The Subjective Belief Factor

TINGYUE CUI, RICARDO DELAO, and SEAN MYERS*

June 2026

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), i.e., a distortion that characterizes many subjective expectations even for non-financial variables. 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. This SBF also accurately replicates differences across the 24 variables in the serial correlation of forecast errors and under/overreaction. Further, the SBF can be used to succinctly incorporate subjective beliefs data into asset pricing models, even those involving many expectations. Applying our measured SBF to a model of cross-sectional stock returns, we estimate that distorted beliefs account for the majority of excess returns for the Fama-French factors and explain about two thirds of the variation in returns across 176 anomalies, while the remaining third is attributed to preferences/risk. Our results support models like robust control and certain versions of diagnostic expectations in which agents' beliefs across different variables are characterized by a single probability distortion.

*Tingyue Cui and Sean Myers are with the Wharton School at the University of Pennsylvania. Ricardo Delao is with the Marshall School of Business, University of Southern California. We thank Aizhan Anarkulova, Adem Atmaz, Nick Barberis, Jaroslav Borovicka, Stefano Cassella, William Cassidy, Mikhail Chernov, Stefano Giglio, Andrei Goncalves, Theis Jensen, Petri Jylha, Mete Kilic, Ralph Koijen, Leonid Kogan, Arthur Korteweg, Johnathan Loudis, Lars Lochstoer, Pooya Molavi, Stefan Nagel, Nikolai Roussanov, Mirela Sandulescu, Tom Sargent, Alireza Tahbaz-Salehi, Sanjay Singh, Dejanir Silva, Andrea Vedolin, as well as seminar participants at Chicago Booth, Baruch College, University of Lausanne, HEC Montreal, University of Arizona, University of Toronto, University of Maryland, Bank of Canada, HKUST, HKU, CityUHK, and numerous conference participants for valuable comments. 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

$$P_t \;\; = \;\; E_t\left[M_{t+1}X_{t+1}\right] \tag{1}$$

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,

$$E_t^*\left[X_{t+1}\right] \;\; = \;\; E_t\left[S_{t+1}X_{t+1}\right]. \tag{2}$$

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}.[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 who overstates the probability of low real GDP growth states will also overstate the probability of high unemployment states, as these states overlap. In this sense, the SBF provides a way to discipline beliefs about unemployment using information about beliefs over real GDP, rather than introducing a separate distortion for each variable.

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 whether asset price anomalies reflect biased beliefs or preferences/risk. Direct data on sub-

[1]For example, even if a researcher has no interest in asset prices and only cares about expectations of output, inflation, and other macroeconomic variables, he/she can still use asset pricing techniques designed to study M_{t+1} from prices in order to study S_{t+1} from subjective expectations data.

jective expectations can help to resolve this. Instead of individually modeling biased beliefs about dozens of variables, our approach summarizes subjective expectations for many variables with a single SBF, S_{t+1}. This allows us to 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.

In this paper, we (i) introduce a framework that jointly models multiple subjective expectations through a single subjective belief factor and (ii) demonstrate its asset pricing applications. 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. Using this SBF, we construct synthetic expectations $E_t\left[S_{t+1}Y_{t+1}\right]$ for an additional nine macroeconomic variables Y_{t+1} and then verify the accuracy of these synthetic expectations using additional survey data from Blue Chip. Along with matching the time series for the 24 subjective expectations, this SBF also closely replicates differences across the variables in two common tests for deviations from Full Information Rational Expectations (FIRE): the predictability of forecast errors using lagged forecast errors (i.e., serially correlated forecast errors) and the predictability of forecast errors using recent changes in expectations (i.e., Coibion-Gorodnichenko regressions). Finally, we show that this SBF accounts for the majority of excess returns for the Fama-French cross-sectional anomalies and accounts for roughly two-thirds of differences in excess returns across 176 anomalies sorted into 22 categories from Chen and Zimmermann (2022).

We view our results as a proof of concept for this framework. Our empirical analysis focuses on macroeconomic expectations of professional forecasters, so the estimated SBF is not intended to capture all belief distortions in all settings. Rather, the exercises show that a low-dimensional belief distortion can summarize a substantial portion of documented biases in macroeconomic expectations and can be carried into asset-pricing applications in a disciplined way.

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 that subjective expectations are "coherent," meaning that they satisfy basic rules regarding addition and multiplication. 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.

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. We study annual forecasts for 15 variables, the maximum number 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 $E_t\left[S_{t+1}X_{t+1}\right]$ for all 15 variables match the actual subjective expectations with an average R^2 of 65.8% and match the differences between subjective and statistical expectations with an average R^2 of 47.4%. This result means that, given an SBF that matches subjective expectations of real GDP growth and the T-bill rate, we can explain subjective expectations for the remaining 13 variables based on their objective covariances 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 Y_{t+1}, such as the mortgage rate and the 10-year Treasury rate, and calculate synthetic expectations $E_t\left[S_{t+1}Y_{t+1}\right]$ based on our estimated SBF. 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 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 similar average R^2 values to the Survey of Professional Forecasters results.

In addition to studying the ability of the estimated SBF to replicate the time series of subjective expectations, we also evaluate its ability to replicate key moments related to deviations from FIRE. Specifically, we study the predictability of forecast errors using lagged forecast errors and the predictability of forecast errors using recent changes in expectations. These moments play a key role in parameterizing models of learning and behavioral biases. Estimating these moments using the Survey of Professional Forecasters and Blue Chip survey data, we find that the serial correlation in forecast errors differs substantially across the 24 variables, ranging from -0.18 to 0.51. Similarly, the coefficients from Coibion-Gorodnichenko regressions also range widely from -0.29 (overreaction) to 1.26 (underreaction).[2] We then repeat these tests of forecast error predictability but replace the predictors with the SBF-implied lagged forecast errors and the change in the SBF-implied expectations. Overall, we find that the SBF-implied predictors closely replicate these moments across the 24 variables and additionally replicate the cross-variable predictability documented in Wu (2023).

