

The Subjective Belief Factor

TINGYUE CUI, RICARDO DE LA O, and SEAN MYERS*

August 2024

ABSTRACT

Subjective expectations and asset prices both revolve around distorted probabilities. Subjective expectations are expectations under biased probabilities, and asset prices are expectations under risk-neutral probabilities. Given this link, asset pricing techniques designed to estimate a Stochastic Discount Factor (SDF) can be used to estimate a Subjective Belief Factor (SBF) – a distortion that characterizes many subjective expectations, even for non-financial variables. Conversely, the Subjective Belief Factor can be used to characterize asset prices, by separating the roles of beliefs and preferences/risk. Using the Survey of Professional Forecasters and Blue Chip, we find that differences between subjective expectations and statistical expectations for 24 macroeconomic variables can be summarized (average R^2 of 50%) by a single SBF related to real GDP growth and the T-bill rate. The results are broadly consistent with extrapolation. Applying this SBF to cross-sectional stock returns, we find it accounts for the majority of excess returns for the Fama-French factors and explains about half the variation in returns across 176 anomalies, while the remaining half is attributed to preferences/risk.

*Tingyue Cui and Sean Myers are with the Wharton School at the University of Pennsylvania. Ricardo De la O is with the Marshall School of Business, University of Southern California. Email addresses: ritacui@wharton.upenn.edu, rdelao@marshall.usc.edu, semyers@wharton.upenn.edu.

Behavioral economics research studying subjective expectations and finance research studying asset prices are solving similar problems. Asset prices can be represented using a distorted probability distribution, specifically a risk-neutral probability distribution, and much of asset pricing research revolves around the equation

$$P_t = E_t [M_{t+1} X_{t+1}] \quad (1)$$

where X_{t+1} is a payoff. Similarly, under general conditions, subjective expectations can be represented as expectations under a distorted probability distribution that differs from the objective distribution due to learning or behavioral biases,

$$E_t^* [X_{t+1}] = E_t [S_{t+1} X_{t+1}]. \quad (2)$$

Because of this similarity, we propose that asset pricing techniques can open novel possibilities for studying subjective expectations, even for *non-financial* variables X_{t+1} . First, rather than studying subjective expectations for each variable individually, we can study a single variable (S_{t+1}) that explains subjective expectations for many variables simultaneously. We refer to this variable as the Subjective Belief Factor (SBF). This is analogous to asset pricing work studying the M_{t+1} that links prices across many assets. Second, we can summarize S_{t+1} using only a small number of subjective expectations, akin to factors in asset pricing. For example, an agent that overstates the probability of low real GDP growth states will also overstate the probability of high unemployment states, as these states overlap, meaning that the researcher does not need to model an additional unemployment-specific bias. Third, by using log-affine functional forms, we can express both the SBF S_{t+1} and subjective expectations in a tractable manner that can be easily related to models of expectation formation, even for a large set of subjective expectations.

Conversely, we argue that framing subjective expectations using a single SBF is beneficial for asset pricing research. Using only information on prices, researchers cannot distinguish if asset price anomalies reflect biased beliefs or preferences/risk. Direct data on subjective expectations can help to resolve this. However, incorporating potentially biased beliefs about

dozens of variables is typically infeasible (e.g., attempting to model overreaction to real GDP, underreaction to inflation, extrapolation of yields, learning about consumption growth, ...). Because of this, asset pricing work is often limited to only using subjective expectations for a small number of variables. By summarizing subjective expectations for many variables with a single SBF, S_{t+1} , we can easily combine S_{t+1} with a distortion exclusively based on preferences/risk, denoted \tilde{M}_{t+1} , to price assets and distinguish the roles of these two variables.¹

In this paper, we provide a theoretical foundation linking subjective expectations, asset pricing techniques, and excess returns. Applying our approach to the Survey of Professional Forecasters, we find that the difference between statistical expectations and subjective expectations for 15 macroeconomic variables can be largely summarized by a single SBF S_{t+1} based on subjective expectations of just two variables: real GDP growth and the T-bill rate. This SBF is broadly consistent with extrapolation. Using this SBF, we construct synthetic expectations $E_t[S_{t+1}Y_{t+1}]$ for an additional 9 macroeconomic variables Y_{t+1} , and then verify the accuracy of these synthetic expectations using additional survey data from Blue Chip. Further, this SBF accounts for the majority of the excess returns for the Fama-French cross-sectional anomalies and accounts for roughly half the variation in the excess returns for 176 anomalies sorted into 22 categories from Chen and Zimmermann (2022).

To formalize our analysis, we first establish the conditions under which subjective expectations for many variables can be characterized by a single SBF. The key condition is the subjective expectations are “coherent”, meaning that they satisfy basic rules regarding addition and multiplication. We argue that this is a reasonable starting condition, given that it is satisfied by many models of expectation formation.² Importantly, agents can disagree and there can be a different SBF for each agent, however, there still exists a single consensus SBF S_{t+1} which explains the consensus forecast for all variables. If variables are

¹Using log-affine functional forms for S_{t+1} and \tilde{M}_{t+1} ensures that their product, $M_{t+1} \equiv S_{t+1}\tilde{M}_{t+1}$, is also log-affine and tractable.

²As a non-exhaustive list, adaptive expectations, sticky expectations, extrapolation, parameter learning, noisy information, and rational inattention can all be represented by a SBF.

normally or log-normally distributed, we show that the SBF has a straightforward log-affine representation and that excess returns on financial assets can be easily decomposed into their comovement with S_{t+1} and their comovement with \tilde{M}_{t+1} . Intuitively, an asset earns high excess returns if it pays off in states of the world that the agent thinks are unlikely or in states where the agent has a low preference for payoffs.

We then evaluate the efficacy of applying asset pricing tools to subjective expectations using the Survey of Professional Forecasters, one of the most commonly used sources of macroeconomic forecasts.³ We study annual forecasts for 15 variables, the maximum amount available, covering a wide range of topics, such as unemployment, housing starts, inflation, and government expenditures. We find that subjective expectations for all 15 variables, as well as the difference between subjective expectations and statistical expectations, can be largely explained by an SBF S_{t+1} based on real GDP growth and the T-bill rate. Specifically, synthetic expectations $E_t[S_{t+1}X_{t+1}]$ for all 15 variables match the actual subjective expectations with an average R^2 of 64.6% and match the difference between subjective and statistical expectations with an average R^2 of 47.2%. This result means that, given an SBF that explains subjective expectations of real GDP growth and the T-bill rate, we can explain subjective expectations for the remaining 13 variables based on the objective covariance of these variables with real GDP growth and the T-bill rate.

Quantitatively, the explanatory power of the estimated SBF is comparable to the upper bound from principal component analysis (PCA). However, unlike PCA, the estimated SBF allows us to extend our results to other variables without needing any additional survey data. We consider nine additional macroeconomic variables, such as the mortgage rate and the 10-year Treasury rate, and calculate synthetic expectations based on our estimated SBF and the objective covariance of these variables with real GDP growth and the T-bill rate. Importantly, we do not use any survey data for these nine variables when constructing the synthetic expectations. We then verify the accuracy of these synthetic expectations by

³In this survey, we can reasonably expect forecasts to be “coherent” since the vast majority of participants use mathematical models to form their expectations.

comparing them to the subjective expectations measured from Blue Chip. Once again, the synthetic expectations closely match the actual subjective expectations and also match the difference between subjective and statistical expectations, with average R^2 's similar to the Survey of Professional Forecasters results.

What do these results mean for models of expectation formation? First, our results indicate that differences between subjective expectations and statistical expectations for many variables can largely be condensed to beliefs about a few key variables, in this case real GDP growth and the T-bill rate. While all 24 subjective expectations differ from their corresponding statistical expectations, researchers do not need 24 biases. This raises the prospect that models of inflation expectations (e.g., Malmendier and Nagel, 2016), consumption expectations (e.g., Collin-Dufresne, Johannes, and Lochstoer, 2016), risk-free rate expectations (e.g., Haubrich, Pennacchi, and Ritchken, 2012), etc. can be represented as different manifestations of one underlying belief distortion. Second, the SBF nests many existing models of expectation formation and the log-affine representation for the SBF makes it straightforward to compare different models. Each model can be summarized by the coefficients it implies for the log-affine formula. For illustration, we consider diagnostic expectations and extrapolation and demonstrate that our estimated coefficients for the SBF are relatively closer to the predictions of extrapolation.

In our last set of tests, we emphasize the benefits of our approach for asset pricing researchers. Intuitively, once subjective expectations are framed using asset pricing tools, they are easy to integrate into asset pricing research. Rather than attempting to incorporate individual biases for each variable, we can focus on a single SBF (S_{t+1}) that summarizes the subjective expectations for many variables and can be easily combined with a distortion related to preferences (\tilde{M}_{t+1}). In particular, if all variables are log-normal, then the historical average excess return \bar{R}_{t+1}^e for an asset can be decomposed into

$$\log(\bar{R}_{t+1}^e) = Cov(-s_{t+1}, r_{t+1}^e) + Cov(-\tilde{m}_{t+1}, r_{t+1}^e) \quad (3)$$

where lowercase denotes log values.

Equation (3) is closely related to asset pricing work relating excess returns to measures of sentiment (Baker and Wurgler, 2006; Huang et al., 2015; Stambaugh, Yu, and Yuan, 2024) or to subjective expectations of returns (Engelberg, McLean, and Pontiff, 2020; Jensen, 2023; De la O, Han, and Myers, 2024) and, in a sense, captures the best of both worlds. Like measures of sentiment, this decomposition relates excess returns for many variables to a single belief-related variable s_{t+1} . Like research on subjective expected returns, this decomposition is quantitative, as it utilizes covariances rather than correlations. Because the decomposition is quantitative, it addresses the common issue of risk and sentiment being correlated.⁴ At the same time, by linking excess returns for many assets to a single SBF, this decomposition removes the need to have subjective return expectations for each individual asset, which relaxes several constraints related to data availability.⁵

Using our SBF estimated based on subjective expectations of real GDP growth and the T-bill rate, we show that the excess returns for the Fama-French size, value, investment, and profitability factors appear to be largely accounted for by their comovement with $-s_{t+1}$. Further, we connect our SBF to the two behavioral factor portfolios of Daniel, Hirshleifer, and Sun (2020), who construct one factor sorting firms by earnings surprises to capture short-term behavioral biases (e.g., weekly-level biases) and a separate factor using share issuances to capture longer term behavioral biases. Given that our SBF is based on one-year subjective expectations, we reassuringly find that our SBF accounts for almost none of the excess return for their short-horizon factor and nearly all (93%) of the excess return for their long-horizon factor. Finally, we study a large set of 176 anomalies sorted into 22 categories

⁴If a researcher has a measure of sentiment that is correlated with future returns, there is always the concern that this measure may simply be correlated with risk. In equation (3), the roles of s_{t+1} and \tilde{m}_{t+1} can be distinguished even if they are correlated. If both s_{t+1} and \tilde{m}_{t+1} negatively comove with the excess return, then $\frac{Cov(-s_{t+1}, r_{t+1}^e)}{\log(\tilde{R}_{t+1}^e)} \in (0, 1)$ will tell us what portion is attributable to s_{t+1} and what portion is attributable to \tilde{m}_{t+1} .

⁵Data on subjective return expectations for the long and short legs of anomalies is often limited to post-2000 and comes almost exclusively from sell-side analyst forecasts, which may be prone to incentive issues such as setting their forecasts to justify current prices.

from Chen and Zimmermann (2022). We find that the estimated SBF accounts for 56.1% of the differences in excess returns across anomalies, while the remaining 43.9% is attributed to the preference-based \tilde{M}_{t+1} . Thus, while the SBF S_{t+1} certainly does not explain all excess returns, we find that it appears similar in importance to \tilde{M}_{t+1} , with roughly a 50-50 split when we consider many anomalies.

Broadly, this paper contributes to and attempts to link two literatures. The first is the literature on subjective expectations. Given the size of this literature, we refer readers to the handbook by Bachmann et al. (2023) for an overview and list of references. To a large extent, these papers focus on expectations for a single variable, such as inflation, output, interest rates, exports, housing, unemployment, etc.⁶ Other papers, such as Coibion and Gorodnichenko (2015) and Bordalo et al. (2020), propose general tests that can be applied to many variables but, importantly, are designed to be applied to each variable separately (e.g., testing for overreaction or underreaction for each individual variable). We build on this work by presenting an approach that allows us to study subjective expectations for many variables jointly and to condense these subjective expectations down to a single SBF based on only a few variables.⁷ Our emphasis on jointly analyzing multiple macroeconomic variables is related to recent work using statistical and machine learning forecasts for large sets of variables (Bianchi, Ludvigson, and Ma, 2024).

Second, there is the asset pricing literature emphasizing that the M_{t+1} that prices assets may contain a belief-based component and a preference-based component, e.g. S_{t+1} and \tilde{M}_{t+1} . The idea of framing subjective expectations as a distortion has been discussed previously. For example, Piazzesi, Salomao, and Schneider (2015) use a probability distortion to characterize subjective expectations related to bond yields. We formalize the idea of representing subjective expectations as a distortion, demonstrating the conditions necessary for S_{t+1} to exist and addressing issues of aggregation and disagreement. Further, we

⁶For some recent examples, see Cieslak (2018), Kuchler and Zafar (2019) and Mueller et al. (2021).

⁷Importantly, we do not claim to be reinventing the wheel. These tools have long-existed in asset pricing. Our goal is to demonstrate their effectiveness in understanding subjective expectations.

demonstrate that this SBF can be used not only to condense existing subjective expectations data but also to form synthetic expectations for other variables. We also emphasize that this approach can be applied outside the context of finance and asset pricing. Even for a researcher purely interested in understanding subjective expectations of the components of output (consumption, investment, government expenditures, and net exports), the tools developed in asset pricing are still relevant.

Given that belief distortions operate in a very similar way to preference-based distortions, Brav and Heaton (2002) and Adam and Nagel (2023) emphasize the difficulty or impossibility of distinguishing S_{t+1} and \tilde{M}_{t+1} solely from asset prices. We provide a method to estimate S_{t+1} from subjective expectations and demonstrate how this can be used to decompose observed excess returns into a belief-based component and a preference/risk-based component. This is closely related to recent work by Chen, Hansen, and Hansen (2020) that uses asset prices to establish bounds on a belief distortion under assumptions regarding relative entropy. While they do not use survey data in their empirical application, their method can incorporate survey data as additional moment conditions. Additionally, our approach is related to Kozak, Nagel, and Santosh (2018) who argue that structural models with specific assumptions about beliefs and preferences can be used to separate S_{t+1} and \tilde{M}_{t+1} using asset price data. We provide a complementary method that extracts the SBF solely from survey data and then evaluates its ability to explain asset prices.

The rest of the paper is organized as follows. Section I establishes the conditions under which subjective expectations for many variables can be described by a single Subjective Belief Factor and demonstrates its key properties. Section II discusses the data on subjective expectations taken from the Survey of Professional Forecasters and Blue Chip. Section III applies asset pricing techniques to understand subjective expectations for financial and non-financial macroeconomic variables. Section IV uses the SBF estimated from subjective expectations to explain excess returns for a wide range of anomalies. Section V concludes.

I. Existence of the Subjective Belief Factor

In this section, we formalize how subjective expectations for multiple variables can be described by a single subjective belief factor. In terms of notation, $E_{i,t}^*[\cdot]$ denotes the subjective expectations of agent i at time t . All other operators use the objective probability distribution. For example, $E_t[\cdot]$ and $Cov(\cdot, \cdot)$ denote the conditional expectation and the unconditional covariance under the objective probability distribution.

Throughout the paper, we assume that subjective expectations are “coherent,” meaning that they satisfy two conditions. First, for any variables Y_{t+1} , Z_{t+1} and constants a, b , $E_{i,t}^*[aY_{t+1} + bZ_{t+1}]$ equals $aE_{i,t}^*[Y_{t+1}] + bE_{i,t}^*[Z_{t+1}]$. Second, $E_{i,t}^*[1]$ equals 1. These conditions are all that is necessary to show that subjective expectations can be represented by a subjective belief factor.

Proposition 1. *There exists $S_{i,t+1}$ such that $E_{i,t}^*[X_{t+1}] = E_t[S_{i,t+1}X_{t+1}]$ for all variables X_{t+1} and $E_t[S_{i,t+1}] = 1$.*

All proofs are provided in Appendix A. To an econometrician that knows the objective probability distribution, the agent’s subjective expectations appear as if she is overweighting some states (i.e., those with $S_{i,t+1} > 1$) and underweighting other states (i.e., those with $S_{i,t+1} < 1$).

