfolios sorted on size and book-to-market using a partial least squares (PLS) approach. We use data on returns and book-to-market ratios of 100 portfolios sorted on size and book-to-market and data on returns of the CRSP value-weighted stock market index from Kenneth French's website.[20] The June market-to-book ratio for each portfolio is the value-weighted average of the book value from fiscal year $t-1$ divided by the market value from June of year t. The log book-to-market ratios in subsequent months for each portfolio are adjusted by subtracting the log of the cumulative value-weighted return on the portfolio since the past June. In the first step of the PLS approach, we estimate the time-series regression

$$\log(BM_{i,t}) = \phi_{i,0} + \phi_{i,1} \log R_{t+1} + e_{i,t}, \tag{A3}$$

for each portfolio i, where $BM_{i,t}$ is the book-to-market of portfolio i in month t and R_{t+1} is the stock market return in month $t+1$. In the second step, we estimate the cross-sectional regression

$$\log(BM_{i,t}) = c_t + ABM_t \hat{\phi}_{i,1} + w_{i,t}, \tag{A4}$$

for each month t. The ABM_t measure is the estimated latent factor in the cross-sectional regressions.[21] We estimate the model using data from January 1927 to January 2022 to obtain a time series of the predictor that spans January 1927 to December 2021.

22. **Tail risk**: Following Kelly and Jiang (2014), we estimate tail risk using Hill's (1975) power law estimator. We use daily returns from CRSP for NYSE-, AMEX-, and Nasdaq-listed firms with share codes of 10 and 11. The tail risk measure in month t is calculated as

$$TAIL_t = \frac{1}{K_t} \sum_{k=1}^{K_t} \log \left(\frac{R_{k,t}}{u_t} \right), \tag{A5}$$

where $R_{k,t}$ is the kth daily return in the pooled cross section in month t that falls below the extreme value threshold (u_t) and K_t is the total number of returns below the threshold in month t. The threshold parameter is the fifth percentile of the pooled cross-sectional distribution of daily returns in month t. We standardize the measure to have mean of zero and standard deviation of one.

23. **Amihud illiquidity measure**: Following Chen, Eaton, and Paye (2018), we measure the

[20] We thank Kenneth French for making these data available (https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html).

[21] We thank Seth Pruitt for making code available (https://sethpruitt.net/research/).

log of the equal-weighted average of the Amihud (2002) illiquidity measure with adjustments for breaks and market volatility. We use daily returns from CRSP for NYSE-listed stocks with share codes of 10, 11, and 12. For each firm i in each month t, the Amihud illiquidity measure is calculated as

$$AMI_{i,t} = \left(\frac{1}{D_{i,t}} \sum_{d=1}^{D_{i,t}} \frac{|R_{i,d,t}|}{DVOL_{i,d,t}} \right) CPI_t, \tag{A6}$$

where $D_{i,t}$ is the number of trading days with positive volume for stock i in month t, $|R_{i,d,t}|$ is the absolute value of the return on stock i on day d of month t, $DVOL_{i,d,t}$ is the dollar trading volume in millions for stock i on day d, and CPI_t is the Consumer Price Index (CPI). CPI data are from the Federal Reserve Bank of St. Louis.[22] We use the real-time break adjustment procedure from Chen, Eaton, and Paye (2018) to adjust the time series for the effects of changes in tick sizes in 1997 and 2001. Finally, we subtract the volatility component, $vol_t = \overline{\log(\hat{\sigma}_{i,t})}$.

24. **Zero return measure**: Following Chen, Eaton, and Paye (2018), we measure the log of the equal-weighted average of the proportion of zero-return days with an adjustment for breaks. We use daily returns from CRSP for NYSE-listed stocks with share codes of 10, 11, and 12. For each firm i in each month t, we calculate the proportion of trading days with a return of zero. We use the real-time break adjustment procedure from Chen, Eaton, and Paye (2018) to adjust the time series for the effects of changes in tick sizes in 1997 and 2001.

25. **Average skewness**: Following Jondeau, Zhang, and Zhu (2019), we measure the value-weighted average of the skewness of excess daily returns within each month. We use daily returns from CRSP for NYSE-, AMEX-, and Nasdaq-listed firms with share codes of 10 and 11. We exclude stocks with an Amihud (2002) illiquidity measure in the top 0.1% or with prices below $1. For a given month, we use all stocks that have at least ten valid return observations.