What do these findings mean for models of expectation formation? Our results emphasize the benefits of modeling subjective expectations for different variables jointly rather than individually. While all 24 subjective expectations differ from their corresponding statistical expectations and even display heterogeneous behavioral patterns such as under/overreaction, they can still be summarized by a low-factor SBF. Importantly, our results are non-causal. It could be that agents form each expectation individually but were also impacted by other biases such as misperceptions of cross-variable relationships, maturities, and persistences which caused the SBF to perform well. Our goal is to highlight that a framework in which

[2]Since we are studying consensus forecasts, the Coibion-Gorodnichenko regressions generally give positive coefficients, indicating underreaction. This is consistent with Bordalo et al. (2020) who show that consensus forecasts tend to exhibit underreaction whereas individual forecasts exhibit overreaction.

expectations are formed jointly using a low-factor SBF provides a parsimonious alternative that successfully matches the time series of survey expectations and reconciles a number of untargeted moments of forecast error predictability.

These results raise the prospect that models of individual variable expectations such as inflation (e.g., Malmendier and Nagel, 2016), consumption (e.g., Collin-Dufresne, Johannes, and Lochstoer, 2016), risk-free rates (e.g., Haubrich, Pennacchi, and Ritchken, 2012), etc. can be represented as different manifestations of one underlying belief distortion. In particular, our findings support models such as robust control (e.g., Hansen and Sargent, 2001; Bhandari, Borovička, and Ho, 2025; Maenhout, Vedolin, and Xing, 2025) and certain forms of diagnostic expectations (e.g., Bordalo et al., 2019; Maxted, 2024; Li, Van Nieuwerburgh, and Renxuan, 2025) in which the agent's beliefs are characterized by a probability distortion that impacts her expectations of multiple variables. As further support for these models of jointly biased expectations, we show that the factor structure of the 24 subjective expectations is often too strong to be explained by models in which the bias for each variable only depends on its own shocks.

In our last set of results, we emphasize the benefits of our approach for incorporating subjective expectations into economic models. Economic models typically depend on expectations for multiple variables, and two common ways to handle this are to (i) specify a complete model of joint belief formation or (ii) collect survey data for every variable of interest. This paper offers a third, middle-ground option: estimate an SBF S_{t+1} from survey expectations of a few key variables and, under a mild coherence assumption, infer expectations for all other variables, $E_t\left[S_{t+1}X_{t+1}\right]$.[3] This provides a minimum-assumptions approach to quantify departures from FIRE and circumvents the need to obtain survey expectations for all variables.

To demonstrate this, we focus on a model of cross-sectional asset pricing with a representative investor who has preferences $\tilde{M}_{t+1}$ for payoffs in different states. In this model, the

[3]Section I.C discusses the role of projections and spanning for the accuracy of these inferred expectations.

key equilibrium condition is that $E_t\left[S_{t+1}\tilde{M}_{t+1}R_{j,t+1}\right]$ equals one for all assets j. The goal of the model is to explain observed differences in average returns across assets. In particular, if all variables are log-normal, then the historical average excess return $\bar{R}^e_{j,t+1}$ for an asset can be decomposed into

$$\log\left(\bar{R}^e_{j,t+1}\right) = Cov\left(-s_{t+1}, r^e_{j,t+1}\right) + Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right) \tag{3}$$

where lowercase denotes log values. 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.

This decomposition is related to work linking excess returns to subjective return expectations[4] but removes the need for survey expectations data on each individual asset, as well as the need for assumptions on $\tilde{M}_{t+1}$, such as $\tilde{M}_{t+1}$ being constant (Bordalo et al., 2025). By linking many excess returns to a single belief-related variable s_{t+1}, this decomposition also relates to work studying measures of "sentiment" (Baker and Wurgler, 2006; Stambaugh, Yu, and Yuan, 2012; Huang et al., 2015). Because this decomposition is quantitative and relies on covariances rather than correlations, it addresses the common issue of risk and sentiment being correlated.[5]

Using our SBF estimated from 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 from sorting firms by earnings surprises to capture short-term behavioral biases (e.g., weekly-level biases) and a separate factor using share

[4] For examples, see Greenwood and Shleifer (2014), Engelberg, McLean, and Pontiff (2020), Giglio, Maggiori, Stroebel, and Utkus (2021), Dahlquist and Ibert (2024), Jensen (2024), Couts, Gonçalves, and Loudis (2025), Décaire and Graham (2026), Delao, Han, and Myers (2026), and Kremens, Martin, and Varela (2025).

[5] If 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 the decomposition presented here, 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 $\frac{Cov\left(-s_{t+1}, r^e_{j,t+1}\right)}{\log\left(\bar{R}^e_{j,t+1}\right)} \in (0,1)$ will tell us what portion is attributable to s_{t+1} and what portion is attributable to $\tilde{m}_{t+1}$.

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 (96%) of the excess return for their long-horizon factor. Finally, we study a large set of 176 anomalies sorted into 22 categories from Chen and Zimmermann (2022). In a two-part decomposition, we find that the belief-based component $Cov\left(-s_{t+1}, r^e_{j,t+1}\right)$ accounts for 64.9% of the differences in excess returns across anomalies, while the remaining 35.1% is attributed to the preference-based component $Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right)$. In a more detailed three-part decomposition, we find that the belief-based component would explain 232% of the differences in excess returns if the preference-based component was fixed across anomalies, but that negative correlation across anomalies between the belief-based component and preference-based component plays an important role in dampening the role of the belief-based component to 64.9%. 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 60-40 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 of macroeconomic variables. Given the size of this literature, we refer readers to the handbook by Bachmann, Topa, and van der Klaauw (2023) for an overview and list of references. To a large extent, these papers focus on individual variable expectations, such as inflation, output, interest rates, exports, housing, unemployment, etc., rather than joint expectations.[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 under/overreaction for each individual variable). We build on this work by presenting an approach that allows us to study subjective expectations for many variables

[6]For some recent examples, see Cieslak (2018), Kuchler and Zafar (2019), Mueller, Spinnewijn, and Topa (2021), Link et al. (2023), and Farmer, Nakamura, and Steinsson (2024). Notable exceptions are Bundick (2015), Kim and Pruitt (2017), Jia, Shen, and Zheng (2023), and Bauer, Pflueger, and Sunderam (2024a,b) who study joint expectations of inflation, yields, and output through a perceived Taylor rule. Our approach takes this even further by showing that there exists an SBF that allows all expectations to be studied jointly.

jointly and to condense these subjective expectations down to a single SBF based on only a few variables. Our approach is heavily indebted to the work of Hansen and Richard (1987) and Hansen and Jagannathan (1991) on the stochastic discount factor. Our emphasis on jointly analyzing multiple macroeconomic variables is related to recent work studying statistical and machine learning forecasts for large sets of variables (e.g., Bianchi, Ludvigson, and Ma, 2026).