This proposition is analogous to the result in asset pricing that the law of one price implies the existence of a stochastic discount factor (SDF) that prices all assets. In asset pricing, the SDF is a powerful, unifying tool. Rather than studying the price of each asset in isolation, researchers study the prices of many assets simultaneously by focusing on a single variable (the SDF). To quote Cochrane (2005), “All asset pricing models amount to alternative ways of connecting the stochastic discount factor to data.” We argue that the SBF $S_{i,t+1}$ is similarly powerful for understanding subjective expectations. Rather than separately studying biases in inflation expectations (Coibion and Gorodnichenko, 2015), interest rate expectations (Cieslak, 2018), unemployment expectations (Link et al., 2023), etc., we can

study the single SBF which jointly explains these expectations.

One might wonder whether it is reasonable to assume subjective expectations are coherent. First, it is useful to note that nearly all models of subjective expectations assume agents can add and multiply (e.g., sticky expectations, extrapolation, adaptive expectations, learning, noisy signals, etc.).⁸ In fact, representing beliefs using a SBF $S_{i,t+1}$ provides a useful way to nest and compare all of these models. Second, for our applications, we focus on the Survey of Professional Forecasters. Stark (2013) shows that 80% of these forecasters make their forecasts using a “mathematical/computer model plus subjective adjustments.” The forecasters have an underlying model and, based on subjective beliefs or intuition, may adjust parameter values or input a certain sequence of residuals (often called “add factors”). While these subjective adjustments will impact the forecasts, they still respect the properties of addition and multiplication and therefore ensure that forecasts are coherent.⁹ In short, it seems reasonable to assume these forecasters are sophisticated enough to understand addition and multiplication.

In the subsections below, we highlight three useful features of the SBF.

A. Aggregation

Given a set of individuals, $i = 1, 2, \dots, I$, there is no requirement that individuals must agree with one another. Each individual may be described by a different SBF $S_{i,t+1}$. Define the consensus expectation as $E_t^* [\cdot] \equiv \frac{1}{I} \sum_i E_{i,t}^* [\cdot]$.

Lemma 1. *There is a consensus SBF $S_{t+1} \equiv \frac{1}{I} \sum_i S_{i,t+1}$ that satisfies $E_t^* [X_{t+1}] = E_t [S_{t+1} X_{t+1}]$ for all variables X_{t+1} .*

⁸Breaking these assumptions within a model is quite difficult as it makes the model predictions highly sensitive to small arbitrary changes, such as measuring outcomes in dollars versus cents.

⁹Klein (2018) provides a useful example of how these subjective adjustments are made. “After the preparation of preliminary predictions from the most recently adjusted Wharton-EFU Model, there is a discussion of the assumptions and properties of the prediction with business and government specialists. A priori information on impending labor disputes, hedge purchasing, production bottlenecks, major economic decisions and similar phenomena are then suggested for further modification of parameter or residual values, and a revised forecast is prepared.”

For example, while each forecaster may have a different model of how GDP, unemployment, and inflation interact, there will be a single SBF S_{t+1} that applies to the consensus forecast for all three variables. While we focus on an equal-weighted average across individuals, this lemma can be trivially extended to any weighted average of individuals (e.g., a wealth-weighted average).

We can also extend this idea of aggregation to asset pricing. Suppose X_{t+1} is the payoff for some asset and P_t is the current price of the asset.

Lemma 2. *If $E_{i,t}^* \left[\tilde{M}_{i,t+1} X_{t+1} \right] = P_t$ for all i , then there is a consensus $\tilde{M}_{t+1} \equiv \frac{\sum_i S_{i,t+1} \tilde{M}_{i,t+1}}{\sum_i S_{i,t+1}}$ that satisfies $E_t^* \left[\tilde{M}_{t+1} X_{t+1} \right] = E_t \left[S_{t+1} \tilde{M}_{t+1} X_{t+1} \right] = P_t$.*

Thus, if each individual has an $\tilde{M}_{i,t+1}$ that prices the asset under her individual SBF $S_{i,t+1}$, then there is also a consensus \tilde{M}_{t+1} that prices the asset under the consensus SBF S_{t+1} . Note that this consensus \tilde{M}_{t+1} does not depend on the specific payoff X_{t+1} or price P_t , meaning that the same \tilde{M}_{t+1} applies for any asset that is priced by each individual. Given this property of aggregation, for the rest of the paper, we will focus on a single aggregated agent.

B. Log-normal representation

A useful feature of the SBF S_{t+1} is that it is highly tractable when variables are normally or log-normally distributed. Let lower case letters denote log values, e.g., $x_{t+1} \equiv \ln(X_{t+1})$.

Proposition 2. *For a multivariate x_{t+1} , if the objective distribution is $x_{t+1} \sim N(E_t[x_{t+1}], \Sigma)$ and the agent subjectively believes $x_{t+1} \sim N(E_t^*[x_{t+1}], \Sigma)$, then, $E_t^*[x_{t+1}] = E_t[S_{t+1}x_{t+1}]$ and $E_t^*[X_{t+1}] = E_t[S_{t+1}X_{t+1}]$ for*

$$\begin{aligned} s_{t+1} &\equiv -\frac{1}{2}\beta_t'\Sigma\beta_t + \beta_t'\eta_{t+1} \\ \beta_t &\equiv \Sigma^{-1}(E_t^*[x_{t+1}] - E_t[x_{t+1}]) \\ \eta_{t+1} &\equiv x_{t+1} - E_t[x_{t+1}]. \end{aligned}$$

In other words, if x_{t+1} is conditionally multivariate normal, then we can write S_{t+1} as conditionally log-normal and the shocks to s_{t+1} as a linear combination of the objective shocks to x_{t+1} , i.e., $\beta'_t \eta_{t+1}$. For an element $\beta_{j,t}$ of the vector β_t , $\beta_{j,t} > 0$ means that it is as if the agent is exaggerating the probability of positive shocks $\eta_{j,t+1}$ and understating the probability of negative shocks. The converse is true for $\beta_{j,t} < 0$.

Note that Proposition 2 allows for a wide range of possible biases in subjective beliefs. The only restrictions placed on subjective beliefs are that the agent knows that x_{t+1} is conditionally normal and the agent knows the objective covariance matrix. Beyond those assumptions, the agent is allowed to make any possible mistake about the conditional mean $E_t^* [x_{t+1}]$. In the subsection below, we show specific examples for biases due to extrapolation and diagnostic expectations, which both generate biased beliefs about the conditional mean but do not affect the believed conditional variance. In the corollary below, we also generalize this proposition to the case where the agent has biased beliefs about the conditional mean and the conditional covariance matrix.

Corollary 1. *For a multivariate x_{t+1} , if the objective distribution is $x_{t+1} \sim N(E_t[x_{t+1}], \Sigma_t)$ and the agent subjectively believes $x_{t+1} \sim N(E_t^*[x_{t+1}], \Sigma_t^*)$, then, $E_t^*[x_{t+1}] = E_t[S_{t+1}x_{t+1}]$ and $E_t^*[X_{t+1}] = E_t[S_{t+1}X_{t+1}]$ for*

$$\begin{aligned} s_{t+1} &\equiv -\frac{1}{2}\beta'_t \Sigma_t^* \beta_t + \beta'_t \eta_{t+1} + \frac{1}{2} \left[\log \left(\frac{\det(\Sigma_t)}{\det(\Sigma_t^*)} \right) + \eta'_{t+1} (\Sigma_t^{-1} - \Sigma_t^{*-1}) \eta_{t+1} \right] \\ \beta_t &\equiv \Sigma_t^{*-1} (E_t^*[x_{t+1}] - E_t[x_{t+1}]) \\ \eta_{t+1} &\equiv x_{t+1} - E_t[x_{t+1}]. \end{aligned}$$

Compared to Proposition 2, the log SBF s_{t+1} now depends not just on the shocks η_{t+1} but also on the squared shocks $\eta'_{t+1} (\Sigma_t^{-1} - \Sigma_t^{*-1}) \eta_{t+1}$ to account for the difference in the believed variance.

B.1. Extrapolation and diagnostic expectations

For intuition, consider the case of AR(1) real GDP (RGDP) growth,

$$g_{t+1} = \rho g_t + \eta_{t+1} \quad (4)$$

where $\rho > 0$ and $\eta_{t+1} \sim N(0, \sigma^2)$. For simplicity, we set the mean of the process to 0.

Suppose the agent is extrapolative, in the sense that the agent believes the conditional distribution is $g_{t+1} \sim N(\rho^* g_t, \sigma^2)$ where $\rho^* > \rho$. Then, from Proposition 2, the agent's beliefs can be represented by the log SBF

$$s_{t+1} = -\frac{1}{2} \frac{\beta_t^2}{\sigma^2} + \beta_t \eta_{t+1} \quad (5)$$

$$\beta_t = \frac{\rho^* - \rho}{\sigma^2} g_t. \quad (6)$$

Note that states with $s_{t+1} > 0$ are states with $S_{t+1} > 1$, meaning that the SBF exaggerates the probability of these states. The converse is true for $s_{t+1} < 0$ and $S_{t+1} < 1$.

The $-\frac{1}{2} \frac{\beta_t^2}{\sigma^2}$ term is simply a normalization term that ensures $E_t[S_{t+1}] = 1$. The important component of s_{t+1} is the loading β_t on the objective shock η_{t+1} . Given $\rho^* > \rho$, the loading is positive when $g_t > 0$. When current growth is high, it is as if the agent is exaggerating the probability of positive η_{t+1} and understating the probability of negative η_{t+1} . When current growth is low, it is as if the agent is exaggerating the probability of negative η_{t+1} and understating the probability of positive η_{t+1} .

Importantly, we do not require that the agent actually thinks in terms of distorted probabilities. In this example, the agent does not intentionally exaggerate or understate certain states of the world. The agent simply believes in a different persistence parameter (ρ^*) than what the econometrician estimates (ρ). The purpose of Proposition 2 is to highlight that the econometrician can represent the agent's beliefs using the SBF S_{t+1} .

As a second example, suppose the agent has diagnostic expectations of future growth. Following Bordalo, Gennaioli, and Shleifer, 2018, the agent believes the conditional distribution is $g_{t+1} \sim N(\rho g_t + \theta \rho \eta_t, \sigma^2)$ where $\theta > 0$ reflects how much the agent overreacts to the

most recent shock η_t . In this case, the agent's beliefs can still be represented by s_{t+1} equal to $-\frac{1}{2}\frac{\beta_t^2}{\sigma^2} + \beta_t\eta_{t+1}$, but the loading is now

$$\beta_t = \frac{\theta\rho}{\sigma^2}\eta_t. \quad (7)$$

When the recent shock η_t is positive, it is as if the agent exaggerates the probability of positive future η_{t+1} and understates the probability of negative η_{t+1} .

Because Proposition 2 is so general, it provides a convenient way to nest and compare extrapolation and diagnostic expectations. Specifically, both extrapolation and diagnostic expectations can be summarized by their implied loading (β_t). For extrapolation, the loading equals $\frac{\rho^* - \rho}{\sigma^2}g_t$, meaning that the loading depends on the current level of g_t and is persistent if growth is persistent. For diagnostic expectations, the loading equals $\frac{\theta\rho}{\sigma^2}\eta_t$ meaning that the loading depends on the current shock to g_t and should have zero persistence, regardless whether growth is persistent or not. In Section III.A, we empirically estimate β_t and test whether it is closer to the predictions of extrapolation or diagnostic expectations.

C. Connection to excess returns

By framing subjective beliefs as a probability distortion, we can easily combine subjective beliefs and preferences about different states of the world. We denote the agent's preference for payoffs in different states of the world with \tilde{M}_{t+1} . For example, in a consumption-based asset pricing model, \tilde{M}_{t+1} would be $\frac{u'(C_{t+1})}{u'(C_t)}$. Let R_t^f denote the one-period return on a risk-free bond and let R_{t+1} denote the one-period return on a risky asset. If the agent prices the risk-free bond and the risky asset, then

$$E_t^* \left[\tilde{M}_{t+1} \right] R_t^f = 1 \quad (8)$$

$$E_t^* \left[\tilde{M}_{t+1} R_{t+1} \right] = 1. \quad (9)$$

Proposition 3. *Assume $R_{t+1}^e \equiv R_{t+1}/R_t^f$ is log-normally distributed, S_{t+1} is conditionally log-normal, and \tilde{M}_{t+1} is conditionally log-normal. If $E_t^* \left[\tilde{M}_{t+1} R_{t+1} \right] = E_t^* \left[\tilde{M}_{t+1} R_t^f \right] = 1$,*

then

$$\log (E [R_{t+1}^e]) = Cov (-s_{t+1}, r_{t+1}^e) + Cov (-\tilde{m}_{t+1}, r_{t+1}^e). \quad (10)$$

Intuitively, Proposition 3 says that if the agent prices the risk-free bond and the risky asset, then a high average R_{t+1}^e must be due to (i) the risky asset paying off in states the agent thinks are unlikely, $Cov (-s_{t+1}, r_{t+1}^e)$, and/or (ii) the risky asset paying off in states the agent does not prefer, $Cov (-\tilde{m}_{t+1}, r_{t+1}^e)$.

Conveniently, equation (10) provides a clean separation of beliefs and preferences. Even if we do not know \tilde{m}_{t+1} , we can still estimate the unconditional mean $E [R_{t+1}^e]$ and the unconditional covariance $Cov (-s_{t+1}, r_{t+1}^e)$ using a time series for R_{t+1}^e and s_{t+1} . For example, we can measure whether subjective beliefs about RGDP growth account for the majority of the investment anomaly by estimating s_{t+1} from survey expectations of RGDP growth and comparing $\log (E [R_{t+1}^e])$ to $Cov (-s_{t+1}, r_{t+1}^e)$.

Further, equation (10) holds even if s_{t+1} and \tilde{m}_{t+1} are correlated. As an example, suppose the agent has biased beliefs about RGDP growth and also has preferences about RGDP growth. Then $Cov (-s_{t+1}, r_{t+1}^e)$ will capture the magnitude of the agent's biased beliefs and $Cov (-\tilde{m}_{t+1}, r_{t+1}^e)$ will capture the magnitude of the agent's preferences. If the agent exaggerates about the probability of low growth states by a factor of 2 relative to high growth states and prefers payoffs in low growth states 10 times as much as she prefers payoffs in high growth states, then $Cov (-s_{t+1}, r_{t+1}^e)$ will be much smaller than $Cov (-\tilde{m}_{t+1}, r_{t+1}^e)$. Even if the researcher does not observe \tilde{m}_{t+1} , the researcher can still correctly conclude that $Cov (-s_{t+1}, r_{t+1}^e)$ only accounts for a small amount of $\log (E [R_{t+1}^e])$.

II. Data

We use two sources of survey data to measure subjective expectations. The main source of survey data used is the Survey of Professional Forecasters, which contains quarterly forecasts for a wide array of macroeconomic variables. Since 1981Q3, the Survey of Professional

Forecasters contains complete coverage for 15 economic variables: real GDP, real consumption, industrial production, real residential investment, real non-residential investment, real federal government spending, real state and local government spending, housing starts, corporate profits, CPI inflation, the 3-month Treasury bill rate, the Aaa rate, the unemployment rate, real change in private inventories, and real net exports. For all variables, we focus on the four quarter ahead forecast.¹⁰ We calculate the consensus forecast as the average across the individual-level forecasts.

For our analysis, we need to convert all 15 variables to stationary processes. For the first nine variables, we calculate the implied forecasted growth by dividing the forecasted future level by the most recently available level. The next four variables are already reported as stationary variables, so we apply no changes. Finally, the last two variables (real change in private inventories and real net exports) can potentially be zero, meaning that the annual growth may not be stationary. Therefore, we use the forecasted future level divided by the most recently reported real GDP to normalize the series.

The secondary source of data is the Blue Chip survey. The Blue Chip sample starts in 1988Q1. We consider the complete list of variables for which we have survey data since 1988Q1 and are not included in the Survey of Professional Forecasters over our sample, which is a total of 9 variables. These are the prime rate, the fed funds rate, the mortgage rate, LIBOR, and Treasury rates for five different maturities (6-months, 1-year, 2-years, 5-years, and 10-years). Just as we did for the Survey of Professional Forecasters, we focus on the four quarter ahead forecasts and calculate the consensus forecast as the average across the individual-level forecasts. Given that all variables are rates, we do not need to renormalize for stationarity.