B Additional results

This appendix contains additional empirical results referenced in the paper.

B.1 Out-of-sample predictability

Table B.I shows results for out-of-sample return predictability using the model forecast combinations. Inclusive of training period, the tests use the sample period of January 1927 to December 2021. The out-of-sample strategies use expanding-window estimation with an initial training period of 120, 180, 240, or 300 months as denoted in the table. The first performance measure is the out-of-sample R^2 statistic of Goyal and Welch (2008), which compares the mean squared error (MSE) of the forecast errors from the regression model with the MSE of the forecast errors from the historical mean. A positive value of the out-of-sample R^2 indicates that the regression model generates small forecast errors in comparison with the historical mean. Statistical significance of the out-of-sample R^2 is based on the mean squared prediction error-adjusted (MSPE-A) method of Clark and West (2007). The second performance measure is the out-of-sample Sharpe ratio of a trading strategy that scales the market return by the return forecast, following Kelly, Malamud, and Zhou (2024). Statistical significance of the difference between the out-of-sample Sharpe ratio and the Sharpe

[22] See https://fred.stlouisfed.org/series/CPIAUCNS.

ratio of a buy-and-hold market strategy is based on the test statistic of Jobson and Korkie (1981) with the correction noted by Memmel (2003).

B.2 Predictor persistence and horizon effects

Our baseline specification has three persistence-based subgroups of predictors to model horizon-specific expected returns. To illustrate horizon effects, Figure B.1 examines horizon-specific contributions to discount-rate news for the models with k ranging from one to three. Panel A shows the proportion of discount-rate news variance that is attributable to variation in expected returns at each horizon from one to 120 months. Specifically, given equation (13), discount-rate news variance is

$$Var(N_{DR,t+1}) = e1'\rho A(I - \rho A)^{-1}\Sigma(I - \rho A)^{-1\prime}A'\rho e1. \tag{B7}$$

We decompose this expression by horizon as

$$Var(N_{DR,t+1}) = \sum_{j=1}^{\infty} e1'\rho^j A^j \Sigma(I - \rho A)^{-1\prime}A'\rho e1, \tag{B8}$$

and we plot the jth term of the sum scaled by $Var(N_{DR,t+1})$ to show the proportion of discount-rate news variance by horizon. Panel B shows the cumulative sum of equation (B8) for horizons of one to 120 months, and the diamond on the right y-axis denotes the total discount-rate news variance from the infinite sum for each model.

With the exception of some short-term dynamics related to the lagged market return predictor, Figure B.1 shows that the single forecast combination model (i.e., $k = 1$) displays something very close to exponential decay for discount-rate effects. These simple dynamics combined with the relatively short half-life of the single predictor (about 16 months) imply that the discount-rate news component from this VAR places heavy emphasis on short-horizon effects. Panel B confirms this intuition. Given the total discount-rate news variance of $31.84(\%^2)$ for this model (denoted by the blue diamond on the right axis), 45% of variance is from discount-rate shocks at horizons of one year or less, 85% is within three years, and almost 100% is within ten years.

It is evident in Figure B.1 that the $k = 2$ model has both short-term and long-term decay, allowing for richer dynamics compared with the $k = 1$ model. The model with k equal to three has a richer pattern still. In untabulated results, we find that including an additional predictor for a $k = 4$ model adds little to our base case of $k = 3$. The figure shows that, of the total discount-rate news variance of $34.84(\%^2)$ for the base model, 26% is from shocks at horizons of one year or less, 61% is within three years, and 93% is within ten years.

B.3 Sensitivity to the number of predictor variables

Figure B.2 demonstrates the stability of discount-rate news estimates across models with a specific number of predictor variables included in each of the three model forecast combinations. The figure shows the histogram of time-series correlations of the discount-rate news estimate from each possible model with the discount-rate news estimate from our baseline model from Section 4.1. Panel A shows results for all models in which each model forecast combination includes a single predictor variable (i.e., a case that is equivalent to a traditional VAR with three predictor variables, but constrained to include one each from each persistence subgroup). Panels B to E examine true model forecast combination models with multiple variables, and the panels incrementally increase the number of included variables from two to five per combination. As such, the total number of predictors used ranges from three in Panel A to 15 in Panel E out of the 25 total return predictors.