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, i.e., S_{t+1} and $\tilde{M}_{t+1}$. The idea of framing subjective expectations as a distortion has been discussed previously. For example, Chernov and Mueller (2012) and Piazzesi, Salomao, and Schneider (2015) use expected and realized inflation and bond yields to study objective, subjective, and risk-neutral probabilities. We demonstrate that this idea of a distortion between the objective and subjective probabilities (i.e., the SBF) can be used not only to describe existing survey expectations data but also to condense the distortion to a few key factors and to form synthetic expectations for other variables. We also emphasize that this approach can be applied outside the context of asset pricing. Even for a researcher purely interested in understanding subjective expectations, 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), Borovička, Hansen, and Scheinkman (2016), and Adam and Nagel (2023) emphasize the difficulty of distinguishing S_{t+1} and $\tilde{M}_{t+1}$ solely from asset prices. We provide a general method to estimate S_{t+1} from survey expectations and demonstrate how this can be used to decompose a wide range of observed excess returns into a belief-based component and a preference/risk-based component. Importantly, this can be done even if we do not have survey expectations data for each specific asset. This is closely related to recent work by Chen, Hansen, and Hansen (2020, 2024) who use 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 demonstrates the conditions under which subjective expectations for many variables can be described by a single Subjective Belief Factor and establishes its key properties. Section II discusses the data on subjective expectations taken from the Survey of Professional Forecasters and Blue Chip. Section III demonstrates that an SBF estimated from two key variables can summarize the full set of expectations for financial and non-financial macroeconomic variables. Section IV shows that this SBF parsimoniously replicates a number of documented moments related to forecast error predictability. Section V uses the estimated SBF in a multi-asset setting to explain excess returns for a wide range of anomalies. Section VI 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.[7] These condi-

[7]For clarity, we impose this condition for the constant 1 as well as the constant random variable 1, which is a variable that takes value 1 with unit probability under the objective probability distribution. This prevents violations of absolute continuity. From an econometrician's perspective, a constant random variable is observationally equivalent to a constant.

tions 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 B. To an econometrician who 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, interest rate expectations, unemployment expectations, etc., we can study the single SBF which jointly explains these expectations.

It is useful to note that many models of subjective expectations assume agents can add and multiply (e.g., sticky expectations, extrapolation, adaptive expectations, learning, noisy signals, etc.). Breaking these assumptions within a model is generally difficult as it makes the model predictions sensitive to small arbitrary changes, such as measuring outcomes in dollars versus cents. Thus, we argue that this assumption of coherence is a useful starting point to study expectations of many variables simultaneously. The fact that the estimated SBF is empirically successful in summarizing survey expectations in Section III also provides support for this approach. Appendix A.4 discusses how coherence can be violated if these mechanisms (e.g., sticky expectations, extrapolation) are applied separately to variables whose shocks are linearly dependent, and how this relates to robust control and different applications of diagnostic expectations.

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

A. Aggregation

Given a set of individuals, $i = 1, 2, ..., n$, 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 $E_t^* [\cdot] \equiv \frac{1}{n} \sum_i E_{i,t}^* [\cdot]$.

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

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, we will focus on a single aggregated agent for the rest of the paper.

B. Log-normal representation

Proposition 1 tells us that an SBF exists. A natural next question is how we can find a variable S_{t+1} that matches a given set of subjective expectations. Given a multivariate X_{t+1} and subjective expectations $E_t^*[X_{t+1}]$, we know that

$$E_t[S_{t+1}X_{t+1}] = E_t[X_{t+1}] + Cov_t(S_{t+1}, X_{t+1}) \tag{4}$$

given the definition of covariance and the fact that $E_t[S_{t+1}] = 1$. Thus, finding an S_{t+1} such that $E_t^*[X_{t+1}] = E_t[S_{t+1}X_{t+1}]$ for all elements of X_{t+1} simply requires finding an S_{t+1} that has the correct objective covariance with X_{t+1}.

One solution is to represent the SBF as a linear projection onto the set of objective shocks.

Lemma 3. *Given a multivariate X_{t+1} and subjective expectations $E_t^*[X_{t+1}]$, we have $E_t^*[X_{t+1}] = E_t[S_{t+1}X_{t+1}]$ for*

$$\begin{aligned} S_{t+1} &= 1 + \beta_t' \varepsilon_{t+1} \\ \beta_t &= \Sigma_t^{-1}\left(E_t^*[X_{t+1}] - E_t[X_{t+1}]\right) \\ \varepsilon_{t+1} &= X_{t+1} - E_t[X_{t+1}] \end{aligned}$$

where Σ_t is the objective covariance matrix of the objective shocks ε_{t+1}.

For an element $\beta_{j,t}$ of the vector β_t, a positive $\beta_{j,t}$ means that it is as if the agent is exaggerating the probability of positive shocks $\varepsilon_{j,t+1}$ and understating the probability of negative shocks. The converse is true for $\beta_{j,t} < 0$. Note that for a finite X_{t+1}, the SBF is not unique, as one could always consider a variable W_{t+1} with shocks ε_{t+1}^w orthogonal to ε_{t+1} and an alternative SBF $\tilde{S}_{t+1} = 1 + \beta_t' \varepsilon_{t+1} + \beta_t^w \varepsilon_{t+1}^w$. This alternative SBF would also match the subjective expectations for X_{t+1}. In this case, the SBF would be based off of observable subjective expectations $E_t^*[X_{t+1}]$ and the researcher's assumption that the agent has distorted beliefs about the unrelated W_{t+1}.