The realized outcomes for all interest rate variables are obtained from the Federal Reserve Bank of St. Louis. The realized outcomes for all other variables are obtained from the real-

¹⁰Note that CPI inflation forecasts are reported as quarter-over-quarter growth rates, so we calculate the forecasted annual growth rate using the geometric mean of forecasted quarterly growth out to t plus four quarters.

time data files maintained by the Federal Reserve Bank of Philadelphia. Refer to Appendix B for full details on each of the surveys and the data construction.

III. Application to macroeconomic forecasts

In this section, our goal is to demonstrate how utilizing tools from asset pricing opens new doors for understanding subjective expectations. First, we are able to largely condense survey expectations for 15 different macroeconomic variables down to a single estimated SBF \hat{S}_{t+1} that is related to RGDP growth and the T-bill rate. Second, after condensing the survey data, we can compare the estimated SBF to the predictions of common expectation-formation models, in this case extrapolation and diagnostic expectations. Third, we show that the SBF allows us to predict subjective expectations for other variables outside of our original 15 variable dataset, and we then confirm the accuracy of these predictions. Broadly, we find that the SBF \hat{S}_{t+1} is successful in unifying expectations across many macroeconomic variables.

A. Condensing macroeconomic expectations

We study the consensus forecasts for the 15 macroeconomic variables contained in the Survey of Professional Forecasters. Let $E_t^*[x_{t+1}]$ denote this 15-variable vector. Figure 1 shows the correlation of these expectations. Importantly, our analysis does not require that variables are uncorrelated. As one would expect, there is a wide range of correlations across these variables. While some expectations are highly correlated (like *rgdp* and *rcon*, $corr = .92$), some of them have no correlation, (like *rgf* and *rnresin*, $corr = -.04$), and some are negatively correlated (like *housing* and *rgsl*, $corr = -.59$). The average pair-wise correlation is 0.16.

We calculate statistical expectations using a VAR(1) model for x_{t+1} . Specifically, we

estimate

$$x_{t+1} = A + B \begin{pmatrix} x_t & E_t^* [x_{t+1}] \end{pmatrix} + \eta_{t+1} \quad (11)$$

and calculate statistical expectations as

$$E_t [x_{t+1}] = A + B \begin{pmatrix} x_t & E_t^* [x_{t+1}] \end{pmatrix}.$$

We include the survey expectations in equation (11) to ensure that our statistical expectations contain any information known to the forecasters. This ensures that any discrepancy between $E_t^* [x_{t+1}]$ and $E_t [x_{t+1}]$ is due to the statistical expectations being a better predictor of x_{t+1} , and not from any informational advantage from the forecasters.¹¹

Given our objective distribution, $x_{t+1} \sim N(E_t [x_{t+1}], \Sigma)$, we assume that forecasters believe $x_{t+1} \sim N(E_t^* [x_{t+1}], \Sigma)$ where $E_t^* [x_{t+1}]$ is measured from the survey forecasts.¹² From Proposition 2, this means that $E_t [S_{t+1} x_{t+1}]$ will perfectly replicate the survey expectations for $s_{t+1} = -\frac{1}{2} \beta_t' \Sigma \beta_t + \beta_t' \eta_{t+1}$ and $\beta_t = \Sigma^{-1} (E_t^* [x_{t+1}] - E_t [x_{t+1}])$. Our goal in this section is to test whether we can replicate the survey expectations using a log SBF \hat{s}_{t+1} based on a smaller number of variables.

Specifically, given any subset of variables $\hat{x}_{t+1} \subset x_{t+1}$, we can estimate the log SBF that perfectly matches the expectations for \hat{x}_{t+1} ,

$$\hat{s}_{t+1} \equiv -\frac{1}{2} \hat{\beta}_t' \hat{\Sigma} \hat{\beta}_t + \hat{\beta}_t' \hat{\eta}_{t+1} \quad (12)$$

where $\hat{\eta}_{t+1}$ is the objective shocks to \hat{x}_{t+1} , $\hat{\Sigma}$ is the covariance matrix of $\hat{\eta}_{t+1}$, and $\hat{\beta}_t = \hat{\Sigma}^{-1} (E_t^* [\hat{x}_{t+1}] - E_t [\hat{x}_{t+1}])$. Then, we can estimate synthetic expectations based on \hat{s}_{t+1} for

¹¹For example, if the survey forecasts are the best possible predictor of future x_{t+1} , then our method will simply give $E_t [x_{t+1}] = E_t^* [x_{t+1}]$.

¹²Appendix F shows that we find very similar results if we assume the variables are objectively and subjectively log-normally distributed rather than normally distributed.

the remaining variables as¹³

$$\begin{aligned}\hat{E}_t^* [x_{t+1}] &\equiv E_t \left[\hat{S}_{t+1} x_{t+1} \right] \\ &= E_t [x_{t+1}] + Cov(\eta_{t+1}, \hat{\eta}_{t+1}) \hat{\beta}_t\end{aligned}\tag{13}$$

We use Figure 1 to gauge the size of our subset \hat{x}_{t+1} . The pair-wise correlations are clustered by a hierarchical tree method, which iteratively merges clusters based on their similarity.¹⁴ Two natural clusters of variables arise with similar correlations. This motivates us to summarize the expectations data using a distortion based on two variables.¹⁵ We choose RGDP growth (*rgdp*) and the T-bill rate (*tbill*) as our two variables, as these two variables lie at the center of the two large clusters in Figure 1. Economically, these two variables are fairly easy to understand given their connection to overall economic activity and monetary policy.¹⁶

Table I evaluates the fit of these synthetic expectations. The first column shows the average R^2 from regressing $E_t^* [x_{j,t+1}]$ on $\hat{E}_t^* [x_{j,t+1}]$ for our 15 variables j . Overall, synthetic expectations based on \hat{s}_{t+1} explain roughly 2/3 (64.6%) of all variation in the 15 macroeconomic expectations. However, this could be due to $E_t^* [x_{j,t+1}]$ being similar to $E_t [x_{j,t+1}]$. As shown in equation (13), even with no distortion, synthetic expectations would still vary due to variation in $E_t [x_{j,t+1}]$.

To better evaluate the role of the distortion, we note that equation (13) can be rewritten as

$$\hat{E}_t^* [x_{t+1}] - E_t [x_{t+1}] = Cov(\eta_{t+1}, \hat{\eta}_{t+1}) \hat{\beta}_t\tag{14}$$

$$= Cov(\eta_{t+1}, \hat{\eta}_{t+1}) \hat{\Sigma}^{-1} (E_t^* [\hat{x}_{t+1}] - E_t [\hat{x}_{t+1}]).\tag{15}$$

¹³The calculation of $E_t \left[\hat{S}_{t+1} x_{t+1} \right]$ utilizes the formula for the mean of a normal log-normal mixture.

¹⁴Given that unemployment is countercyclical, we use negative subjective expectations of unemployment when clustering the correlation matrix.

¹⁵The figure also suggests that one could potentially consider a third, smaller cluster represented by either housing starts or real residential investment. For parsimony, we focus on the two larger clusters.

¹⁶In Appendix D, we show that RGDP growth and the T-bill rate are not only intuitive variables for understanding beliefs about the broader macroeconomy, but are also quantitatively quite close to best possible pair of variables for condensing the Survey of Professional Forecasters data.




Figure 1. Correlation matrix of expectations. This figure shows the pairwise correlation of the survey expectations of 15 variables in the Survey of Professional Forecasters for 1981Q3-2022Q2. The variables are grouped using a simple hierarchical tree method. We use negative unemployment in the correlation matrix so that it is positively related to other typical business-cycle variables.

The second column of Table I tests how well the synthetic bias, measured as $\hat{E}_t^* [x_{t+1}] - E_t [x_{t+1}]$, matches the survey bias, $E_t^* [x_{t+1}] - E_t [x_{t+1}]$. Across the 15 variables, we find that nearly half (47.2%) of the variation in survey bias can be explained by biased expectations of just two variables, RGDP growth and the T-bill rate ($E_t^* [\hat{x}_{t+1}] - E_t [\hat{x}_{t+1}]$), and the objective covariance of shocks to the remaining 13 variables with shocks to RGDP growth and the T-bill rate ($Cov(\eta_{t+1}, \hat{\eta}_{t+1}) \hat{\Sigma}^{-1}$).

To gauge the performance of our method, we compare our results to less restrictive versions of equation (15). The most straightforward test is a regression

$$E_t^* [x_{t+1}] - E_t [x_{t+1}] = \alpha + \Gamma (E_t^* [\hat{x}_{t+1}] - E_t [\hat{x}_{t+1}]) + \varepsilon_t. \quad (16)$$

This is identical to equation (15), except that the matrix of coefficients Γ is now flexible, rather than being determined by the objective covariance of shocks. This specification represents the best linear prediction one can achieve using RGDP growth and T-bill rate biases,

Table I

Condensing the Survey of Professional Forecasters

This table evaluates the ability of the synthetic expectations formed from \hat{s}_{t+1} to explain movements in the 15 Survey of Professional Forecasters variables. Column 1 shows the average R^2 of the regression of survey expectations ($E_t^*[x_{j,t+1}]$) on synthetic expectations ($\hat{E}_t^*[x_{j,t+1}]$) across the 15 different variables. Column 2 shows the average R^2 of the regression of $E_t^*[x_{j,t+1}] - E_t[x_{j,t+1}]$ on $\hat{E}_t^*[x_{j,t+1}] - E_t[x_{j,t+1}]$ where $E_t[x_{j,t+1}]$ is the statistical expectation. For comparison, Column 3 shows the average R^2 of the best linear predictor of $E_t^*[x_{j,t+1}] - E_t[x_{j,t+1}]$ using the individual biases in *rgdp* and *tbill* coming from equation (16) and Columns 4 and 5 show the explanatory power of the first two principal components of the 15 biases $E_t^*[x_{j,t+1}] - E_t[x_{j,t+1}]$.

	$E_t^*[x_{t+1}]$	$E_t^*[x_{t+1}] - E_t[x_{t+1}]$			
	$\hat{E}_t^*[x_{t+1}]$	$\hat{E}_t^*[x_{t+1}] - E_t[x_{t+1}]$	Best Linear Predictor	PC-1	PC-2
$R^2(\%)$	64.6	47.2	52.3	43.0	63.8

and thus, provides an upper bound of how much can one explain *given these two variables*. Averaging across all variables, we find that this regression approach gives an R^2 of 52.3%. Therefore, our estimated SBF \hat{s}_{t+1} gives an average R^2 (47.2%) that is quite close to this upper bound (52.3%).

Let's look at an even more challenging test. While the first benchmark evaluates the distortion relative to the best linear predictor of using RGDP growth and T-bill rate biases, we can also show that the distortion performs well *relative to any two arbitrary time series* Λ_t . Our synthetic expectations attempt to characterize the 15 biases $E_t^*[x_{t+1}] - E_t[x_{t+1}]$ using two time series $E_t^*[\hat{x}_{t+1}] - E_t[\hat{x}_{t+1}]$ and a fixed matrix $Cov(\eta_{t+1}, \hat{\eta}_{t+1}) \hat{\Sigma}^{-1}$ based on the objective covariances. The final two columns in Table I show a more general upper bound based on principal components analysis (PCA),

$$E_t^*[x_{t+1}] - E_t[x_{t+1}] = \alpha + \Gamma \Lambda_t. \quad (17)$$

The first two principal components of the 15 biases explain 63.8% of the variation, and the first principal component explains 43.0%. This means that the estimated SBF performs better than the first principal component of these series and captures roughly four fifths of the maximum possible R^2 , 63.8%.

Table II shows a detailed comparison between the synthetic expectations and the survey expectation expectations at the individual variable level. The correlation between $\hat{E}_t^*[x_{j,t+1}]$

Table II

Comparing synthetic expectations and the Survey of Professional Forecasters

This table shows the correlation of each of the Survey of Professional Forecasters survey expectations and biases with their respective synthetic expectations and biases formed from the log SBF \hat{s}_{t+1} . The first column shows the correlation of survey expectations with synthetic expectations for the 15 variables. The second column shows the correlation of the survey biases with the synthetic biases for the same 15 variables. Note that the log SBF \hat{s}_{t+1} is formed only using RGDP growth and the T-Bill rate biases.

$x_{j,t+1}$	$\text{Corr} \left(E_t^* [x_{j,t+1}], \hat{E}_t^* [x_{j,t+1}] \right)$	$\text{Corr} \left(E_t^* [x_{j,t+1}] - E_t [x_{j,t+1}], \hat{E}_t^* [x_{j,t+1}] - E_t [x_{j,t+1}] \right)$
<i>rgdp</i>	1	1
<i>rcon</i>	0.8891	0.8695
<i>cpi</i>	0.6122	0.3941
<i>unempl</i>	0.8973	0.6131
<i>indp</i>	0.7141	0.8324
<i>tbill</i>	1	1
<i>aaa</i>	0.9692	0.4792
<i>rnresin</i>	0.7214	0.8007
<i>rresinv</i>	0.7143	0.6287
<i>rgf</i>	0.8464	0.5976
<i>rgsl</i>	0.8165	0.4270
<i>housing</i>	0.5784	0.3515
<i>rcbi</i>	0.4977	0.7479
<i>rexpport</i>	0.9367	0.5234
<i>cprof</i>	0.6273	0.5744

and $E_t^* [x_{j,t+1}]$ is larger than 0.50 for all variables, and larger than 0.70 for nine out of the 13 independent variables. Perhaps more relevant is the ability of the synthetic biases $\hat{E}_t^* [x_{j,t+1}] - E_t [x_{j,t+1}]$ to match the survey biases $E_t^* [x_{j,t+1}] - E_t [x_{j,t+1}]$, as this comparison directly captures the ability of the SBF to capture deviations from the statistical expectations. The correlation between $\hat{E}_t^* [x_{j,t+1}] - E_t [x_{j,t+1}]$ and $E_t^* [x_{j,t+1}] - E_t [x_{j,t+1}]$ is also high, with an average correlation across variables of 0.66 and nine out of the 13 independent variables having a correlation larger than 0.50.

It is important to emphasize the key benefits of the SBF relative to the two benchmarks. First, PCA and best linear prediction require the use of all 15 biases $E_t^* [x_{t+1}] - E_t [x_{t+1}]$ as Γ and/or Λ_t are directly estimated to fit these biases. Our method only depends on biases in two variables and the objective covariance matrix. The loadings β_t are not fitted to any of the expectations outside of the two main biases. This simplifies the translation to

economic models, as one only needs to understand two expectations. For empirical work, this parsimony makes our approach more interpretable and easier to implement in various economic contexts. For example, Section III.B shows that we can study expectations beyond the 15 variables in the Survey of Professional Forecasters without needing to re-estimate \hat{s}_{t+1} , whereas the principal components and best linear coefficients would both need to be re-estimated for these new variables.

Further, the link from these two variables to the other variables is entirely determined by the objective covariance matrix. This means that if we can understand the biases in these two variables, then we can readily explain the biases in the other variables. For example, overstating RGDP growth by 1 percentage point causes agents to understate unemployment by 0.51pp because of the empirical covariance between RGDP growth and unemployment. This contrasts with other methods like best linear estimates, which would require a specific mechanism to explain why agents perceive a certain link between RGDP growth and unemployment. Similarly, using PCA in the context of a model would require interpreting the principal components and understanding the weight matrix (i.e., why does the unemployment bias load on the first PC with a certain coefficient).

Finally, our method also allows for the study of biased beliefs about correlated variables. For instance, in cases where the biases in RGDP growth and T-bill rate are correlated, traditional methods like the best linear prediction may struggle to distinguish between them. For instance, if an agent is understating future unemployment, it may be unclear whether this is due to overstating the chance of an expansion (RGDP growth) or overstating the likelihood of accommodative monetary policy (T-bill rate). Our method captures magnitudes, allowing us to differentiate between these scenarios. Even if the biases in RGDP growth and the T-bill rate are correlated, we can discern their relative importance based on the magnitude of each bias and the comovement between each variable and unemployment. This enables a more nuanced understanding of the underlying drivers of forecast biases across various economic indicators.

To highlight that this scenario of correlated biases is not just a hypothetical, the top panel of Figure A.1 plots the biases $E_t^*[\hat{x}_{t+1}] - E_t[\hat{x}_{t+1}]$ for these two variables. Appendix C discusses these series in more detail along with the estimated time series \hat{s}_{t+1} .