Relative to the bleak picture painted by Figure 1 for traditional VARs, Figure B.2 provides encouraging support for our approach. Panel A, which uses only one predictor per persistence group, does not gain the benefits of model forecast combinations, and many of the models have low or negative correlation with our baseline model. Including two or three variables in each forecast combination begins to improve the outlook, as shown in Panels B and C.

When four or five variables are included in each model combination (i.e., Panels D and E of Figure B.2), the VAR models consistently generate discount-rate news estimates that have large, positive correlations with our baseline approach that uses all 25 variables. Including four variables in each of the three persistence-based groups uses fewer than half of the predictor variables (i.e., 12 of 25), but this approach has stable performance across variable choices. The median correlation with our baseline model is 0.84, and 90% of models have correlations above 0.61 with our base approach.

B.4 Sensitivity to the persistence cutoffs

Figure B.3 demonstrates the stability of discount-rate news estimates across models with different persistence cutoffs for the three model forecast combinations. The figure shows the histogram of time-series correlations of the discount-rate news estimate from each possible model with the discount-rate news estimate from our baseline model from Section 4.1. Panel A shows results for all models in which the low-persistence threshold (which is three months in the base case) varies between one month and 36 months, while maintaining a high-persistence threshold of 48 months. Panel B maintains the low-persistence threshold at three months and varies the high-persistence threshold (which is 48 months in the base case) between 12 months and 120 months. The figure considers the set of all possible models that move predictors across the persistence groups one at a time based on changing the persistence threshold. Figure B.3 shows that the discount-rate news estimates are insensitive to the persistence thresholds, as all correlations are near one.

Table B.I
Out-of-sample strategy performance.

The table shows out-of-sample R^2 values and strategy Sharpe ratios for the mean combination approach with alternative initial training windows. The out-of-sample (OOS) R^2 is computed following Goyal and Welch (2008). The OOS strategy returns use the fitted return prediction as the state variable for market timing. The sample period is January 1927 to December 2021. Statistical significance for the OOS R^2 uses the mean squared prediction error-adjusted (MSPE-A) method of Clark and West (2007). Statistical significance for the difference between the Sharpe ratios of the OOS strategy and the market uses the Jobson and Korkie (1981) test statistic with the correction from Memmel (2003). An * (**) [***] indicates statistical significance at the 10% (5%) [1%] level.

Training period	OOS R^2	OOS strategy SR	Market SR
120 months	0.41^{**}	0.58^{*}	0.51
180 months	0.64^{***}	0.67^{***}	0.60
240 months	0.49^{***}	0.64^{***}	0.55
300 months	0.38^{**}	0.59^{**}	0.52