The benefit of Lemma 3 is that it requires no assumptions about the distribution of

X_{t+1}. A potential limitation of Lemma 3 is that the projected SBF may be negative for large magnitude shocks ε_{t+1}. Fortunately, equation (4) shows that if we have information about the *objective* distribution of X_{t+1}, then we can estimate an SBF that is always nonnegative. In particular, if variables are objectively normally distributed, then we can specify an SBF that is both tractable and always nonnegative. Let $s_{t+1} \equiv \log(S_{t+1})$.

Proposition 2. *Given a multivariate X_{t+1} and subjective expectations $E_t^*[X_{t+1}]$, if the objective conditional distribution is $X_{t+1} \sim N(E_t[X_{t+1}], \Sigma)$ then $E_t^*[X_{t+1}] = E_t[S_{t+1}X_{t+1}]$ for*

$$\begin{aligned} s_{t+1} &= -\frac{1}{2}\beta_t'\Sigma\beta_t + \beta_t'\varepsilon_{t+1} \\ \beta_t &= \Sigma^{-1}\left(E_t^*[X_{t+1}] - E_t[X_{t+1}]\right) \\ \varepsilon_{t+1} &= X_{t+1} - E_t[X_{t+1}]. \end{aligned}$$

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., $\beta_t'\varepsilon_{t+1}$. The $-\frac{1}{2}\beta_t'\Sigma\beta_t$ term in s_{t+1} is simply a Jensen's term to ensure $E_t[S_{t+1}] = 1$. The result is very similar to Lemma 3 but ensures that $S_{t+1} > 0$.

Importantly, Proposition 2 does not require any assumptions about the agent's subjective beliefs about the distribution (e.g., assuming the agent believes X_{t+1} is normally distributed). As shown in equation (4), we only need to know the *objective* covariance between S_{t+1} and X_{t+1} in order to ensure that $E_t[S_{t+1}X_{t+1}]$ matches the subjective $E_t^*[X_{t+1}]$.

The same logic holds if X_{t+1} is not normal but is a function of normal variables. Suppose $X_{t+1} = f_t(\varepsilon_{t+1})$ where $f_t(\cdot)$ is a potentially time-varying function and ε_{t+1} is objectively multivariate standard normal.[8] For example, X_{t+1} could be a CES aggregator, $X_{t+1} = \left(a_t^\rho + (B_t\varepsilon_{t+1})^\rho\right)^{1/\rho}$, or an indicator variable, $X_{t+1} = \mathbb{1}\{a_t + B_t\varepsilon_{t+1} > 0\}$.

Proposition 3. *Given a multivariate $X_{t+1} = f_t(\varepsilon_{t+1})$ and subjective expectations $E_t^*[X_{t+1}]$,*

[8] Assuming that ε_{t+1} is standard normal rather than simply normal is WLOG. If $X_{t+1} = f_t(\eta_{t+1})$ where $\eta_{t+1} \sim N(\mu_t, \Sigma_t)$, then this can always be expressed as $X_{t+1} = \tilde{f}_t(\varepsilon_{t+1})$ where $\tilde{f}_t(\varepsilon_{t+1}) \equiv f_t\left(\mu_t + \Sigma_t^{0.5}\varepsilon_{t+1}\right)$ and ε_{t+1} is standard multivariate normal.

if the objective conditional distribution is $\varepsilon_{t+1} \sim N(0, I)$ *then* $E_t^*[X_{t+1}] = E_t[S_{t+1}X_{t+1}]$ *for*

$$\begin{aligned} s_{t+1} &= -\frac{1}{2}\beta_t'\beta_t + \beta_t'\varepsilon_{t+1} \\ \beta_t &= h_t^{-1}\left(E_t^*[X_{t+1}]\right) \\ h_t(\beta) &\equiv E_t\left[f_t(\beta + \varepsilon_{t+1})\right]. \end{aligned}$$

In words, if X_{t+1} is a function of normal shocks, then there is a log-normal SBF that matches the subjective expectations, $E_t^*[X_{t+1}] = E_t[S_{t+1}X_{t+1}]$, and the shock to this log-normal SBF is a linear combination of the objective shocks, $\beta_t'\varepsilon_{t+1}$. The details of the function $f_t(\cdot)$ only affect the loadings β_t. Specifically, the loadings depend on the inverse of the $f_t(\cdot)$ function. For example, if X_{t+1} is log-normal, $X_{t+1} = \exp\left(\mu_t + \Sigma^{0.5}\varepsilon_{t+1}\right)$, then $\beta_t = \Sigma^{-1}\left[\log\left(E_t^*[X_{t+1}]\right) - \log\left(E_t[X_{t+1}]\right)\right]$.

Note that Propositions 2 and 3 allow for a wide range of possible subjective expectations $E_t^*[X_{t+1}]$. Given an objective process for X_{t+1} that is conditionally normal or a function of normal shocks, different subjective expectations $E_t^*[X_{t+1}]$ simply appear as different loadings β_t in the log-normal SBF. The subsection below provides an example. In Appendix A.2, we discuss an extension of Propositions 2-3 that can be used when we want to not only match the subjective expectations of the mean $E_t^*[X_{t+1}]$ but also want to match additional data on expected variances and covariances.

B.1. The SBF in simple expectation formation models

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

$$g_{t+1} = \rho g_t + \varepsilon_{t+1} \tag{5}$$

where $\rho > 0$ and $\varepsilon_{t+1} \sim N(0, \sigma^2)$. For simplicity, we set the mean of the process equal to 0. From Proposition 2, subjective expectations of g_{t+1} can be represented by the log SBF $s_{t+1} = -\frac{1}{2}\beta_t^2\sigma^2 + \beta_t\varepsilon_{t+1}$. 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 $-\frac{1}{2}\beta_t^2\sigma^2$ 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 ε_{t+1}. Different models of expectation formation will imply different β_t. Below, we show the β_t for a model of extrapolation. Appendix A.3 shows the loading β_t for diagnostic expectations, Bayesian learning about the mean, and robust control.