A.1. Connection to expectation-formation models

As mentioned above, a key benefit of our SBF \hat{s}_{t+1} relative to PCA or best linear predictions is that \hat{s}_{t+1} can be easily translated into models of expectation formation.¹⁷ To do that, we can simply compare the estimated loadings β_t with the loadings predicted by the common models.

In Section I.B.1, we discussed how two common expectation-formation mechanisms, extrapolation and diagnostic expectations, translate into implied loadings β_t . By equation (6), the implied loadings of the extrapolation model are proportional to the current value of \hat{x}_t and their persistence is thus equal the persistence of \hat{x}_t . Similarly, by equation (7), the loadings implied by diagnostic expectations should be proportional to the most recent shocks $\hat{\eta}_t$, and they should have 0 persistence.

While neither model performs perfectly, our estimated loadings are overall closer to the predictions of extrapolation than the predictions of diagnostic expectations. First, we find that the *rgdp* loading is significantly correlated with current RGDP growth (0.29*) and the *tbill* loading is significantly correlated with the current T-bill rate (0.45***). In comparison, the correlation of each loading with the *rgdp* shock and the *tbill* shock are smaller and insignificant, 0.20 and 0.12, respectively. Both loadings have significant persistence, 0.32*** annually for the *rgdp* loading and 0.46*** annually for the T-bill rate. This does not align with the prediction of diagnostic expectations and moderately aligns with the prediction of extrapolation. The persistence of RGDP growth and the T-bill rate are 0.08 and 0.88***, meaning that the persistences of the loadings are quantitatively different from what extrapolation

¹⁷Most models of expectation-formation satisfy the assumptions required for the existence of the SBF. Further, many of these models also satisfy the conditions for Proposition 2, which assumes people make mistakes about conditional means but not about conditional variances.

olation would imply. But, the qualitative prediction of extrapolation that the persistence of $\hat{\beta}_t$ should be increasing in the persistence of \hat{x}_t does hold (i.e., the *tbill* loading is more persistent than the *rgdp* loading).

B. Forming expectations for other variables

Beyond condensing existing survey data, a benefit of our estimated log SBF \hat{s}_{t+1} is that we can form synthetic expectations for other variables even if we do not have survey expectations available for those variables. For example, the Survey of Professional Forecasters contains relatively few financial variables. Despite this, we can use our estimated \hat{s}_{t+1} to form synthetic expectations for a variety of financial variables, such as the mortgage rate or the long-term risk-free rate (e.g., the 10-year Treasury rate). In this subsection, we form synthetic expectations for nine financial variables without using any survey expectations for these variables. We then evaluate these synthetic expectations by comparing them to Blue Chip survey expectations for these nine variables.

Let y_{t+1} represent the nine financial variables. These variables are the prime rate, the fed funds rate, the mortgage rate, LIBOR, and Treasury rates for five different maturities (6-months, 1-year, 2-years, 5-years, and 10-years). We choose these variables as this is the maximum set of variables that are (i) not included in the Survey of Professional Forecasters data and (ii) are included in the Blue Chip data, which will allow us to evaluate the synthetic expectations.

We calculate synthetic expectations $\hat{E}_t^*[y_{t+1}]$ solely using realized data and the Survey of Professional Forecasters survey data. First, we estimate a VAR(1) model for y_{t+1} ,¹⁸

$$y_{t+1} = A_y + B_y \begin{pmatrix} y_t & x_t & E_t^*[x_{t+1}] \end{pmatrix} + \eta_{y,t+1} \quad (18)$$

$$E_t[y_{t+1}] = A_y + B_y \begin{pmatrix} y_t & x_t & E_t^*[x_{t+1}] \end{pmatrix}. \quad (19)$$

¹⁸Appendix F shows that we find similar results if we estimate a VAR(1) process for the log of the financial variables rather than the level.

To ensure our statistical expectations contain as much current information as possible, we include the current value x_t and the survey expectations $E_t^*[x_{t+1}]$ for the Survey of Professional Forecasters variables in equation (18). Then, using our log SBF \hat{s}_{t+1} estimated from the Survey of Professional Forecasters RGDP growth and T-bill rate expectations, we calculate our synthetic expectations for y_{t+1} as

$$\begin{aligned}\hat{E}_t^*[y_{t+1}] &\equiv E_t[\hat{S}_{t+1}y_{t+1}] \\ &= E_t[y_{t+1}] + Cov(\eta_{y,t+1}, \hat{\eta}_{t+1})\hat{\beta}_t.\end{aligned}\tag{20}$$

Importantly, we do not use any survey expectations of y_{t+1} in the construction of $\hat{E}_t^*[y_{t+1}]$.

How well do these synthetic expectations match actual subjective expectations for these nine variables? In Table III, we compare the synthetic expectations $\hat{E}_t^*[y_{t+1}]$ to subjective expectations measured from Blue Chip, $E_t^*[y_{t+1}]$. The first column of Table III shows that the synthetic expectations are highly correlated with the actual survey expectations. For all nine variables, the correlation is above 0.90. When we focus specifically on the bias in expectations, $E_t^*[y_{j,t+1}] - E_t[y_{j,t+1}]$, we still find a notable correlation, on average 0.74. This is similar to the performance of the synthetic expectations for the Survey of Professional Forecasters variables in Table II, indicating that the log SBF \hat{s}_{t+1} is similarly effective in summarizing biases in both the Survey of Professional Forecasters and the Blue Chip data. This extension is non-trivial as the two groups of forecasters could potentially have different beliefs, which would generally make it harder to find a single SBF that accurately summarizes the expectations of both groups.

To push this idea of jointly explaining the Survey of Professional Forecasters and the Blue Chip forecasts further, Table IV tests how well our log SBF \hat{s}_{t+1} based on two variables summarizes the combined 24 forecasts from both groups. Let z_{t+1} be the union of the 15 Survey of Professional Forecasters variables x_{t+1} and the 9 Blue Chip variables y_{t+1} . Overall, we find that the synthetic expectations account for the majority of the variation in

Table III

Synthetic expectations out of sample - Blue Chip survey expectations

This table evaluates the ability of the log SBF \hat{s}_{t+1} formed from only the RDGP growth and T-Bill rate biases to generate synthetic expectations for Blue Chip variables. The first column shows the correlation of the time series of the Blue Chip survey expectations $E_t^*[y_{j,t+1}]$ with the synthetic expectations $\hat{E}_t^*[y_{j,t+1}]$ constructed from \hat{s}_{t+1} . The second column shows the correlation of the survey biases with the synthetic biases for the same variables.

$y_{j,t+1}$	$\text{Corr} \left(E_t^* [y_{j,t+1}], \hat{E}_t^* [y_{j,t+1}] \right)$	$\text{Corr} \left(E_t^* [y_{j,t+1}] - E_t [y_{j,t+1}], \hat{E}_t^* [y_{j,t+1}] - E_t [y_{j,t+1}] \right)$
<i>prime</i>	0.9434	0.8573
<i>fedfunds</i>	0.9455	0.8488
<i>mortgage</i>	0.9304	0.4869
<i>tbill-1yr</i>	0.9309	0.8427
<i>tbill-6m</i>	0.9322	0.8585
<i>tnote-10yr</i>	0.9479	0.5229
<i>tnote-5yr</i>	0.9434	0.6784
<i>tnote-2yr</i>	0.9343	0.7942
<i>libor-3m</i>	0.9110	0.8134

the survey expectations, with an average R^2 of 72.0%. If we focus on biases in expectations, $E_t^*[z_{t+1}] - E_t[z_{t+1}]$, the synthetic expectations account for half of all variation, 50.3%.

The last three columns of Table IV compare our results to the upper bounds implied by the best linear predictor and PCA. Just as in equation (15), the synthetic bias for each of our 24 variables is equal to a linear combination of the bias in RGDP growth expectations and T-bill rate expectations, where the coefficients are determined entirely by the objective covariance of shocks. We can compare this to the best linear predictor from a regression,

$$E_t^*[z_{t+1}] - E_t[z_{t+1}] = \alpha_z + \Gamma_z (E_t^*[\hat{x}_{t+1}] - E_t[\hat{x}_{t+1}]) + \varepsilon_{z,t} \quad (21)$$

where the coefficient matrix Γ_z is unrestricted. We find that the average R^2 produced by the synthetic expectations (50.3%) is quite close to the upper bound implied by the best linear predictor of 55.0%.

Similarly, the synthetic expectations perform well even when compared to the more generalized upper bound implied by PCA. By applying PCA to the expanded set of all 24 variables, the biases are now allowed to depend on any time series $\Lambda_{z,t}$ and use any coefficient matrix

Table IV

Condensing Blue Chip and the Survey of Professional Forecasters

This table evaluates the ability of the synthetic expectations $\hat{E}_t^*[z_{t+1}]$ formed from \hat{s}_{t+1} to explain the 24 Blue Chip + Survey of Professional Forecasters variables $E_t^*[z_{t+1}]$. Column 1 shows the average R^2 of the regression of the survey expectations $E_t^*[z_{j,t+1}]$ on $\hat{E}_t^*[z_{j,t+1}]$ across the 24 different variables. Column 2 shows the average R^2 of the regression of the survey biases $E_t^*[z_{j,t+1}] - E_t[z_{j,t+1}]$ on the synthetic biases $\hat{E}_t^*[z_{j,t+1}] - E_t[z_{j,t+1}]$. For comparison, Column 3 shows the average R^2 of the best linear predictor of $E_t^*[z_{j,t+1}] - E_t[z_{j,t+1}]$ using the individual *rgdp* and *tbill* biases coming from equation (21) and Columns 4 and 5 show the explanatory power of the first two principal components of the 24 biases $E_t^*[z_{j,t+1}] - E_t[z_{j,t+1}]$.

	$E_t^*[z_{t+1}]$	$E_t^*[z_{t+1}] - E_t[z_{t+1}]$			
	$\hat{E}_t^*[z_{t+1}]$	$\hat{E}_t^*[z_{t+1}] - E_t[z_{t+1}]$	Best Linear Predictor	PC-1	PC-2
$R^2(\%)$	72.0	50.3	55.0	47.4	66.4

Γ_z ,

$$E_t^*[z_{t+1}] - E_t[z_{t+1}] = \alpha_z + \Gamma_z \Lambda_{z,t}. \quad (22)$$

The fourth column of Table IV shows that the synthetic expectations outperforms the first principal component and achieves more than three fourths of the maximum possible R^2 of 66.4%.

Overall, we find that the log SBF \hat{s}_{t+1} manages to (i) accurately predict subjective expectations and biases for other variables without using any survey data on those variables and (ii) performs nearly as well as theoretical upper bounds in condensing biases in many expectations down to just biases in two expectations. The first item effectively acts as an out of sample test for the SBF. The second item reinforces the finding from Section III.A that the distinction between subjective expectations and statistical expectations for many variables – in this case 24 variables – can largely be explained beliefs about a few key variables and the objective relationships between variables.

IV. Decomposing excess returns

In this section, we demonstrate how framing subjective expectations using an SBF is useful for studying asset prices. In particular, by summarizing subjective expectations for many

variables into a single distortion, we can easily integrate subjective expectations into asset pricing equations, which allows us to distinguish the roles of beliefs versus preferences. This section applies this methodology to the Fama-French factors from Fama and French (2015), to the behavioral factors constructed in Daniel, Hirshleifer, and Sun (2020), and also to a comprehensive set of anomalies compiled in Chen and Zimmermann (2022).¹⁹ In each case, we provide evidence that the belief-based component plays an important role in the excess returns of these anomalies.

Proposition 3 shows a simple way to decompose the excess returns of an asset into the belief-based component – measured through the SBF – and the preference-based or risk-based component. By assuming log-normality, we can express the average excess return as

$$\log (E [R_{t+1}^e]) = Cov (-s_{t+1}, r_{t+1}^e) + Cov (-\tilde{m}_{t+1}, r_{t+1}^e). \quad (23)$$

To give an analogy for this decomposition, this exercise is similar to asset pricing papers that estimate an SDF to match one set of assets and then measure whether that SDF accounts for the excess returns of some other set of test assets. In this case, we are estimating a distortion s_{t+1} that accurately describes one set of data – subjective expectations of macroeconomic variables – and then testing whether this distortion can quantitatively account for the excess returns of the test assets.

At first glance, this approach may seem reminiscent of research arguing that future returns are correlated with qualitative measures of “sentiment.”²⁰ While our approach is related to this sentiment literature, it importantly addresses issues with sentiment and risk being correlated. Given a measure of sentiment that is correlated with future returns, there is always a concern whether sentiment is simply correlated with risk, e.g., sentiment may be low in a recession while risk is high.

Fortunately, decomposition (23) holds even if s_{t+1} and \tilde{m}_{t+1} are correlated. This is be-

¹⁹Details on the construction of annual returns are shown in Appendix E.

²⁰See Baker and Wurgler (2006) and Kumar and Lee (2006) for early examples in this literature and Bordalo et al. (2023) as a recent example arguing that *LTG* is correlated with future returns.

cause $Cov(-s_{t+1}, r_{t+1}^e)$ captures the magnitude of the agent’s biased beliefs and $Cov(-\tilde{m}_{t+1}, r_{t+1}^e)$ captures the magnitude of the agent’s preferences. Taking the example from Section I.C, if the agent exaggerates about the probability of low RGDP growth states by a factor of 2 relative to high growth states and prefers payoffs in low RGDP growth states 10 times as much as she prefers payoffs in high growth states, then $Cov(-s_{t+1}, r_{t+1}^e)$ will be much smaller than $Cov(-\tilde{m}_{t+1}, r_{t+1}^e)$. In this case, a researcher who only knows s_{t+1} and r_{t+1}^e would still correctly conclude that $Cov(-\tilde{m}_{t+1}, r_{t+1}^e)$ accounts for most of the excess return.

By addressing this issue of correlated beliefs and risk, our approach highlights the benefits of a quantitative SBF compared to qualitative measures of sentiment. Additionally, our estimated SBF does not depend on a specific model of sentiment. The SBF is directly measured from survey data and the objective covariance matrix and can potentially encompass two or more behavioral models of sentiment.

Our approach also builds on the previous work using analyst expectations for individual stock returns to study anomalies (Engelberg, McLean, and Pontiff, 2020; Jensen, 2023; De la O, Han, and Myers, 2024), and provides two useful features. First, since we do not require specific survey expectations about the assets in question, we remove the sample limitations associated with stock-level analyst return expectations, which are generally only available post-2000. Second, rather than using survey expectations for different sets of stocks for each anomaly, we show that a single SBF based on beliefs about just two macroeconomic variables is able to account for approximately half of all excess returns across many anomalies.

In a sense, our SBF captures the best of both worlds. It provides a single variable that can potentially account for many different anomalies, similar to the sentiment literature. At the same time, it maintains the quantitative benefits of previous work on analyst return expectations, as risk and biased beliefs can be distinguished through the use of covariance decompositions rather than correlations.

A. Fama-French factors

We first evaluate the four Fama-French anomalies.²¹ While it is hard to directly measure the risk component involving \tilde{m}_{t+1} , we can measure the belief-based component using our estimated SBF \hat{s}_{t+1} . We use the same estimated SBF \hat{s}_{t+1} as the previous sections, which is derived from the subjective expectations for RGDP growth and the T-bill rate and it is available from 1981Q1 to 2022Q2.

Figure 2 shows the average excess returns $\log(E[R_{j,t+1}^e])$ and the belief component denoted as $Cov(-\hat{s}_{t+1}, r_{j,t+1}^e)$ for each the four anomalies over the same sample. Despite not using any information about anomaly returns in the construction of \hat{s}_{t+1} , we see that $Cov(-\hat{s}_{t+1}, r_{j,t+1}^e)$ gives reasonable values for annual anomaly returns of 3.7pp to 8.8pp. For each anomaly, we find a positive $Cov(-\hat{s}_{t+1}, r_{j,t+1}^e)$ that is quantitatively large enough to explain the entire observed average excess return. In other words, the excess returns on these anomalies can be explained by the fact that these anomalies tend to pay off in states of the world that forecasters appear to underestimate.

Notably, we do find a mismatch between the average excess return and the belief component for the size anomaly, where $Cov(-\hat{s}_{t+1}, r_{j,t+1}^e)$ is several times larger than $\log(E[R_{j,t+1}^e])$. This may be due to the average return on the size anomaly being notably smaller than the other anomalies over this sample. Or, this could indicate a negative risk component $Cov(-\tilde{m}_{t+1}, r_{j,t+1}^e)$.