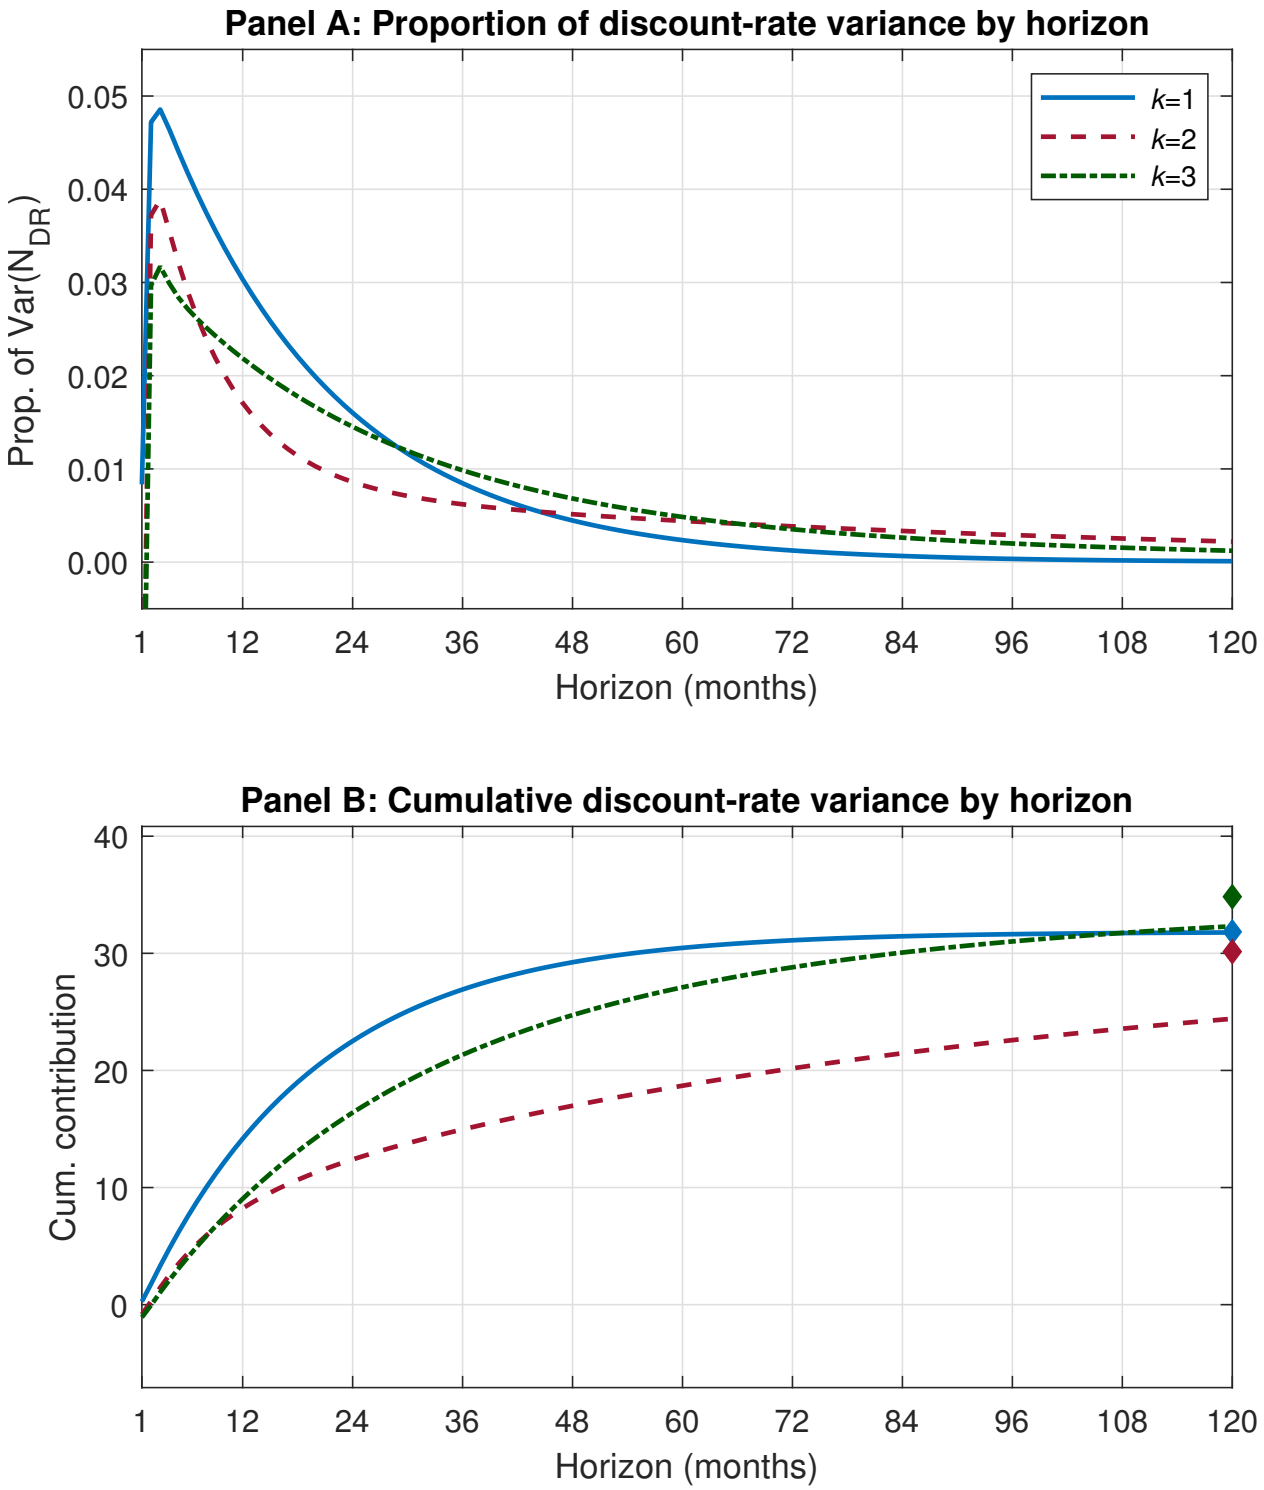


Figure B.1. Discount-rate news variance and horizon for VARs with model combination state variables. The figure summarizes horizon effects on discount-rate news variance for vector autoregressions with the model forecast combinations for alternative values of k (i.e., the number of persistence-based subgroups) ranging from one to three. Panel A shows the proportion of discount-rate news variance that is attributable to each horizon and Panel B shows the cumulative discount-rate news variance across horizons. The diamonds on the right axis in Panel B show total discount-rate news variance. The sample period is January 1927 to December 2021.