Suppose the agent is extrapolative, meaning that the agent believes the persistence is ρ^* rather than ρ, and expects next period growth to be $\rho^* g_t$. The loading is then

$$\beta_t^{EX} = \frac{\rho^* - \rho}{\sigma^2} g_t. \tag{6}$$

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 ε_{t+1} and understating the probability of negative ε_{t+1}. When current growth is low, it is as if the agent is exaggerating the probability of negative ε_{t+1} and understating the probability of positive ε_{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). Proposition 2 allows the econometrician to represent this belief using the SBF S_{t+1}.

C. *Synthetic expectations and spanning*

Since having expectations for all variables of interest X_{t+1} is typically infeasible, it is useful to know how well an SBF estimated from the expectations of a subset of variables or a completely different set of variables $\hat{X}_{t+1} \neq X_{t+1}$ approximates the full set of expectations $E_t^*[X_{t+1}]$.

Let S_{t+1} be the SBF from Proposition 1 and Lemma 1 that matches subjective expectations $E_t^*[X_{t+1}] = E_t[S_{t+1}X_{t+1}]$ for all variables X_{t+1}. Given a set of variables $\hat{X}_{t+1}$ and subjective expectations $E_t^*\left[\hat{X}_{t+1}\right]$, let $\hat{S}_{t+1}$ be such that $E_t^*\left[\hat{X}_{t+1}\right] = E_t\left[\hat{S}_{t+1}\hat{X}_{t+1}\right]$. For clarity, we will refer to $\hat{S}_{t+1}$ as the "estimated SBF."

Given an estimated SBF $\hat{S}_{t+1}$, one can construct synthetic expectations $E_t\left[\hat{S}_{t+1}X_{t+1}\right]$ for other variables X_{t+1}. Proposition 4 tells us that the accuracy of these synthetic expectations is directly tied to how well X_{t+1} and/or S_{t+1} is objectively spanned by the $\hat{X}_{t+1}$ used to estimate $\hat{S}_{t+1}$.

Proposition 4. *The difference between subjective expectations $E_t^*[X_{t+1}]$ and synthetic expectations $E_t\left[\hat{S}_{t+1}X_{t+1}\right]$ only depends on u_{t+1} and $u_{t+1}^{S-\hat{S}}$,*

$$E_t^*[X_{t+1}] = E_t\left[\hat{S}_{t+1}X_{t+1}\right] + Cov_t\left(u_{t+1}, u_{t+1}^{S-\hat{S}}\right), \tag{7}$$

where u_{t+1} is the unspanned residual from projecting X_{t+1} onto $\hat{X}_{t+1}$, i.e., $X_{t+1} = a_t + b_t'\hat{X}_{t+1} + u_{t+1}$, and $u_{t+1}^{S-\hat{S}}$ is similarly the unspanned residual from projecting $S_{t+1} - \hat{S}_{t+1}$ onto $\hat{X}_{t+1}$.

In words, the synthetic expectations using the estimated SBF $\hat{S}_{t+1}$ contain all biases in expected X_{t+1} that are spanned by $\hat{X}_{t+1}$. These synthetic expectations will only differ from the subjective expectations if (i) there is unspanned variation in the variable, u_{t+1}, (ii) there is unspanned variation in the true SBF, $u_{t+1}^{S-\hat{S}}$, and (iii) the unspanned u_{t+1} and $u_{t+1}^{S-\hat{S}}$ are correlated. For example, this proposition tells us that a low-factor $\hat{S}_{t+1}$ can provide accurate synthetic expectations even for X_{t+1} that are not well-spanned by $\hat{X}_{t+1}$, so long as the true SBF is well-spanned by $\hat{X}_{t+1}$.

If the variables are objectively normal, we have an even simpler expression.

Lemma 4. *If X_{t+1} and $\hat{X}_{t+1}$ are conditionally normal and $\hat{S}_{t+1}$ is defined by Proposition 2 to match $E_t^*\left[\hat{X}_{t+1}\right] = E_t\left[\hat{S}_{t+1}\hat{X}_{t+1}\right]$, then equation (7) reduces to:*

$$E_t^*[X_{t+1}] = E_t\left[\hat{S}_{t+1}X_{t+1}\right] + Cov_t\left(u_{t+1}, u_{t+1}^{S}\right) \tag{8}$$

where u_{t+1} is the unspanned residual from projecting X_{t+1} onto $\hat{X}_{t+1}$ and u_{t+1}^{S} is the unspanned residual from projecting S_{t+1} onto $\hat{X}_{t+1}$.

Appendix A.1 shows that equations (7)-(8) can be used to estimate a range for synthetic expectations, rather than a point estimate $E_t\left[\hat{S}_{t+1}X_{t+1}\right]$, where the range is wider for variables that have more unspanned variation u_{t+1}.

D. Application to economic models

Here, we sketch a general outline on how the Subjective Belief Factor can be embedded in standard economic analyses. Section V then provides a detailed example. In many environments, equilibrium conditions can be expressed as

$$A_t = E_t^*\left[X_{t+1}\right],$$

where A_t is some observable outcome of interest and $E_t^*\left[\cdot\right]$ is the model agent's expectation. For example, in models of investment with capital adjustment costs, we would have $A_t = I_t/K_t$ as the investment-to-capital ratio and $X_{t+1} = MPK_{t+1}$ as the marginal product of capital.

The conventional approach for researchers studying these types of settings is to impose FIRE, estimate the objective process for X_{t+1}, and then compare the objective expectations $E_t[X_{t+1}]$ to the empirically observed A_t. A limitation of this approach is that if the estimated $E_t[X_{t+1}]$ differs noticeably from the empirical A_t, then the researcher does not know if this is due to agents in the data violating the equilibrium condition $A_t = E_t^*\left[X_{t+1}\right]$ or if this is due to agents having expectations that differ from the objective $E_t[X_{t+1}]$. One way to address this issue is to obtain subjective expectations $E_t^*[X_{t+1}]$ and compare them to the empirical A_t. However, in practice, this is often infeasible. In the example above, it would require survey forecasts of MPK_{t+1}, an object that is not covered in most surveys of expectations. Conveniently, the SBF provides a straightforward way to utilize survey forecasts for related variables (e.g., output growth, employment) and integrate them into an economic analysis.