B. Behavioral factors

Since the above results suggest that some of the Fama-French factors can be explained with current survey evidence, it is worth illustrating the ability of our SBF \hat{s}_{t+1} to explain behavioral factors. Daniel, Hirshleifer, and Sun (2020) propose two behavioral factors that can span a significant subset of the Fama-French factors. They show that the FIN factor,

²¹All four long-short return portfolios are taken from Ken French's website and compounded into annual returns for the same sample as the distortion, 1981-2022. Appendix E contains details on the portfolio construction.

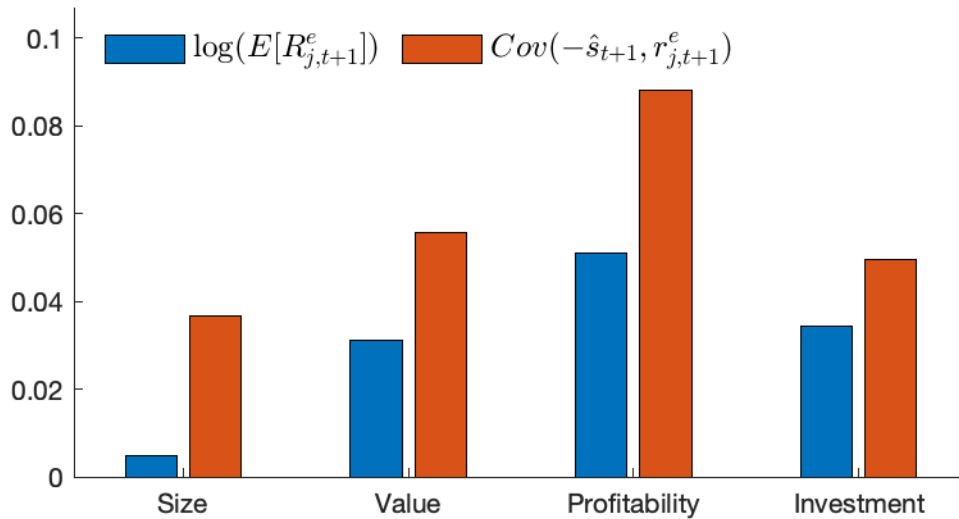


Figure 2. Fama-French factors and their covariance with the SBF. This figure shows the decomposition (23) of excess returns of each of the Fama-French cross-sectional factors into the belief-based component using the log SBF \hat{s}_{t+1} . The blue bars show the average excess return $\log(E[R_{t+1}^e])$ of the size, value, profitability, and investment factors respectively. The red bars show the portion of the average return attributable to the covariance of the excess return with the log SBF $Cov(-\hat{s}_{t+1}, r_{t+1}^e)$.

constructed as the return spread between recent issuers and repurchasers, together with the PEAD factor, constructed as the returns spread between firms with positive earnings surprises and firms with negative earnings surprises, together account for a wide range of anomalies. Both factors are motivated by behavioral stories, one related to the four-day response to earnings surprises and one related to long-horizon mispricing.

Given the behavioral motivation of these factors, we are interested in how well our SBF \hat{s}_{t+1} can explain these excess returns. In particular, our SBF is based on one-year expectations, which Daniel, Hirshleifer, and Sun (2020) would denote as “long-horizon.” The short-horizon factor (PEAD) is intended to capture “high frequency” biases.

Figure 3 shows the average excess returns of the behavioral factors FIN and PEAD and their covariance with our SBF $-\hat{s}_{t+1}$. We observe that our SBF explains the majority of the excess returns from the financing factor, accounting for 93% of the average excess returns generated by the factor. This evidence supports a behavioral explanation for the financing factor and provides evidence against the idea that the excess return on FIN is

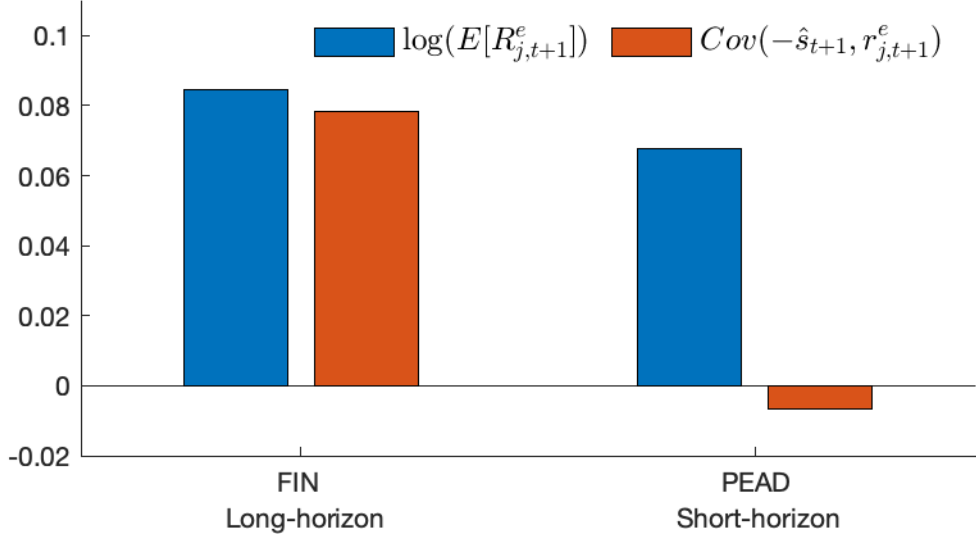


Figure 3. Behavioral factors and their covariance with the SBF. This figure shows the decomposition (23) of excess returns of each of the behavioral factors from Daniel, Hirshleifer, and Sun (2020) into the belief-based component using the log SBF \hat{s}_{t+1} . The blue bars show the average excess return $\log(E[R_{t+1}^e])$ of the PEAD (Post-Earning Announcement Drift) and FIN (Financing) factors. The red bars show the portion of the average return attributable to the covariance of the excess return with the log SBF $Cov(-\hat{s}_{t+1}, r_{t+1}^e)$.

mainly due to a rational risk premium. On the other hand, Figure 3 shows that our SBF has nearly no comovement with the short-horizon PEAD factor, with $Cov(-\hat{s}_{t+1}, r_{j,t+1}^e)$ being -0.67pp compared to the observed 6.8pp average excess return. This is consistent with the fact that our SBF \hat{s}_{t+1} is estimated from one-year expectations rather than high-frequency expectations. In this case, the average excess return would need to be explained by belief distortions not captured by our one-year subjective expectations or by a risk premium $Cov(-\tilde{m}_{t+1}, r_{j,t+1}^e)$.

C. Large set of anomalies

We next consider how the belief distortion performs when taken to a wide range of anomalies as constructed in Chen and Zimmermann (2022). There are 22 categories of anomalies available for our 1981Q1 to 2022Q2 sample. Each category generally contains multiple individual anomalies as there are typically multiple ways to measure the underlying economic

variable, e.g., four different ways to measure leverage. To reduce noise, we calculate the average $\log(E[R_{j,t+1}^e])$ and the average $Cov(-\hat{s}_{t+1}, r_{j,t+1}^e)$ across all anomalies j in each category.²² Figure 4 shows the results.²³

How can one summarize these findings across categories? One useful method is to measure how much differences in $\log(E[R_{j,t+1}^e])$ across categories is associated with differences in $Cov(-\hat{s}_{t+1}, r_{j,t+1}^e)$. Consider a simple two-equation regression framework,

$$Cov(-s_{t+1}, r_{j,t+1}^e) = \alpha_s + \gamma_s \log(E[R_{j,t+1}^e]) + \varepsilon_{j,s} \quad (24)$$

$$Cov(-\tilde{m}_{t+1}, r_{j,t+1}^e) = \alpha_m + \gamma_{\tilde{m}} \log(E[R_{j,t+1}^e]) + \varepsilon_{j,\tilde{m}}. \quad (25)$$

From equation (23), we know that

$$\gamma_s + \gamma_{\tilde{m}} = 1.$$

In other words, a one unit increase in $\log(E[R_{j,t+1}^e])$ must correspond to a one unit increase in $Cov(-s_{t+1}, r_{j,t+1}^e) + Cov(-\tilde{m}_{t+1}, r_{j,t+1}^e)$, and γ_s and $\gamma_{\tilde{m}}$ tell us whether it is primarily $Cov(-s_{t+1}, r_{j,t+1}^e)$ that increases or primarily $Cov(-\tilde{m}_{t+1}, r_{j,t+1}^e)$ that increases. Using our estimated $-\hat{s}_{t+1}$, Figure 4 shows that regression (24) gives a coefficient γ_s equal to 0.561 when each category is weighted equally. If we use a weighted regression based on the number of anomalies in each category, we find γ_s equal to 0.726. Both results point to the belief component and the risk component playing a non-trivial role, with either a 50/50 split or a 70/30 split between the two. This means that our single SBF \hat{s}_{t+1} based on beliefs for just two macroeconomic variables does appear to account for a notable amount of cross-sectional

²²To avoid categories having only a single anomaly, we combine some categories based on conceptual similarity. For example, we combine “Profitability” and “Profitability alt” into a single category given that “Profitability Alt” contains only a single anomaly. The only exception is “Size”, which only contains a single anomaly but is not combined with any other category given its historical importance.

²³Note that the trading strategy details differ between Fama and French (2015) and Chen and Zimmermann (2022). For example, Fama and French (2015) measure their profitability factor RMW using a bivariate sort on profitability and size and averaging the results across size. In comparison, Chen and Zimmermann (2022) use a univariate sort based solely on profitability. Because of implementation differences and the fact that Chen and Zimmermann (2022) provides multiple ways to define profitability, the results for categories like “profitability” and “investment” in Figure 4 may differ from the values for the individual anomalies in Figure 2. Importantly, both results use the same log SBF \hat{s}_{t+1} . Given that both methods are plausible ways to measure the profitability anomaly, we show both sets of results and highlight that under both methods we find a positive and quantitatively large comovement between the excess return and $-\hat{s}_{t+1}$.

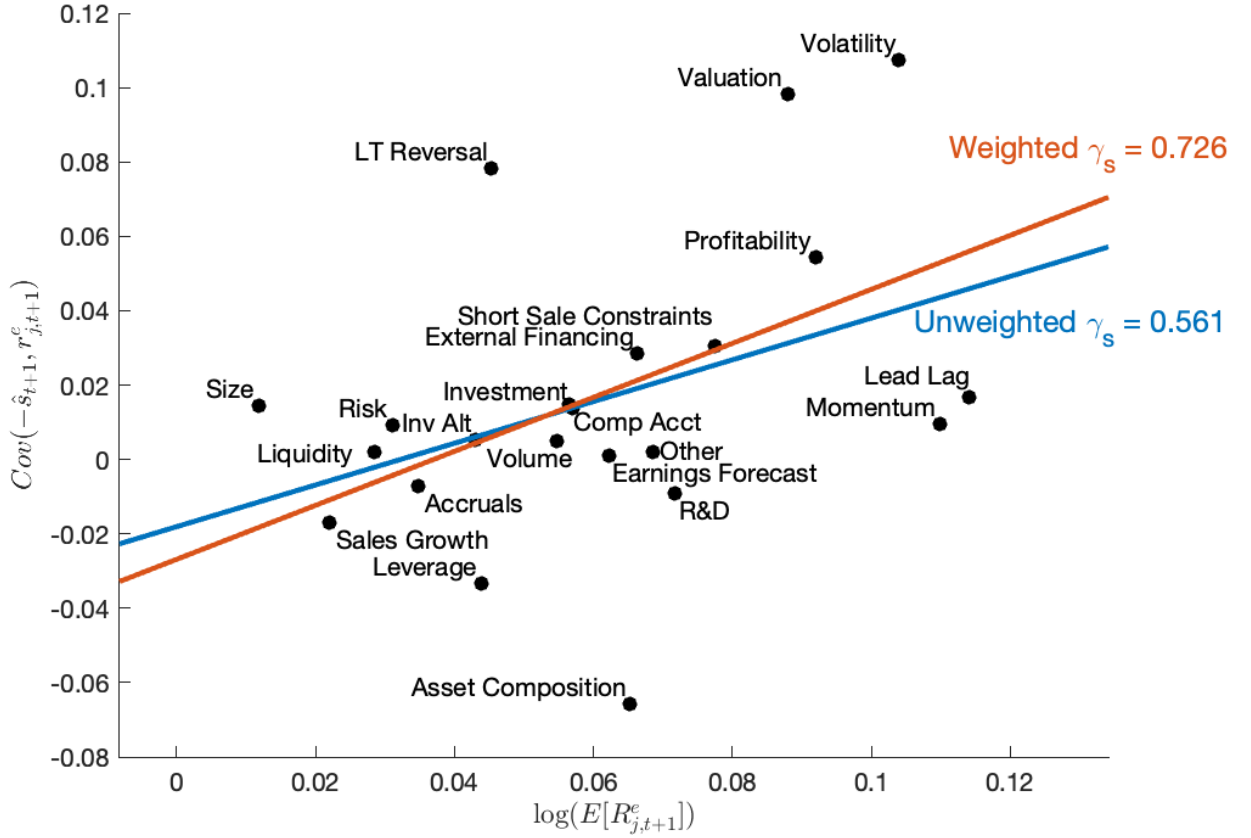


Figure 4. Anomalies and their covariance with the SBF. This figure shows the decomposition of excess returns of the Chen and Zimmermann (2022) anomalies into the belief-based component using the log SBF \hat{s}_{t+1} . The x-axis shows the average excess return $\log(E[R_{t+1}^e])$ of the anomaly portfolios. The y-axis shows $Cov(-\hat{s}_{t+1}, r_{t+1}^e)$, which measures how much of the average excess returns is attributable to the covariance with the log SBF \hat{s}_{t+1} . The line in blue shows the slope of a linear regression of $Cov(-\hat{s}_{t+1}, r_{t+1}^e)$ on $\log(E[R_{t+1}^e])$ as in equation (24). The line in red shows the slope of the same regression weighted by the number of anomalies on each category.

variation in excess returns, but there is still an equally large amount of variation remaining that can be attributed to risk or preferences.

V. Conclusion

Under general conditions, there is no mathematical distinction between behavioral economists attempting to characterize subjective expectations and financial economists attempting to characterize asset prices. Both problems can be framed as explaining expected outcomes

under a distorted probability distribution. While these two fields have largely developed separately, we argue that there is substantial potential for tools and data from each field to be applied to the other.

In this paper, we utilize tools from asset pricing to characterize subjective expectations. Just as financial economists are able to link many asset prices to a single SDF, we demonstrate that subjective expectations for many variables can be linked to a single SBF S_{t+1} . We then utilize data on subjective expectations to better understand asset prices. While there is only limited data on subjective expectations of returns for cross-sectional anomalies, we demonstrate that a SBF based on subjective expectations for RGDP growth and the T-bill rate goes a substantial way towards accounting for a wide range of cross-sectional anomaly returns.

Future work can continue to merge tools and data across these two fields. Techniques used to study asset prices and the SDF, such as Fama-MacBeth regressions, Hansen-Jagannathan bounds, and affine term-structure models, can potentially provide important insights for subjective expectations. Conversely, models of expectation formation, tests for biases such as Coibion and Gorodnichenko (2015), and additional data on subjective expectations can potentially be translated into wide-ranging implications for asset prices by characterizing the belief component S_{t+1} of the SDF that prices assets, $M_{t+1} = S_{t+1}\tilde{M}_{t+1}$.

References

- Adam, Klaus, and Stefan Nagel, 2023, Expectations data in asset pricing.
- Bachmann, Rüdiger, Giorgio Topa, and Wilbert van der Klaauw, 2023, Handbook of economic expectations, *Elsevier* .
- Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *The Journal of Finance* 61, 1645–1680.
- Bianchi, Francesco, Sydney Ludvigson, and Sai Ma, 2024, What hundreds of economic news events say about belief overreaction in the stock market, *Working Paper*, Johns Hopkins University.
- Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer, 2020, Overreaction in macroeconomic expectations, *American Economic Review* 110, 2748–2782.
- Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2023, Belief overreaction and stock market puzzles, *Journal of Political Economy* .
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2018, Diagnostic expectations and credit cycles, *Journal of Finance* 73, 199–227.
- Brav, Alon, and J B Heaton, 2002, Competing theories of financial anomalies, *The Review of Financial Studies* 15, 575–606.
- Chen, Andrew Y., and Tom Zimmermann, 2022, Open source cross-sectional asset pricing, *Critical Finance Review* 27, 207–264.
- Chen, Xiaohong, Lars Peter Hansen, and Peter G Hansen, 2020, Robust identification of investor beliefs, *Proceedings of the National Academy of Sciences* 117, 33130–33140.
- Cieslak, Anna, 2018, Short-rate expectations and unexpected returns in treasury bonds, *Review of Financial Studies* 31, 3265–3306.