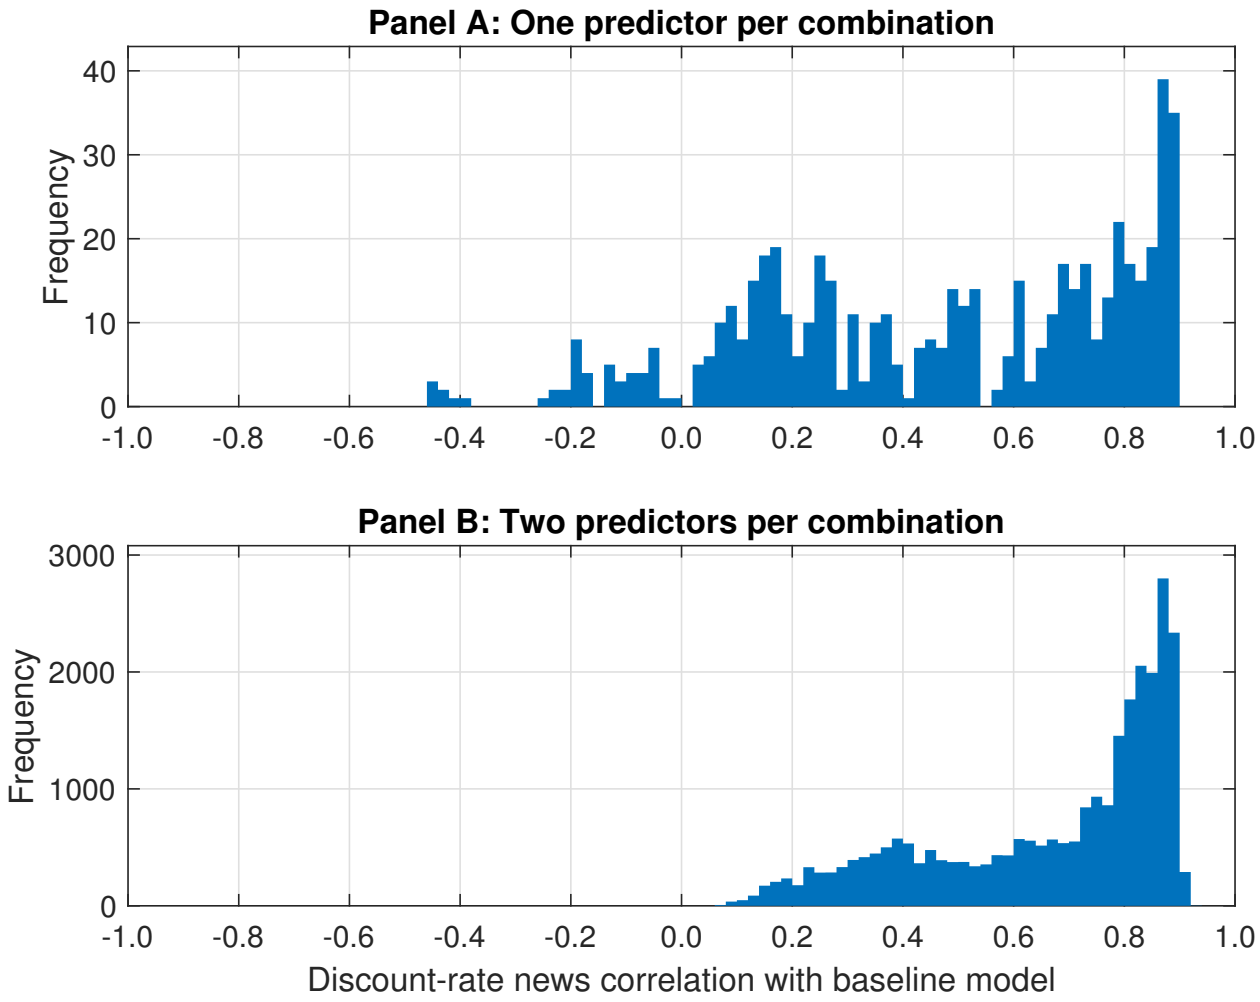