Let $\hat{S}_{t+1}$ denote an estimated SBF constructed from different variables $\hat{X}_{t+1}$ economically related to X_{t+1}, but not necessarily including X_{t+1} itself. One can then compare the synthetic expectation $E_t[\hat{S}_{t+1}X_{t+1}]$ to the observed A_t to assess whether incorporating distorted beliefs helps to explain the empirically observed A_t relative to the FIRE benchmark $E_t\left[X_{t+1}\right]$. Thus, even when direct survey measures of $E_t^*[X_{t+1}]$ are unavailable, the SBF allows one to construct synthetic expectations that capture all biases in expected X_{t+1} that are spanned by $\hat{X}_{t+1}$.

In Section V, we present an application of this outline to a common model in finance, a representative agent investing in multiple assets. While the model equilibrium conditions involve expectations for hundreds of assets, we show that a single SBF based on subjective expectations of just two macroeconomic variables (real output and interest rates) explains a substantial amount of the cross-section of asset prices.

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. We calculate the consensus forecast as the average across the individual-level forecasts. Given that we use lagged forecast errors in part of our analysis, our main sample starts in 1982Q4.

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.

The secondary source of data is the Blue Chip survey. The Blue Chip sample starts in 1988Q1. We consider the complete list of variables that (i) are not included in the Survey of Professional Forecasters over our sample and (ii) have survey data since 1988Q1, 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. To have lagged forecast errors, our main Blue Chip sample starts in 1989Q2.

The realized outcomes for all interest rate variables are obtained from the Federal Reserve Bank of St. Louis and Federal Reserve Board, as well as the Intercontinental Exchange for LIBOR. The realized outcomes for all other variables are obtained from the real-time data files maintained by the Federal Reserve Bank of Philadelphia. Refer to Appendix C for full details on each of the surveys and the data construction.

III. The SBF and Macroeconomic Expectations

In Sections III and IV, our goal is to demonstrate how utilizing tools from asset pricing opens new doors for understanding subjective expectations. First, given a dataset of survey expectations for multiple variables, we can condense these expectations down to a single SBF $\hat{S}_{t+1}$ based on a few key variables. Specifically, we show that survey expectations for 15 macroeconomic variables from the Survey of Professional Forecasters can largely be explained by an SBF related to RGDP growth and the T-bill rate. Second, we can use the estimated SBF to predict subjective expectations for variables outside of this dataset (i.e., variables for which we do not have survey data). Using the estimated SBF from the Survey of Professional Forecasters data, we predict subjective expectations for 9 financial variables and then confirm the accuracy of these predictions by comparing them to Blue Chip financial forecasts. Third, using R^2 bounds, we show that the success of this two-factor SBF in explaining biases in the 24 subjective expectations is too high to be explained by models in which the bias for each variable only depends on its own shocks. Fourth, we show that this SBF replicates a range of untargeted moments of forecast error predictability. To better understand these results, Section III.C provides more details on the performance of the SBF and Section III.D provides direct tests of coherence.

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*, $corr = 0.92$), some of them have no correlation, (like *rcbi* and *rresinv*, $corr = -0.02$), and some are negatively correlated (like *housing* and *rgsl*, $corr = -0.58$). The average pairwise correlation is 0.16.

We calculate statistical expectations using an autoregressive model for X_{t+1}. Specifically, we estimate

$$X_{t+1} \quad = \quad a + B \begin{pmatrix} X_t & E_t^*[X_{t+1}] \end{pmatrix} + \varepsilon_{t+1} \tag{9}$$

where ε_{t+1} is a multivariate Gaussian shock with covariance matrix Σ.[9] The statistical

[9] We find very similar results if we assume X_{t+1} is conditionally log-normal rather than conditionally normal. Figure A10 in Appendix J shows that our estimate of the SBF is almost identical.

expectations are then

$$E_t\left[X_{t+1}\right] = a + B\left(\begin{array}{cc} X_t & E_t^*\left[X_{t+1}\right] \end{array}\right). \tag{10}$$

We include the survey expectations in equation (9) to ensure that our statistical expectations incorporate any information known to the forecasters. This ensures that any discrepancy between $E_t^*\left[X_{t+1}\right]$ and $E_t\left[X_{t+1}\right]$ is due to the statistical expectations being a better predictor of X_{t+1}, and not from any informational advantage of the forecasters.[10]

From Proposition 2, we know that there exists an SBF S_{t+1} that perfectly replicates the survey expectations, $E_t^*\left[X_{t+1}\right] = E_t\left[S_{t+1}X_{t+1}\right]$. Our goal in this section is to test how well we can replicate the survey expectations using an estimated SBF $\hat{S}_{t+1}$ based on a small number of variables. Specifically, given any subset of variables $\hat{X}_{t+1} \subset X_{t+1}$, we know from Proposition 2 that we can estimate a log SBF that perfectly matches the expectations for $\hat{X}_{t+1}$,

$$\hat{s}_{t+1} = -\frac{1}{2}\hat{\beta}_t'\hat{\Sigma}\hat{\beta}_t + \hat{\beta}_t'\hat{\varepsilon}_{t+1} \tag{11}$$

where $\hat{\varepsilon}_{t+1}$ is the vector of objective shocks to $\hat{X}_{t+1}$, $\hat{\Sigma}$ is the covariance matrix of $\hat{\varepsilon}_{t+1}$, and $\hat{\beta}_t = \hat{\Sigma}^{-1}\left(E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]\right)$. Then, we can estimate synthetic expectations based on $\hat{s}_{t+1}$ for the remaining variables as[11]

$$\begin{aligned}
\hat{E}_t^*\left[X_{t+1}\right] &\equiv E_t\left[\hat{S}_{t+1}X_{t+1}\right] \\
&= E_t\left[X_{t+1}\right] + Cov\left(\varepsilon_{t+1}, \hat{\varepsilon}_{t+1}\right)'\hat{\beta}_t \\
&= E_t\left[X_{t+1}\right] + Cov\left(\varepsilon_{t+1}, \hat{\varepsilon}_{t+1}\right)'\hat{\Sigma}^{-1}\left(E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]\right).
\end{aligned} \tag{12}$$

It is useful to highlight two points for understanding these synthetic expectations $\hat{E}_t^*\left[X_{t+1}\right]$. First, there is a straightforward intuition for equation (12). As discussed in Section I.C, we can always represent X_{t+1} as a projection onto $\hat{X}_{t+1}$. The term $Cov\left(\varepsilon_{t+1}, \hat{\varepsilon}_{t+1}\right)'\hat{\Sigma}^{-1}$ in equation (12) is simply the slope coefficients for this projection. So, the synthetic expecta-

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

[11] The 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.