- Cochrane, John H, 2005, *Asset Pricing*, revised edition edition (Princeton University Press).
- Coibion, Olivier, and Yuriy Gorodnichenko, 2015, Information rigidity and the expectations formation process: A simple framework and new facts, *American Economic Review* 105, 2644–2678.
- Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer, 2016, Parameter learning in general equilibrium, *American Economic Review* 106, 664–698.
- Daniel, Kent, David Hirshleifer, and Lin Sun, 2020, Short- and long-horizon behavioral factors.
- De la O, Ricardo, Xiao Han, and Sean Myers, 2024, The cross-section of subjective expectations: Understanding prices and anomalies, *Working Paper*, SSRN.
- Engelberg, Joseph, R. David McLean, and Jeffrey Pontiff, 2020, Analysts and anomalies, *Journal of Accounting and Economics* 69, 101249.
- Fama, Eugene F, and Kenneth R French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.
- Haubrich, Joseph, George Pennacchi, and Peter Ritchken, 2012, Inflation expectations, real rates, and risk premia: Evidence from inflation swaps, *The Review of Financial Studies* 25, 1588–1629.
- Huang, Dashan, Fuwei Jiang, Jun Tu, and Guofu Zhou, 2015, Investor sentiment aligned: A powerful predictor of stock returns, *The Review of Financial Studies* 28, 791–837.
- Jensen, Theis Ingerslev, 2023, Subjective risk and return, *Working Paper*, Yale University.
- Klein, Lawrence R, 2018, *Klein's Last Quarterly Econometric Model Of The United States - Wharton Quarterly Econometric Model: Mark 10*, volume 6 (World Scientific).

- Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2018, Interpreting factor models, *The Journal of Finance* 73, 1183–1223.
- Kuchler, Theresa, and Basit Zafar, 2019, Personal experiences and expectations about aggregate outcomes, *The Journal of Finance* 74, 2491–2542.
- Kumar, Alok, and Charles M C Lee, 2006, Retail investor sentiment and return comovements, *The Journal of Finance* 61, 2451–2486.
- Link, Sebastian, Andreas Peichl, Christopher Roth, and Johannes Wohlfart, 2023, Information frictions among firms and households, *Journal of Monetary Economics* 135, 99–115.
- Malmendier, Ulrike, and Stefan Nagel, 2016, Learning from inflation experiences, *Quarterly Journal of Economics* 131, 53–87.
- Mueller, Andreas I, Johannes Spinnewijn, and Giorgio Topa, 2021, Job seekers’ perceptions and employment prospects: Heterogeneity, duration dependence, and bias, *American Economic Review* 111, 324–363.
- Piazzesi, Monika, Juliana Salomao, and Martin Schneider, 2015, Trend and cycle in bond premia, *Working Paper*, Stanford University.
- Stambaugh, Robert F, Jianfeng Yu, and Yu Yuan, 2024, The short of it: Investor sentiment and anomalies, *Journal of Financial Economics* 144, 84–112.
- Stark, Tom, 2013, Spf panelists’ forecasting methods: A note on the aggregate results of a november 2009 special survey.

Appendix

A. Proofs

Proof of Proposition 1: Because subjective expectations are coherent, we know that $E_{i,t}^*[\cdot]$ is a continuous linear functional. We define the inner product operator as $\langle Y_{t+1}, Z_{t+1} \rangle \equiv E_t[Y_{t+1}Z_{t+1}]$. By the Reisz representation theorem, there exists $S_{i,t+1}$ such that $E_{i,t}^*[X_{t+1}] = \langle X_{t+1}, S_{i,t+1} \rangle = E_t[S_{i,t+1}X_{t+1}]$. To show that $E_t[S_{i,t+1}] = 1$, we just use the fact that $E_{i,t}^*[1] = E_t[S_{i,t+1}] = 1$.

Proof of Lemma 1: For each individual i , there exists $S_{i,t+1}$ such that for any X_{t+1} ,

$$E_{i,t}^*[X_{t+1}] = E_t[S_{i,t+1}X_{t+1}].$$

Our goal is to show that $E_t^*[X_{t+1}] = E_t[S_{t+1}X_{t+1}]$. Given the definition of $E_t^*[X_{t+1}]$ and S_{t+1} , we have

$$\begin{aligned} E_t^*[X_{t+1}] &= \frac{1}{I} \sum_i E_{i,t}^*[X_{t+1}] = \frac{1}{I} \sum_i E_t[S_{i,t+1}X_{t+1}] \\ &= E_t \left[\frac{1}{I} \sum_i S_{i,t+1}X_{t+1} \right] = E_t[S_{t+1}X_{t+1}]. \end{aligned}$$

Proof of Lemma 2: As in the last lemma, for each individual i , there exists there exists $S_{i,t+1}$ such that for any X_{t+1} ,

$$E_{i,t}^*[X_{t+1}] = E_t[S_{i,t+1}X_{t+1}].$$

Given that $E_{i,t}^*[\tilde{M}_{t+1}X_{t+1}] = P_t$ for all i , our goal is to show that $E_t^*[\tilde{M}_{t+1}X_{t+1}] = E_t[S_{t+1}\tilde{M}_{t+1}X_{t+1}] = P_t$. By the definition of \tilde{M}_{t+1} and S_{t+1} , we know that

$$\begin{aligned} P_t &= \frac{1}{I} \sum_i E_{i,t}^*[\tilde{M}_{t+1}X_{t+1}] = \frac{1}{I} \sum_i E_t[S_{i,t+1}\tilde{M}_{i,t+1}X_{t+1}] \\ &= E_t \left[\frac{1}{I} \sum_i S_{i,t+1}\tilde{M}_{i,t+1}X_{t+1} \right] = E_t \left[\left(\frac{1}{I} \sum_i S_{i,t+1} \right) \tilde{M}_{t+1}X_{t+1} \right] \\ &= E_t[S_{t+1}\tilde{M}_{t+1}X_{t+1}]. \end{aligned}$$

From Lemma 1, we know that $E_t^* [\tilde{M}_{t+1} X_{t+1}] = E_t [S_{t+1} \tilde{M}_{t+1} X_{t+1}]$.

Proof of Proposition 2: Let $\phi(\cdot)$ denote the standard normal pdf. The objective distribution is $x_{t+1} \sim N(E_t[x_{t+1}], \Sigma)$ and the agent subjectively believes $x_{t+1} \sim N(E_t^*[x_{t+1}], \Sigma)$. This means that the ratio of the subjective probability to the objective probability can be defined as

$$\begin{aligned}
S_{t+1} &= \frac{\phi((x_{t+1} - E_t^*[x_{t+1}])' \Sigma^{-1} (x_{t+1} - E_t^*[x_{t+1}]))}{\phi((x_{t+1} - E_t[x_{t+1}])' \Sigma^{-1} (x_{t+1} - E_t[x_{t+1}]))} \\
\Rightarrow s_{t+1} &= -\frac{1}{2} (x_{t+1} - E_t^*[x_{t+1}])' \Sigma^{-1} (x_{t+1} - E_t^*[x_{t+1}]) + \frac{1}{2} (x_{t+1} - E_t[x_{t+1}])' \Sigma^{-1} (x_{t+1} - E_t[x_{t+1}]) \\
&= -\frac{1}{2} (\eta_{t+1} + E_t[x_{t+1}] - E_t^*[x_{t+1}])' \Sigma^{-1} (\eta_{t+1} + E_t[x_{t+1}] - E_t^*[x_{t+1}]) + \frac{1}{2} \eta_{t+1}' \Sigma^{-1} \eta_{t+1} \\
&= -\frac{1}{2} (E_t[x_{t+1}] - E_t^*[x_{t+1}])' \Sigma^{-1} (E_t[x_{t+1}] - E_t^*[x_{t+1}]) - (E_t[x_{t+1}] - E_t^*[x_{t+1}])' \Sigma^{-1} \eta_{t+1} \\
&= -\frac{1}{2} \beta_t' \Sigma \beta_t + \beta_t' \eta_{t+1}
\end{aligned}$$

given the definition of β_t .

Proof of Corollary 1: Again, let $\phi(\cdot)$ denote the standard normal pdf. The logic of the proof is virtually identical to the proof of Proposition 1. The objective distribution is $x_{t+1} \sim N(E_t[x_{t+1}], \Sigma_t)$ and the agent subjectively believes $x_{t+1} \sim N(E_t^*[x_{t+1}], \Sigma_t^*)$. This means that the ratio of the subjective probability to the objective probability can be defined

as

$$\begin{aligned}
S_{t+1} &= \frac{\phi\left((x_{t+1} - E_t^*[x_{t+1}])' \Sigma_t^{*-1} (x_{t+1} - E_t^*[x_{t+1}])\right)}{\phi\left((x_{t+1} - E_t[x_{t+1}])' \Sigma_t^{-1} (x_{t+1} - E_t[x_{t+1}])\right)} \\
\Rightarrow s_{t+1} &= -\frac{1}{2} \log(\det(\Sigma_t^*)) - \frac{1}{2} (x_{t+1} - E_t^*[x_{t+1}])' \Sigma_t^{*-1} (x_{t+1} - E_t^*[x_{t+1}]) \\
&\quad + \frac{1}{2} \log(\det(\Sigma_t)) + \frac{1}{2} (x_{t+1} - E_t[x_{t+1}])' \Sigma_t^{-1} (x_{t+1} - E_t[x_{t+1}]) \\
&= \frac{1}{2} \log\left(\frac{\det(\Sigma_t)}{\det(\Sigma_t^*)}\right) - \frac{1}{2} (\eta_{t+1} + E_t[x_{t+1}] - E_t^*[x_{t+1}])' \Sigma_t^{*-1} (\eta_{t+1} + E_t[x_{t+1}] - E_t^*[x_{t+1}]) \\
&\quad + \frac{1}{2} \eta_{t+1}' \Sigma_t^{-1} \eta_{t+1} \\
&= \frac{1}{2} \log\left(\frac{\det(\Sigma_t)}{\det(\Sigma_t^*)}\right) + \frac{1}{2} \eta_{t+1}' (\Sigma_t^{-1} - \Sigma_t^{*-1}) \eta_{t+1} \\
&\quad - \frac{1}{2} (E_t[x_{t+1}] - E_t^*[x_{t+1}])' \Sigma_t^{*-1} (E_t[x_{t+1}] - E_t^*[x_{t+1}]) - (E_t[x_{t+1}] - E_t^*[x_{t+1}])' \Sigma_t^{*-1} \eta_{t+1} \\
&= \frac{1}{2} \log\left(\frac{\det(\Sigma_t)}{\det(\Sigma_t^*)}\right) - \frac{1}{2} \beta_t' \Sigma_t^* \beta_t + \frac{1}{2} \eta_{t+1}' (\Sigma_t^{-1} - \Sigma_t^{*-1}) \eta_{t+1} + \beta_t' \eta_{t+1}
\end{aligned}$$

given the definition of β_t .

Proof of Proposition 3: The distributions for the three variables are

$$\begin{aligned}
r_{t+1}^e &= \mu_r - \frac{1}{2} \sigma_r^2 + \sigma_r \varepsilon_{t+1}^r \\
s_{t+1} &= \mu_{s,t} - \frac{1}{2} \sigma_{s,t}^2 + \sigma_{s,t} \varepsilon_{t+1}^s \\
\tilde{m}_{t+1} &= \mu_{m,t} - \frac{1}{2} \sigma_{m,t}^2 + \sigma_{m,t} \varepsilon_{t+1}^m
\end{aligned}$$

where $(\varepsilon_{t+1}^r, \varepsilon_{t+1}^s, \varepsilon_{t+1}^m)$ are potentially correlated Gaussian shocks. We have that

$$\begin{aligned}
1 &= E_t^* \left[\tilde{M}_{t+1} R_{t+1} \right] \\
&= E_t \left[S_{t+1} \tilde{M}_{t+1} R_{t+1} \right] \\
&= E_t \left[S_{t+1} \tilde{M}_{t+1} R_{t+1}^e \right] R_t^f \\
&= E_t \left[S_{t+1} \tilde{M}_{t+1} \right] E_t \left[R_{t+1}^e \right] \exp\left(\text{Cov}_t(s_{t+1} + \tilde{m}_{t+1}, r_{t+1}^e)\right) R_t^f \\
&= \exp\left(\mu_r + \text{Cov}_t(s_{t+1} + \tilde{m}_{t+1}, r_{t+1}^e)\right)
\end{aligned}$$

where the second to last line uses the fact that $E_t^* \left[\tilde{M}_{t+1} R_t^f \right] = E_t \left[S_{t+1} \tilde{M}_{t+1} \right] R_t^f = 1$.

Taking logs and rearranging shows that

$$\begin{aligned}
 \mu_r &= Cov_t(-s_{t+1} - \tilde{m}_{t+1}, r_{t+1}^e) \\
 &= E[Cov_t(-s_{t+1} - \tilde{m}_{t+1}, r_{t+1}^e)] \\
 &= Cov(-s_{t+1} - \tilde{m}_{t+1}, r_{t+1}^e).
 \end{aligned}$$

B. Details on Data Construction

The Survey of Professional Forecasters is a quarterly survey currently administered by the Federal Reserve Bank of Philadelphia. The survey elicits forecasts for a host of economic variables from professional forecasters and reports individual-level forecasts. In each survey, forecasters are asked to provide point forecasts for the current quarter and for each of the next four quarters. Variables are sometimes added to the survey following changes in the macroeconomy. The earliest survey dates back to 1968:Q4, which included forecasts for *ngdp*, *pgdp*, *cprof*, *unemp*, *indprod*, *housing*, and *rgdp*. A round of significant changes occurred in the third quarter of 1981, when the NBER added ten more variables to the survey. The new survey variables included since the 1981:Q3 survey are *tbill*, *aaa*, *rconsum*, *rnresin*, *rresinv*, *rgf*, *rgsl*, *rcbi*, *rexporth*, and *cpi*. In our analysis, we use all variables for which we have survey data since 1981:Q3. Note that there is some redundancy in the variables due to the fact the nominal GDP, real GDP, and the GDP deflator are all forecasted variables. To ensure that our ability to explain many economic forecasts with a single SBF is not due to redundancy in the economic variables, we exclude both nominal GDP and the GDP deflator and only keep real GDP and a single inflation measure (CPI). Table AI reports the resulting 15 variables.

The second source of survey data we use is the Blue Chip Financial Forecasts. The survey elicits forecasts for economic variables from top analysts from manufacturers, banks, insurance companies, and brokerage firms. In each survey, forecasters are similarly asked to provide point forecasts for interest rate variables and quarter-over-quarter growth rate forecasts for economic variables. Similar to the Survey of Professional Forecasters, the Blue

Table AI

List of Variables from the Survey of Professional Forecasters

This table lists the 15 Survey of Professional Forecasters variables for which one-year survey forecast data exists since 1981Q3.

Variable Name	Variable Description
<i>rgdp</i>	Real GDP
<i>rcon</i>	Real personal consumption expenditures
<i>cpi</i>	CPI inflation rate
<i>unempl</i>	Unemployment rate
<i>indp</i>	Index of Industrial Production
<i>tbill</i>	Three-month treasury bill rate
<i>aaa</i>	Moody's Aaa corporate bond yield
<i>rnresin</i>	Real nonresidential fixed investment
<i>rresinv</i>	Real residential fixed investment
<i>rgf</i>	Real federal government consumption and gross investment
<i>rgsl</i>	Real state and local government consumption and gross investment
<i>housing</i>	Level of housing starts
<i>rcbi</i>	Real net change in private inventories
<i>rexport</i>	Chain-weighted real net exports
<i>cprof</i>	Level of nominal corporate profits

Chip survey has variables added, and sometimes removed, particularly in the earlier years of the survey. A round of significant changes occurred in 1988:Q1 when a host of interest rates were added to the survey: *tbill-1yr*, *tbill-6m*, *tnote-10yr*, *tnote-2yr*, and *libor-3m*. Earlier variables available include the prime rate (*prime*), the fed funds rate (*fedfunds*), *munis*, *aaa*, mortgage rate (*mortgage*), among others. In our analysis, we use all variables for which we have survey data since 1988:Q1. We exclude from our analysis any variables which were previously available, but subsequently removed from the survey (for instance *munis* forecasts are no longer solicited as a part of the Blue Chip Survey).²⁴ In order to illustrate that our SBF – which is estimated from data on the Survey of Professional Forecasters – can explain additional subjective expectations, we exclude any variables which are also included in the Survey of Professional Forecasters. Table AII shows the resulting 9 Blue Chip variables.