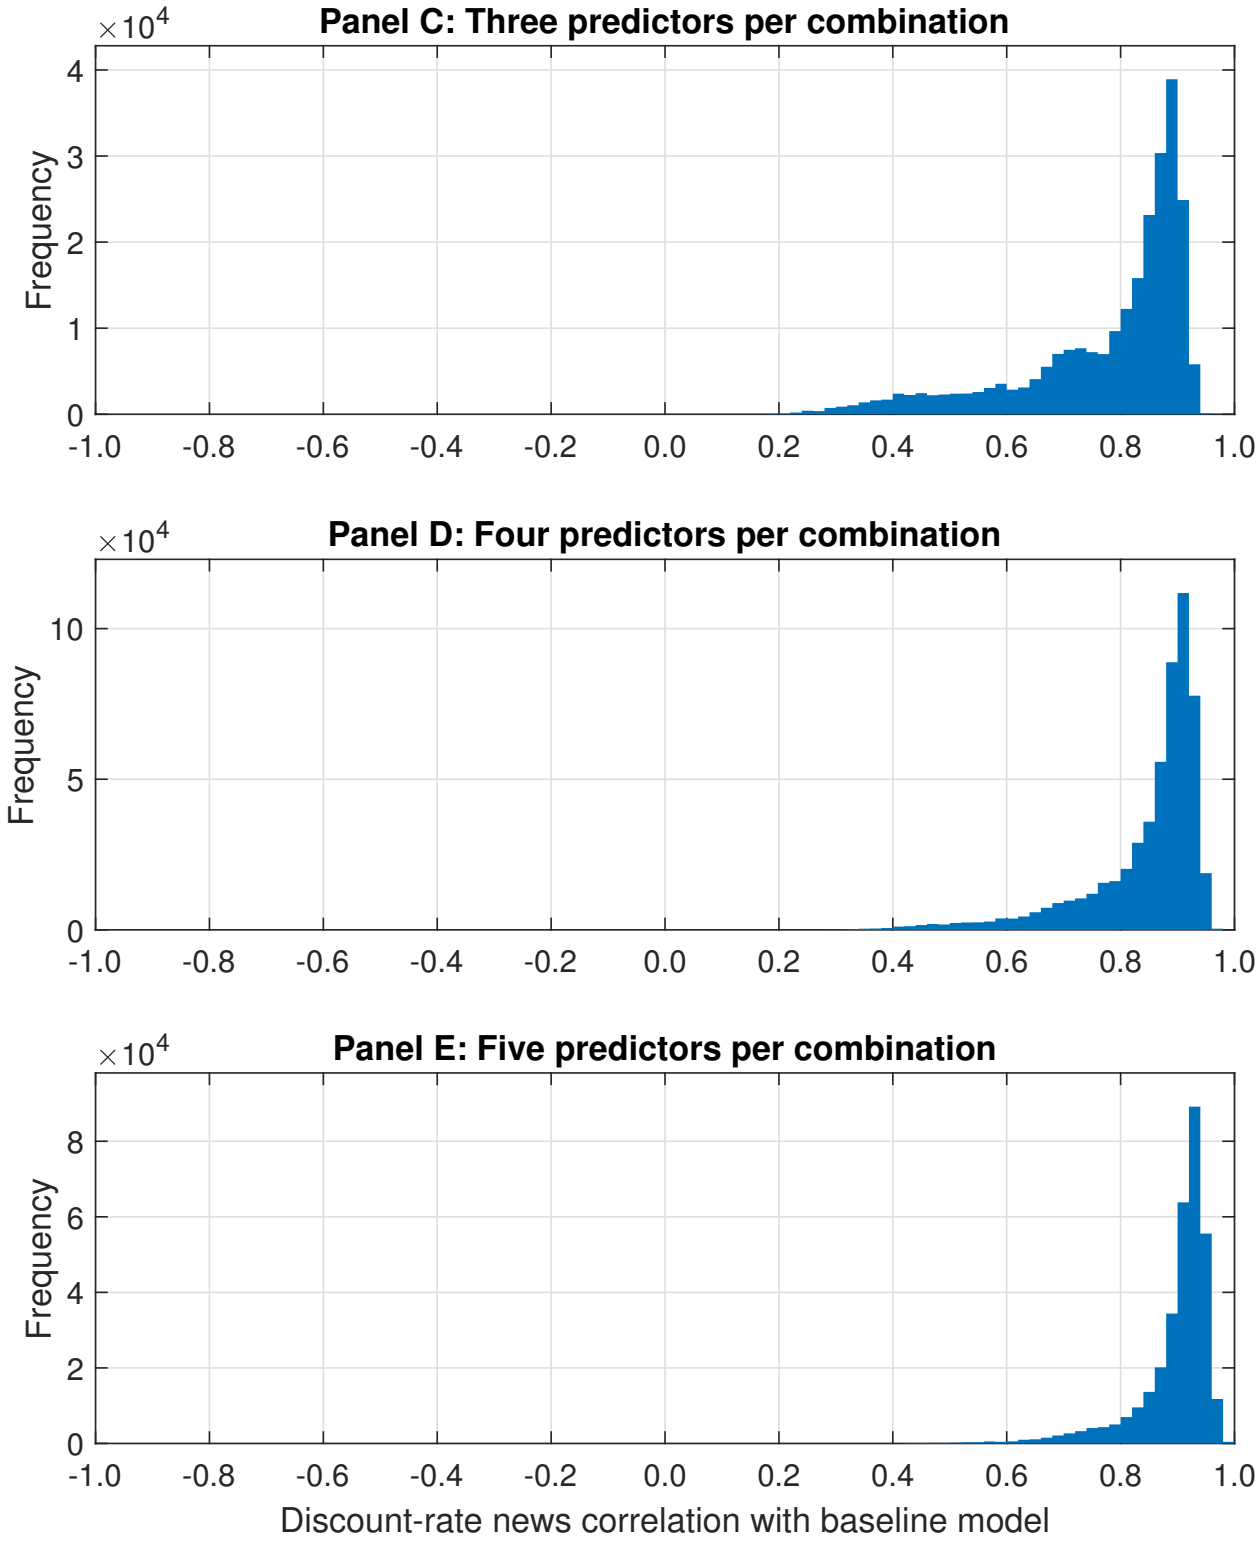