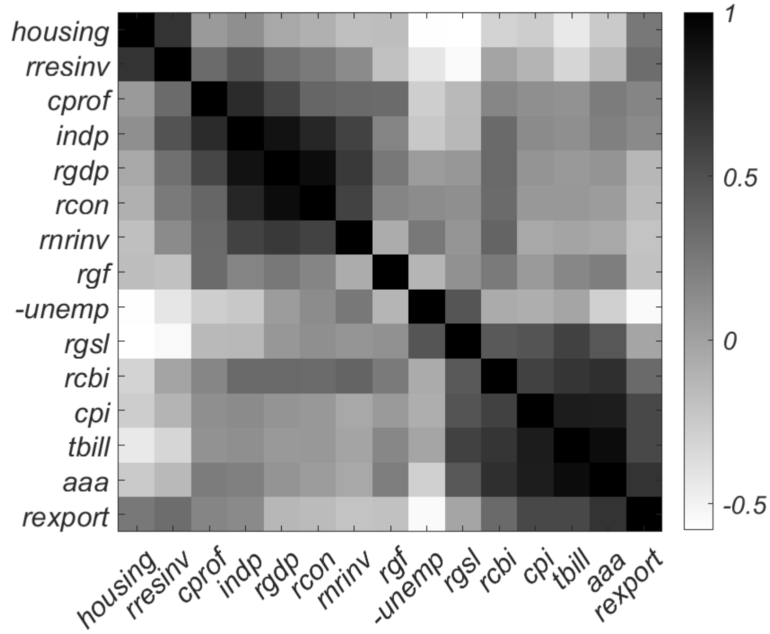

Figure 1. Correlation matrix of expectations. This figure shows the pairwise correlation of the survey expectations of 15 variables in the Survey of Professional Forecasters for 1982Q4-2022Q2. The variables are grouped using a simple hierarchical tree method. We use negative unemployment in the correlation matrix so that it is positively related to other typical business-cycle variables.

tions estimate the bias for X_{t+1} based on the objective projection of X_{t+1} onto $\hat{X}_{t+1}$ (i.e., $Cov\left(\varepsilon_{t+1}, \hat{\varepsilon}_{t+1}\right)' \hat{\Sigma}^{-1}$) and the observed bias for $\hat{X}_{t+1}$ (i.e., $E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]$). Second, according to Lemma 4, the difference between synthetic expectations and survey expectations depends on the unspanned variation u_{t+1} in X_{t+1}. For elements of X_{t+1} that are largely spanned by $\hat{X}_{t+1}$, synthetic expectations will closely mirror the survey expectations. This means that survey expectations for large sets of variables X_{t+1} can be explained by a low-dimensional SBF $\hat{S}_{t+1}$, so long as $\hat{S}_{t+1}$ is based on variables $\hat{X}_{t+1}$ that span most of the variation in X_{t+1}.

We use Figure 1 to gauge the size of our subset $\hat{X}_{t+1}$. The pairwise correlations are clustered by a hierarchical tree method, which iteratively merges clusters based on their similarity.[12] Two natural clusters of variables arise with similar correlations. This motivates

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

us to summarize the expectations data using a distortion based on two variables.[13] 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.[14] Figure A5 shows the estimated biases for RGDP and T-bill and the associated log SBF $\hat{s}_{t+1}$, and Appendix E provides a discussion of the properties of these series.

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 (65.8%) 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 (12), 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, note that equation (12) can be rewritten as

$$\hat{E}_t^*[X_{t+1}] - E_t[X_{t+1}] \quad = \quad Cov(\varepsilon_{t+1}, \hat{\varepsilon}_{t+1})' \hat{\Sigma}^{-1} \left(E_t^* \left[\hat{X}_{t+1} \right] - E_t \left[\hat{X}_{t+1} \right] \right). \tag{13}$$

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 (47.4%) of the variation in survey biases can be explained by biased expectations of just two variables, RGDP growth and the T-bill rate $\left(E_t^* \left[\hat{X}_{t+1} \right] - E_t \left[\hat{X}_{t+1} \right] \right)$, and the objective covariance of shocks to the remaining 13 variables with shocks to RGDP growth and the T-bill rate $\left(Cov(\varepsilon_{t+1}, \hat{\varepsilon}_{t+1})' \hat{\Sigma}^{-1} \right)$. If we focus on the 13 variables other than RGDP and T-bill, we still explain 39.3% of the variation in survey biases and 60.6% of the variation in the survey expectations. Figure A1 shows the detailed correlations between synthetic expectations and survey expectations for each variable, as well as the correlations between the synthetic bias and the survey bias.

[13]The 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.

[14]In Section III.C, we discuss alternative pairs of variables that perform similarly well. The key is not the specific variables themselves, but what the variables span. Because of this, replacing either RGDP or T-bill with a highly correlated alternative gives similar results.

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 forecasts for 15 variables from the Survey of Professional Forecasters. Column 1 shows the average R^2 from regressions of survey expectations ($E_t^*[X_{j,t+1}]$) on synthetic expectations ($\hat{E}_t^*[X_{j,t+1}]$) for each of the 15 different variables. Column 2 shows the average R^2 from regressions 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 (14) and Columns 4 and 5 show the explanatory power of the first two principal components of the 15 biases $E_t^*[X_{t+1}] - E_t[X_{t+1}]$.