²⁴We exclude 1-month commercial paper given the lack of available realized data before 1997. We exclude the 30-year treasury rate because the survey removed this variable from 2002Q1-2006Q1 and instead asked for forecasts of “long-term average” Treasury yields. This change not only introduces potential inconsistency in the forecast term structure but also complicates the identification of an appropriate realized counterpart for such forecasts.

Table AII

List of Variables from Blue Chip

This table lists all Blue Chip variables for which one-year survey forecast data exists since 1988Q1 and are not sampled in our Survey of Professional Forecasts dataset shown in Table AI. We consider survey forecasts for these variables from 1988:Q1 to 2022:Q2.

Variable Name	Variable Description
<i>prime</i>	Prime Bank Rate
<i>fedfunds</i>	Federal Funds Rate
<i>mortgage</i>	Home Mortgage Rate
<i>tbill-1yr</i>	1 Year Treasury Bill Rate
<i>tbill-6m</i>	6 Months Treasury Bill Rate
<i>tnote-10yr</i>	10 Year Treasury Note Rate
<i>tnote-5yr</i>	5 Year Treasury Note Rate
<i>tnote-2yr</i>	2 Year Treasury Note Rate
<i>libor-3m</i>	3 Month Libor Rate (SOFR Rate from 2022:Q1)

We use four-quarter ahead survey forecasts. For interest rate variables and the unemployment rate, we consider the point forecasts made for quarter $t+4$ at time t when analysts are filling out the survey. For *cpi*, we consider the annual forecasted growth by calculating the geometric mean of forecasted quarterly inflation over the next four quarters. For *rgdp*, *rcon*, *indp*, *rnresin*, *rresinv*, *rgf*, *rgsl*, *housing*, *cprof*, we convert point forecasts into implied growth rates. For example, for *rgdp*, we compute the forecaster's point forecast for $rgdp_{t+4}$ divided by the initial release of the $rgdp_{t-1}$ available in quarter t . The choice of using the first release of the value of $rgdp_{t-1}$ in quarter t ensures we align with the information set of forecasters as they are filling out the survey. For net exports (*rexp*) and change in private inventories (*rcbi*), which can potentially be zero or negative, we divide the forecasted value by $rgdp_{t-1}$ to make the variable stationary.

We use standard sources for realized values of macroeconomic and financial variables. For US interest rate variables, we use data reported by the Federal Reserve Bank of St. Louis. For the 3-month Libor series, we source from the ICE Benchmark Administration. For each interest rate variable, the realized value is calculated as the average value over the quarter that is being forecasted (i.e., it is a quarterly average rather than the interest rate on the last

day of the quarter). For the remaining variables, we use real-time data maintained by the Federal Reserve Bank of Philadelphia. Specifically, we use the initial release of the realized value of the macroeconomic variable in $t + 4$ made available in quarter $t + 5$. To compute the realized growth rate for the macroeconomic variable, we divide the $t + 4$ realized value by the $t - 1$ value, also released in quarter $t + 5$ for the sake of consistency. For example, we calculate the realized growth as the BEA's estimate in 2001Q3 for 2001Q2 real GDP divided by the BEA's estimate in 2001Q3 for 2000Q1 real GDP. This is equivalent to saying that we use the BEA's estimate in 2001Q3 for real GDP growth from 2000Q1 to 2002Q2.

C. *The behavior of the main biases and the estimated log SBF over time*

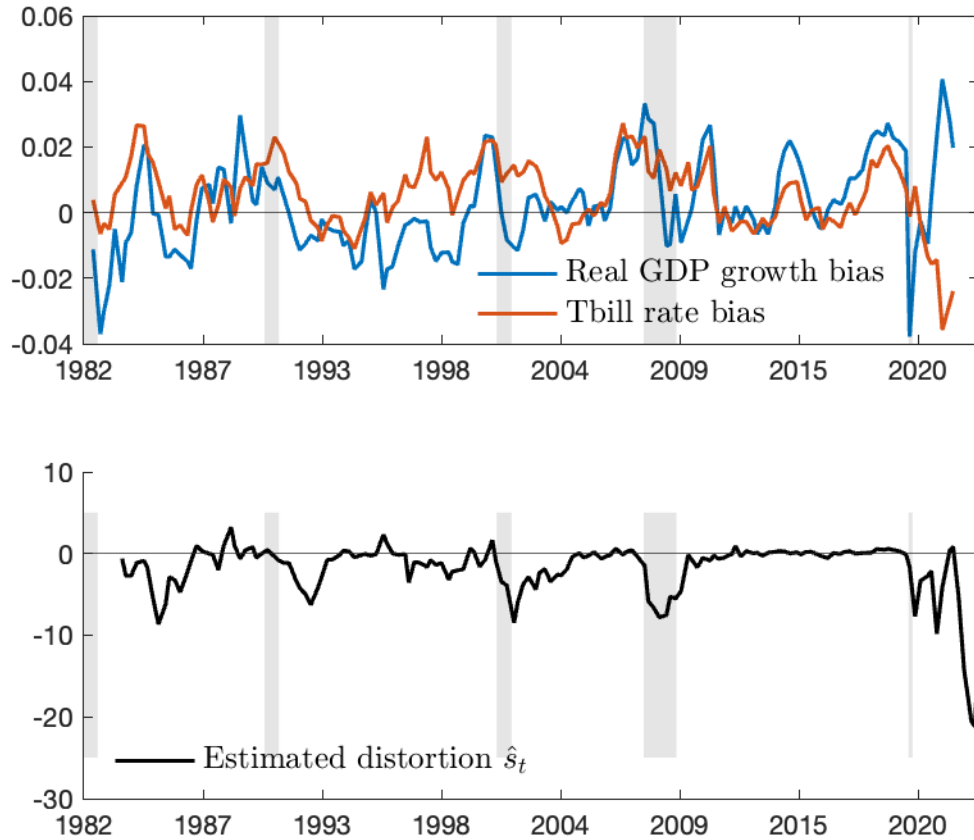


Figure A1. Biases and estimated log SBF over time. The top figure shows the RGDP growth and T-bill rate biases $E_t^*[\hat{x}_{t+1}] - E_t[\hat{x}_{t+1}]$ over the main sample. The bottom figure shows the log SBF \hat{s}_t estimated from these two biases using equation (12). Note that the sample for \hat{s}_t starts one year after the sample for $E_t^*[\hat{x}_{t+1}] - E_t[\hat{x}_{t+1}]$, as $E_t^*[\hat{x}_{t+1}] - E_t[\hat{x}_{t+1}]$ are combined with the future shocks $\hat{\eta}_{t+1}$ to calculate the realized value for the log SBF.

The top panel of Figure A1 shows the time series for the gap between statistical and subjective expectations for RGDP growth and the T-bill rate ($E_t^*[\hat{x}_{t+1}] - E_t[\hat{x}_{t+1}]$). In general, the gap between statistical and subjective expectations for RGDP growth and the T-bill rate are weakly correlated (0.28). However, they diverge in the post-Covid period, with forecasters being pessimistic about the recovery of output (i.e., RGDP growth) and understating interest rate increases. Compared to statistical expectations, subjective expectations of RGDP

growth and the T-bill rate are generally too high leading into recessions.

The bottom panel of the figure shows the realized time series for the estimated log SBF \hat{s}_t . We see that \hat{s}_t drops during most recessions, highlighting that these events are largely unexpected. Further, we see that \hat{s}_t is highly negative during the post-Covid period. Again, this highlights that forecasters understated the probability of a rapid economic recovery coupled with large interest rate increases.

D. *Alternative variable pairs to estimate the distortion*

To estimate the log SBF \hat{s}_{t+1} , we choose a subset of the variables in Survey of Professional Forecasters. Specifically, we include two variables in \hat{x}_{t+1} , RGDP growth and the T-bill rate, because of their economic significance and because they lie near the center of the two large clusters of variables in Figure 1. However, we could consider alternative pairs of variables.

In this section, we consider all possible pairs of variables \hat{x}_{t+1} . For each pair of variables, we estimate the log SBF \hat{s}_{t+1} that explains the subjective expectations for those two variables and then evaluate how well the synthetic biases $\hat{E}_t^* [x_{t+1}] - E_t [x_{t+1}]$ based on \hat{s}_{t+1} explain the survey biases $E_t^* [x_{t+1}] - E_t [x_{t+1}]$ for the full set of 15 variables. Table AIII shows the 10 pairs of variables that deliver the highest average R^2 . Out of the 210 possible pairs, we find that only three pairs perform better than RGDP growth and the T-bill rate and that the differences are quite small, i.e. an average R^2 of 48.2% instead of 47.2%.

Similar to our main exercise, can compare the results of each pair of variables to the regression

$$E_t^* [x_{t+1}] - E_t [x_{t+1}] = \alpha + \Gamma (E_t^* [\tilde{x}_{t+1}] - E_t [\tilde{x}_{t+1}]) + \varepsilon_t. \quad (\text{A1})$$

This specification represents the best linear prediction one can achieve using each pair of variables, and thus, provides an upper bound of how much one can explain *given these two variables*. Column 2 shows the R^2 of those linear predictions and shows that the average R^2 produced by the estimated \hat{s}_{t+1} is quite close to these upper bounds. Columns 3 and 4 show

Table AIII

Synthetic expectations using an alternative pair of variables.

This table evaluates the ability of synthetic expectations to explain movements in the Survey of Professional Forecasters variables. For each row, a set of synthetic expectations is constructed using a log SBF \hat{s}_{t+1} formed from a different pair of variables. The table shows the top ten pairs of variables ranked according to the average R^2 of the regression of the survey biases on the synthetic biases $\hat{E}_t^*[x_{j,t+1}] - E_t[x_{j,t+1}]$, shown in Column 1. For comparison, Column 2 shows the average R^2 of the best linear predictor of $E_t^*[x_{j,t+1}] - E_t[x_{j,t+1}]$ using the individual biases coming from equation (16) and Columns 3 and 4 show the explanatory power of the first two principal components of the 15 biases $E_t^*[x_{j,t+1}] - E_t[x_{j,t+1}]$. Columns 5 and 6 show the pair of variables used for each alternative log SBF \hat{s}_{t+1} .

Rank	$E_t^*[x_{t+1}] - E_t[x_{t+1}]$				Variables	
	$\hat{E}_t^*[x_{t+1}] - E_t[x_{t+1}]$	Best Linear Predictor	PC-1	PC-2		
1	48.2	54.0	43.0	63.8	<i>indp</i>	<i>rresinv</i>
2	48.0	50.3	43.0	63.8	<i>rcon</i>	<i>rcbi</i>
3	47.5	49.6	43.0	63.8	<i>rgdp</i>	<i>cpi</i>
4	47.2	52.3	43.0	63.8	<i>rgdp</i>	<i>tbill</i>
5	46.5	48.6	43.0	63.8	<i>indp</i>	<i>rexpport</i>
6	46.2	48.8	43.0	63.8	<i>rgdp</i>	<i>rcbi</i>
7	46.0	54.4	43.0	63.8	<i>unemp</i>	<i>rresinv</i>
8	45.4	50.4	43.0	63.8	<i>rnresin</i>	<i>rresinv</i>
9	45.2	51.7	43.0	63.8	<i>rgdp</i>	<i>rresinv</i>
10	45.0	48.3	43.0	63.8	<i>unemp</i>	<i>rexpport</i>

the general upper bound based on principal components analysis (PCA). For any of the ten variable pairs, the log SBF performs better than the first principal component of these series and captures roughly three fourths of the maximum possible R^2 , 63.8%.

E. Treatment of anomaly returns

For the three sets of anomalies (Fama and French, 2015; Daniel, Hirshleifer, and Sun, 2020; Chen and Zimmermann, 2022) analyzed in Section IV, we obtain the monthly returns from each of the authors' websites. Using geometric averages, we annualize the excess returns for each anomaly and annualize the risk-free rate provided by Fama.

Given that these data represent excess level returns and our theory deals with R_{t+1}/R_t^f (i.e., the exponential of excess log returns), we measure R_{t+1} as the return on a strategy that invests \$1 in the risk-free bond, \$1 in the long end of the anomaly and -\$1 in the short end of the anomaly. This gives $R_{t+1} = R_t^f + R_{anom,t+1}^e$, where $R_{anom,t+1}^e$ is the excess level return directly taken from the authors, and ensures that our measured $R_{t+1}^e \equiv R_{t+1}/R_t^f$ is simply $1 + R_{anom,t+1}^e/R_t^f$.

For the DHS behavioral anomalies, we only have returns data up to 2018 and therefore evaluate the performance of our SBF over this sample (1981-2018). For the Chen-Zimmerman anomalies, we consider all anomalies with returns over the sample of our SBF (1981-2022). This results in 176 anomalies assigned to 32 broader anomaly categories by the authors' original classification. Given our goal is to calculate the average $\log(E[R_{j,t+1}^e])$ and the average $Cov(-\hat{s}_{t+1}, r_{j,t+1}^e)$ across all anomalies j in each category, we reassign anomalies which belong to groups with three or fewer anomalies.²⁵ For example, *Cash Flow Risk* and *Default Risk*, which both contain one anomaly, are grouped together with the *Risk* category. Table AIV shows the categories of anomalies in the original paper as well as the number of anomalies available for our sample which fall under each category. Column 3 of Table AIV shows the assigned category based on conceptual similarity.

²⁵We set the threshold at 3 or fewer anomalies as 5 of the 32 categories have exactly 4 anomalies and we

Table AIV

Chen-Zimmerman Anomalies and Categories.

This table shows the broader set of Chen and Zimmermann (2022) anomalies which overlap with the sample of our log SBF \hat{s}_{t+1} from 1981 to 2022. Column 1 shows the anomaly categories in Chen and Zimmermann (2022). Column 2 shows the number of anomaly returns assigned to those categories. Column 3 shows the reassignment of anomalies which belong to categories with three or fewer anomaly returns. This results in a total of 176 anomalies assigned to 22 categories.

CZ2022 Category	# Anomalies	Assigned Category	# Anomalies
Accruals	4	Accruals	4
Asset Composition	5	Asset Composition	5
Risk	6	Risk	8
Cash Flow Risk	1	=	
Default Risk	1	=	
Composite Accounting	4	Composite Accounting	4
Earnings Forecast	5	Earnings Forecast	9
Earnings Event	1	=	
Earnings Growth	3	=	
External Financing	12	External Financing	12
Investment Alt	9	Investment Alt	9
Investment	8	Investment	11
Investment Growth	3	=	
Lead Lag	4	Lead Lag	4
Leverage	4	Leverage	4
Liquidity	8	Liquidity	8
Long Term Reversal	6	Long Term Reversal	6
Momentum	9	Momentum	9
Profitability	8	Profitability	9
Profitability Alt	1	=	
R&D	5	R&D	5
Sales Growth	6	Sales Growth	6
Short Sale Constraints	5	Short Sale Constraints	5
Size	1	Size	1
Valuation	17	Valuation	17
Volatility	5	Volatility	5
Volume	4	Volume	4
Other	24	Other	31
Info Proxy	1	=	
Ownership	2	=	
Payout Indicator	3	=	
Short Term Reversal	1	=	
Total	176	Total	176

F. Using log-normal distributions rather than normal distributions

For our main results, we assume that the variables covered in the Survey of Professional Forecasters and the Blue Chip survey are normally distributed. However, our approach also can be fruitfully applied when variables are log-normally distributed, meaning that future researchers have the freedom to choose whichever assumption is best-suited for their setting. Overall, we find almost identical results if we assume the variables are log-normally distributed. To highlight this, Figure A2 shows the SBF estimated in this section assuming log-normality (red) is nearly identical to the SBF estimated assuming normality from our main analysis (blue).

want to limit the number of categories that are being reassigned.