Figure B.2. Correlation of discount-rate news with our baseline model across predictor group sizes. The figure shows a histogram of the time-series correlation between the discount-rate news estimates of our baseline model and all possible models that include the given number of predictor variables in each of the persistence-based model forecast combinations. The VAR models under consideration are all possible combinations of the market return and three persistence-based predictor variables. Panel A shows results for one predictor variable per persistence-based subgroup, Panel B for two, Panel C for three, Panel D for four, and Panel E for five. The sample period is January 1927 to December 2021.


Panel C: Three predictors per combination
×10^4
Frequency
0
1
2
3
4
-1.0
-0.8
-0.6
-0.4
-0.2
0.0
0.2
0.4
0.6
0.8
1.0
Panel D: Four predictors per combination
×10^4
Frequency
0
5
10
Panel E: Five predictors per combination
×10^4
Frequency
0
2
4
6
8
Discount-rate news correlation with baseline model

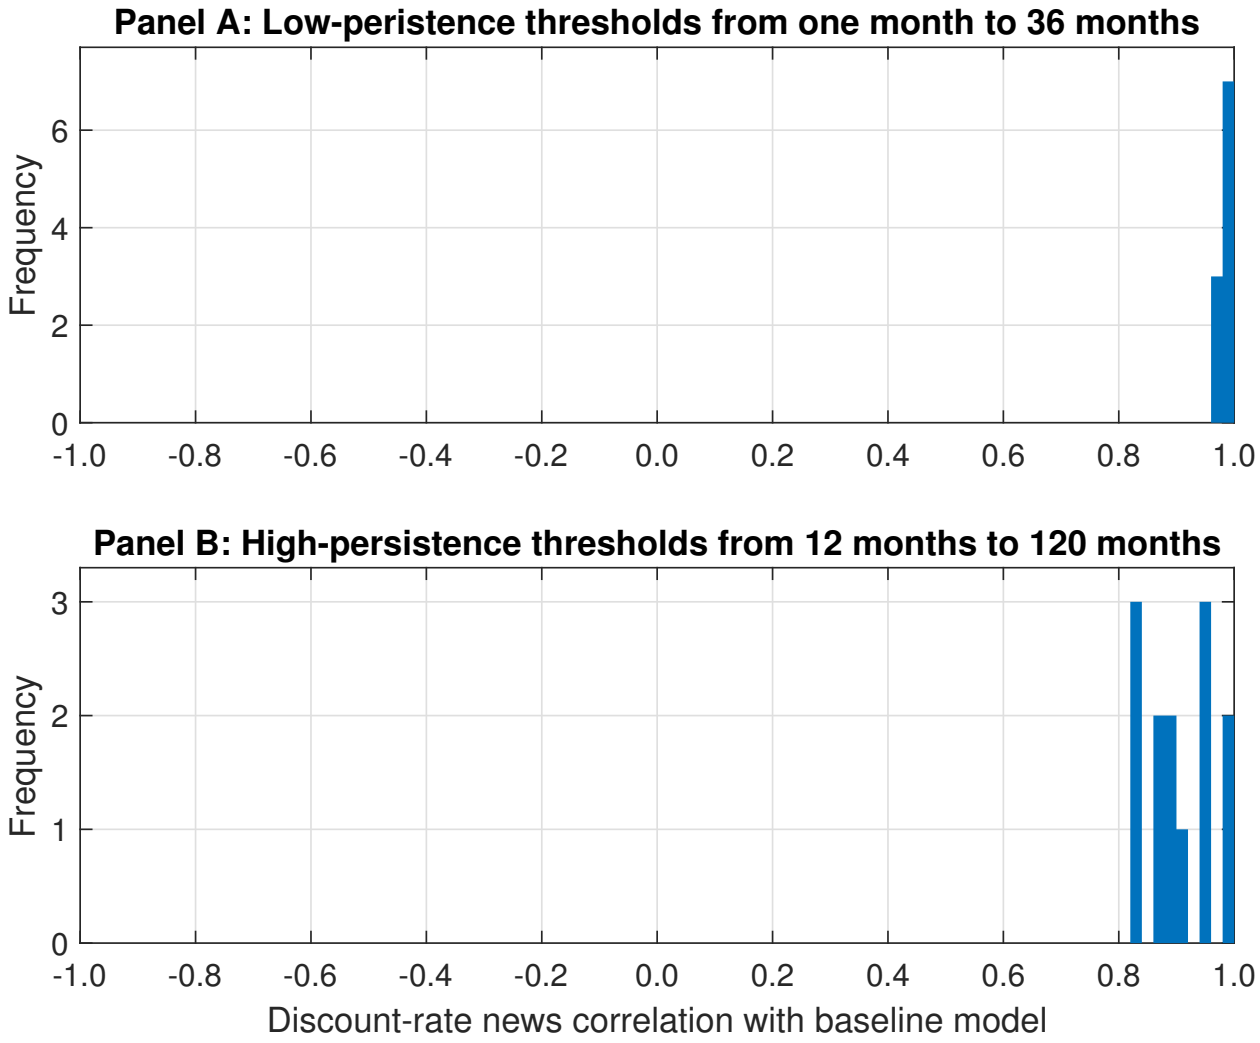


Figure B.3. Correlation of discount-rate news with our baseline model across persistence thresholds. The figure shows a histogram of the time-series correlation between the discount-rate news estimates of our baseline model and models with alternative thresholds for the persistence-based subgroups. The VAR models under consideration are all possible combinations of the market return and three persistence-based predictor variables. Panel A shows results for all models with a persistence threshold of 48 months for the high-persistence subgroup and a persistence threshold ranging from one month to 36 months for the low-persistence subgroup. Panel B shows results for all models with a persistence threshold of three months for the low-persistence subgroup and a persistence threshold ranging from 12 months to 120 months for the high-persistence subgroup. The sample period is January 1927 to December 2021.