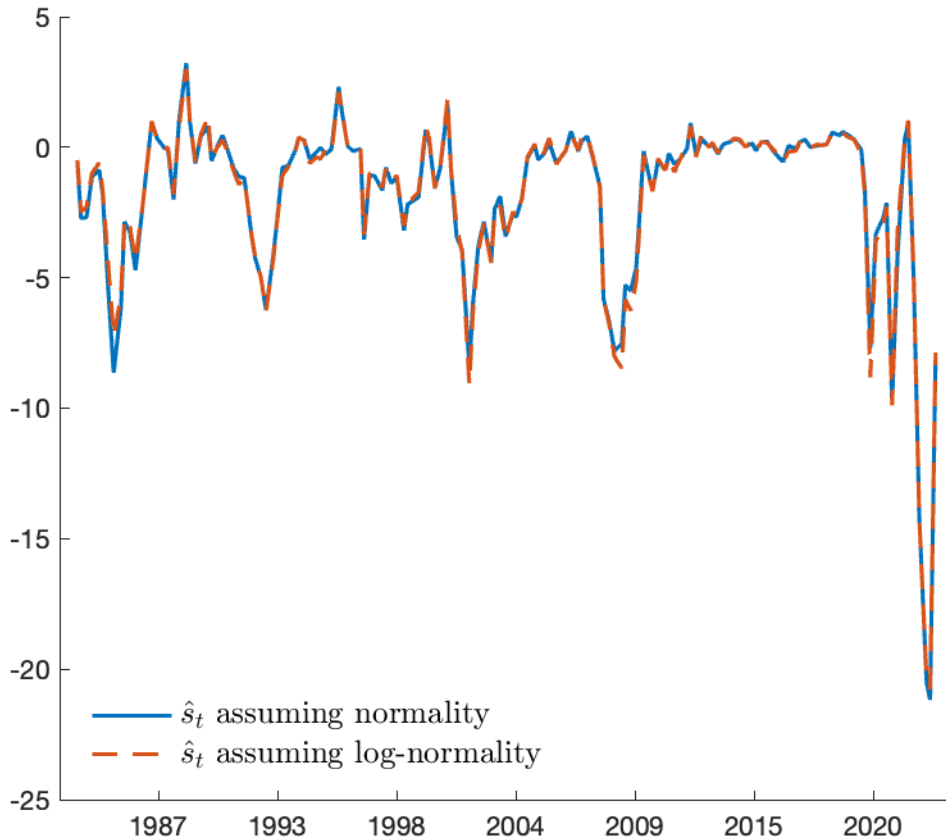


Figure A2. Estimated SBF. This figure compares the estimated SBF when we assume that the Survey of Professional Forecasters variables are objectively and subjectively normally distributed and when we assume they are objectively and subjectively log-normally distributed. The blue solid line shows the estimated SBF assuming normality, which is the SBF used in the main body of the paper. The red dashed line shows the estimated SBF assuming log-normality.

Just as in the main analysis, we start by attempting to condense the Survey of Professional Forecasters data. Of the 15 variables, most are growth rates or interest rates, meaning that they are already well-suited for a log-normal distribution. For net exports (*rexport*) and change in private inventories (*rcbi*), which can potentially be zero or negative, we divide the forecasted value by $rgdp_{t-1}$ and add 1 to make the variable stationary and suitable to the log-normal assumptions. For unemployment, we consider 1 minus the unemployment rate, which we simply refer to as *empl*, so that the value is closer to one and less impacted by the log transformation.

Given the 15 log-normally distributed variables X_{t+1} , we calculate statistical expectations using a VAR(1) model for $x_{t+1} \equiv \log(X_{t+1})$,

$$x_{t+1} = A + B \begin{pmatrix} x_t \\ \log(E_t^*[X_{t+1}]) \end{pmatrix} + \eta_{t+1} \quad (\text{A2})$$

$$E_t[X_{t+1}] = \exp\left(E_t[x_{t+1}] + \frac{1}{2}\Sigma\right) \quad (\text{A3})$$

where Σ is the estimated covariance matrix of the shocks η_{t+1} . We include the log of the survey expectations in equation (A2) to ensure that our statistical expectations contain any information known to the forecasters. Given our objective distribution, $x_{t+1} \sim N(E_t[x_{t+1}], \Sigma)$, we assume that forecasters believe $x_{t+1} \sim N(E_t^*[x_{t+1}], \Sigma)$ where $E_t^*[x_{t+1}] \equiv \log(E_t^*[X_{t+1}]) - \frac{1}{2}\Sigma$. We choose RGDP growth and the T-bill rate as the two variables in our subset $\hat{X}_{t+1} \subset X_{t+1}$ and estimate the log SBF that perfectly matches the subjective expectations $E_t^*[\hat{X}_{t+1}]$,

$$\hat{s}_{t+1} \equiv -\frac{1}{2}\hat{\beta}'_t\hat{\Sigma}\hat{\beta}_t + \hat{\beta}'_t\hat{\eta}_{t+1} \quad (\text{A4})$$

where $\hat{\eta}_{t+1}$ is the objective shocks to \hat{x}_{t+1} , $\hat{\Sigma}$ is the covariance matrix of $\hat{\eta}_{t+1}$, and $\hat{\beta}_t = \hat{\Sigma}^{-1}(E_t^*[\hat{x}_{t+1}] - E_t[\hat{x}_{t+1}])$.

Then, we can estimate synthetic expectations based on \hat{s}_{t+1} for the remaining variables as

$$\begin{aligned} \hat{E}_t^*[X_{t+1}] &\equiv E_t[\hat{S}_{t+1}X_{t+1}] \\ &= E_t[X_{t+1}] \exp\left(\text{Cov}(\eta_{t+1}, \hat{\eta}_{t+1})\hat{\beta}_t\right). \end{aligned} \quad (\text{A5})$$

Just as in the main analysis, we can compare our results to more generalized estimates based on the best linear predictor and PCA. Equation (A5) can be rewritten as

$$\log\left(\frac{\hat{E}_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right) = \text{Cov}(\eta_{t+1}, \hat{\eta}_{t+1})\hat{\beta}_t \quad (\text{A6})$$

$$= \text{Cov}(\eta_{t+1}, \hat{\eta}_{t+1})\hat{\Sigma}^{-1}(E_t^*[\hat{x}_{t+1}] - E_t[\hat{x}_{t+1}]). \quad (\text{A7})$$

The best linear predictor is estimated from

Table AV

Condensing the Survey of Professional Forecasters

This table evaluates the ability of the synthetic expectations formed from \hat{s}_{t+1} to explain movements in the 15 Survey of Professional Forecasters variables. Column 1 shows the average R^2 of the regression of $\log(E_t^*[X_{j,t+1}])$ on $\log(\hat{E}_t^*[x_{j,t+1}])$ across the 15 different variables. Column 2 shows the average R^2 of the regression of $\log\left(\frac{E_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right)$ on $\log\left(\frac{\hat{E}_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right)$. For comparison, Column 3 shows the average R^2 of the best linear predictor of $\log\left(\frac{E_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right)$ using the individual biases in *rgdp* and *tbill* coming from equation (A8) and Columns 4 and 5 show the explanatory power of the first two principal components of the 15 biases $\log\left(\frac{E_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right)$.

	$\log(E_t^*[X_{t+1}])$	$\log\left(\frac{E_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right)$			
	$\log\left(\frac{\hat{E}_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right)$	$\log\left(\frac{\hat{E}_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right)$	Best Linear Predictor	PC-1	PC-2
$R^2(\%)$	64.0	48.4	53.8	45.2	65.2

$$\log\left(\frac{\hat{E}_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right) = \alpha + \Gamma(E_t^*[\hat{x}_{t+1}] - E_t[\hat{x}_{t+1}]) + \varepsilon_t. \quad (\text{A8})$$

The PCA results are estimated from

$$\log\left(\frac{\hat{E}_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right) = \alpha + \Gamma\Lambda_t. \quad (\text{A9})$$

Table AV evaluates the fit of our synthetic expectations. The results are quite close to the results in I. The synthetic $\log\left(\frac{\hat{E}_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right)$ accounts for the majority of variation in the survey $\log(E_t^*[X_{t+1}])$. In terms of accounting for the gap between survey and statistical expectations, the synthetic $\log\left(\frac{\hat{E}_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right)$ accounts for roughly half of the variation in $\log\left(\frac{E_t^*[X_{t+1}]}{E_t[X_{t+1}]}\right)$, is close to the upper bound implied by the best linear predictor, out-performs the first principal component, and captures three fourths of the variation explained by the first two principal components. Table AVI shows the correlations for each individual variable.

We can also extend our results to the Blue Chip data. Let Y_{t+1} represent the nine financial variables in the Blue Chip data. We calculate synthetic expectations $\hat{E}_t^*[Y_{t+1}]$ solely using realized data and the Survey of Professional Forecasters survey data. First, we estimate a

Table AVI

Comparing synthetic expectations and the Survey of Professional Forecasters

This table shows the correlation of each of the Survey of Professional Forecasters expectations and biases with their respective synthetic expectations and biases formed from the log SBF \hat{s}_{t+1} . The first column shows the correlation of survey expectations with synthetic expectations for the 15 variables. The second column shows the correlation of the survey biases with the synthetic biases for the same 15 variables. Note that the log SBF \hat{s}_{t+1} is formed only using RGDP growth and the T-Bill rate biases.

	$\text{Corr} \left(\log \left(E_t^* [X_{j,t+1}] \right), \log \left(\hat{E}_t^* [X_{j,t+1}] \right) \right)$	$\text{Corr} \left(\log \left(\frac{E_t^* [X_{t+1}]}{E_t [X_{t+1}]} \right), \log \left(\frac{\hat{E}_t^* [X_{t+1}]}{\hat{E}_t [X_{t+1}]} \right) \right)$
<i>rgdp</i>	1	1
<i>rcon</i>	0.8847	0.8724
<i>cpi</i>	0.6136	0.3761
<i>empl</i>	0.8967	0.6321
<i>indp</i>	0.7088	0.8523
<i>tbill</i>	1	1
<i>aaa</i>	0.9695	0.4783
<i>rnresin</i>	0.7182	0.8118
<i>rresinv</i>	0.7200	0.6682
<i>rgf</i>	0.8426	0.5940
<i>rgsl</i>	0.8168	0.4314
<i>housing</i>	0.5384	0.3899
<i>rcbi</i>	0.4823	0.7466
<i>rexpport</i>	0.9349	0.5071
<i>cprof</i>	0.6204	0.6279

VAR(1) model for $y_{t+1} \equiv \log(Y_{t+1})$,

$$y_{t+1} = A_y + B_y \begin{pmatrix} y_t & x_t & \log(E_t^*[X_{t+1}]) \end{pmatrix} + \eta_{y,t+1} \quad (\text{A10})$$

$$E_t[Y_{t+1}] = \exp\left(E_t[y_{t+1}] + \frac{1}{2}\Sigma_y\right) \quad (\text{A11})$$

where Σ_y is the estimated covariance matrix of $\eta_{y,t+1}$. To ensure our statistical expectations contain as much current information as possible, we include the current value x_t and the survey expectations $\log(E_t^*[X_{t+1}])$ for the Survey of Professional Forecasters variables in equation (A10).

Then, using our log SBF \hat{s}_{t+1} estimated from the Survey of Professional Forecasters RGDP growth and T-bill rate expectations, we calculate our synthetic expectations for Y_{t+1} as

$$\begin{aligned} \hat{E}_t^*[Y_{t+1}] &\equiv E_t\left[\hat{S}_{t+1}Y_{t+1}\right] \\ &= E_t[Y_{t+1}] \exp\left(\text{Cov}(\eta_{y,t+1}, \hat{\eta}_{t+1})\hat{\beta}_t\right). \end{aligned} \quad (\text{A12})$$

Importantly, we do not use any survey expectations of Y_{t+1} in the construction of $\hat{E}_t^*[Y_{t+1}]$. In Table AVII shows how well the synthetic expectations match the survey expectations and shows how well the gap between synthetic expectations and statistical expectations $\log\left(\frac{\hat{E}_t^*[Y_{t+1}]}{E_t[Y_{t+1}]}\right)$ matches the gap between survey expectations and statistical expectations $\log\left(\frac{E_t^*[Y_{t+1}]}{E_t[Y_{t+1}]}\right)$. The results are similar to Table III.

Finally, Table AVIII tests how well our log SBF \hat{s}_{t+1} based on two variables summarizes the combined 24 forecasts from both groups. Let Z_{t+1} be the union of the 15 Survey of Professional Forecasters variables X_{t+1} and the 9 Blue Chip variables Y_{t+1} . Again, the results are quite close to Table IV. We compare our results to the upper bounds implied by the best linear predictor and PCA. Just as in equation (15), the synthetic bias for each of our 24 variables is equal to a linear combination of the bias in log RGDP growth expectations and log T-bill rate expectations, where the coefficients are determined entirely by the objective

Table AVII

Predicting Blue Chip survey expectations

This table evaluates the log SBF \hat{s}_{t+1} formed from only the RDGP growth and T-Bill rate biases in generating synthetic expectations for Blue Chip variables. The first column shows the correlation of the time series of the Blue Chip survey expectations $E_t^*[Y_{j,t+1}]$ with the synthetic expectations $\hat{E}_t^*[Y_{j,t+1}]$ constructed from \hat{s}_{t+1} . The second column shows the correlation of the survey biases with the synthetic biases for the same variables.

	$\text{Corr} \left(\log \left(E_t^* [Y_{j,t+1}] \right), \log \left(\hat{E}_t^* [Y_{j,t+1}] \right) \right)$	$\text{Corr} \left(\log \left(\frac{E_t^* [Y_{t+1}]}{E_t [Y_{t+1}]} \right), \log \left(\frac{\hat{E}_t^* [Y_{t+1}]}{E_t [Y_{t+1}]} \right) \right)$
<i>prime</i>	0.9495	0.8778
<i>fedfunds</i>	0.9514	0.8717
<i>mortgage</i>	0.9326	0.5150
<i>tbill-1yr</i>	0.9387	0.8683
<i>tbill-6m</i>	0.9403	0.8827
<i>tnote-10yr</i>	0.9498	0.5535
<i>tnote-5yr</i>	0.9473	0.7099
<i>tnote-2yr</i>	0.9411	0.8230
<i>libor-3m</i>	0.9182	0.8359

covariance of shocks. We can compare this to the best linear predictor from a regression,

$$\log \left(\frac{E_t^* [Z_{t+1}]}{E_t [Z_{t+1}]} \right) = \alpha_z + \Gamma_z (E_t^* [\hat{x}_{t+1}] - E_t [\hat{x}_{t+1}]) + \varepsilon_{z,t} \quad (\text{A13})$$

and the PCA

$$\log \left(\frac{E_t^* [Z_{t+1}]}{E_t [Z_{t+1}]} \right) = \alpha_z + \Gamma_z \Lambda_{z,t}. \quad (\text{A14})$$

Table AVIII

Condensing Blue Chip and the Survey of Professional Forecasters

This table evaluates the ability of the synthetic expectations $\hat{E}_t^* [Z_{t+1}]$ formed from \hat{s}_{t+1} to explain the 24 Blue Chip + Survey of Professional Forecasters variables $E_t^* [Z_{t+1}]$. Column 1 shows the average R^2 of the regression of the survey expectations $\log(E_t^* [Z_{j,t+1}])$ on $\log(\hat{E}_t^* [Z_{j,t+1}])$ across the 24 different variables. Column 2 shows the average R^2 of the regression of the survey biases $\log\left(\frac{E_t^* [Z_{j,t+1}]}{E_t [Z_{j,t+1}]}\right)$ on the synthetic biases $\log\left(\frac{\hat{E}_t^* [Z_{j,t+1}]}{E_t [Z_{j,t+1}]}\right)$. For comparison, Column 3 shows the average R^2 of the best linear predictor of $\log\left(\frac{\hat{E}_t^* [Z_{j,t+1}]}{E_t [Z_{j,t+1}]}\right)$ using the individual *rgdp* and *tbill* biases coming from equation (A13) and Columns 4 and 5 show the explanatory power of the first two principal components of the 24 biases $\log\left(\frac{\hat{E}_t^* [Z_{j,t+1}]}{E_t [Z_{j,t+1}]}\right)$.

	$\log(E_t^* [Z_{t+1}])$	$\log\left(\frac{E_t^* [Z_{t+1}]}{E_t [Z_{t+1}]}\right)$			
	$\log(\hat{E}_t^* [Z_{t+1}])$	$\log\left(\frac{\hat{E}_t^* [Z_{t+1}]}{E_t [Z_{t+1}]}\right)$	Best Linear Predictor	PC-1	PC-2
$R^2(\%)$	72.1	52.7	57.5	48.8	68.3