	$E_t^*[X_{t+1}]$	$E_t^*[X_{t+1}] - E_t[X_{t+1}]$			
	$\hat{E}_t^*[X_{t+1}]$	$\hat{E}_t^*[X_{t+1}] - E_t[X_{t+1}]$	Best Linear Predictor	PC-1	PC-2
$R^2(\%)$	65.8	47.4	52.7	43.2	64.5

To gauge the performance of our synthetic expectations, we compare our results to less restrictive versions of equation (13). The most straightforward test is a regression

$$E_t^*[X_{t+1}] - E_t[X_{t+1}] \quad = \quad \alpha + \Gamma\left(E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]\right) + \eta_t. \tag{14}$$

This is identical to equation (13), 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, and thus, provides an upper bound of how much one can explain *given these two variables*. Averaging across all variables, we find that this regression approach gives an R^2 of 52.7%. Therefore, our estimated SBF $\hat{S}_{t+1}$ gives an average R^2 (47.4%) that is quite close to this upper bound (52.7%).

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 $E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]$ and a fixed matrix $Cov(\varepsilon_{t+1}, \hat{\varepsilon}_{t+1})' \hat{\Sigma}^{-1}$ based on the objective covariances. The final two columns in Table I show a more general upper bound based on principal components analysis (PCA),

$$E_t^*[X_{t+1}] - E_t[X_{t+1}] \quad = \quad \alpha + \Gamma\Lambda_t. \tag{15}$$

The first two principal components of the 15 biases explain 64.5% of the variation, and the first principal component explains 43.2%. This means that the estimated SBF performs better than the first principal component of these series and captures roughly three fourths of the maximum possible R^2, 64.5%.

Before moving to a more detailed analysis of these 15 synthetic expectations, it is important to emphasize the key benefits of the SBF relative to the benchmark of PCA. First, there is the typical trade-off of performance versus economic interpretability. The SBF depends on biases in two familiar variables (RGDP and T-bill) whereas the Λ_t of equation (15) are artificially constructed combinations of all 15 biases. Similarly, the SBF-related loadings in equation (13) simply depend on the objective covariance of RGDP and T-bill with the other variables. For example, a positive 1 pp bias in expected RGDP growth causes a negative 0.51 pp synthetic bias in expected unemployment because of the objective relationship between higher RGDP growth and lower unemployment. In comparison, the PCA loadings Γ in equation (15) are less straightforward. Thus, while PCA by definition gives the highest possible R^2, the SBF is arguably easier to translate to an economic model.

The second and more important benefit is that once $\hat{S}_{t+1}$ is estimated, it can be applied to new variables for which the researcher does not have survey data. In Section III.B, we show that even if we do not have survey data for new variables Y_{t+1}, we can form statistical expectations $E_t\left[Y_{t+1}\right]$ and then use the estimated SBF to form synthetic expectations $E_t\left[\hat{S}_{t+1}Y_{t+1}\right]$. Similarly, Section V considers a model in which outcomes depend on subjective expectations of complicated objects such as the product of returns and preferences $\tilde{M}_{t+1}R^e_{j,t+1}$ for different assets j. Once again, we can use the estimated SBF to calculate synthetic expectations $E_t\left[\hat{S}_{t+1}\tilde{M}_{t+1}R^e_{j,t+1}\right]$.

Table II and Figure 2 show a more detailed assessment of how well the synthetic biases align with the survey biases for our 15 variables. For each variable j, Table II shows the regression

$$E^*_t\left[X_{j,t+1}\right] - E_t\left[X_{j,t+1}\right] = a_j + b_j\left(\hat{E}^*_t\left[X_{j,t+1}\right] - E_t\left[X_{j,t+1}\right]\right) + \eta_{j,t} \qquad (16)$$

Table II

Matching the Survey of Professional Forecasters

This table evaluates the ability of each individual synthetic bias to match the survey bias of the 15 variables in the Survey of Professional Forecasters. The table shows the intercept and slope coefficients from regressing $E_t^*[X_{j,t+1}] - E_t[X_{j,t+1}]$ on $\hat{E}_t^*[X_{j,t+1}] - E_t[X_{j,t+1}]$, where $\hat{E}_t^*[X_{j,t+1}]$ is the synthetic expectation constructed using the log SBF $\hat{s}_{t+1}$. The last row shows a joint regression across all variables except for *rgdp* and *tbill*. In the joint regression, the survey and synthetic bias for each variable are scaled by the standard deviation of the survey bias to ensure comparability across variables. The fifth column shows the R^2 of the regression, and the last column shows the R^2 when the slope of the regression b is constrained to be 1. The row labeled "Average" shows the average R^2 and constrained R^2 across the first 15 rows. Superscripts indicate significance at the 1% (∗∗∗), 5% (∗∗), and 10% (∗) level. The sample period is 1982Q4-2022Q2.

$X_{j,t+1}$	a	se_a	b	se_b	R^2	constrained R^2
rgdp	0.000	—	1.000	—	1.000	1.000
rcon	-0.002**	(0.001)	0.841***	(0.071)	0.748	0.722
cpi	-0.002	(0.003)	1.847**	(0.823)	0.156	0.123
unemp	0.000	(0.001)	0.709***	(0.151)	0.377	0.314
indp	0.000	(0.002)	1.235***	(0.078)	0.765	0.737
tbill	0.000	—	1.000	—	1.000	1.000
aaa	0.004***	(0.001)	0.935***	(0.308)	0.249	0.247
rnrinv	-0.021***	(0.006)	1.443***	(0.126)	0.622	0.564
rresinv	0.004	(0.010)	2.290***	(0.474)	0.316	0.215
rgf	0.009**	(0.004)	2.672***	(0.577)	0.350	0.213
rgsl	0.002	(0.001)	1.147***	(0.267)	0.182	0.179
housing	0.039*	(0.022)	1.411***	(0.331)	0.158	0.144
rcbi	-0.000	(0.000)	1.031***	(0.151)	0.533	0.533
rexport	0.003***	(0.001)	2.064***	(0.387)	0.287	0.211
cprof	-0.025*	(0.014)	1.255***	(0.161)	0.360	0.345
Average					0.474	0.437
Joint	0.065	(0.049)	1.033***	(0.061)	0.334	0.334

Figure 2 visualizes the regression outcomes, plotting the survey bias and the fitted bias using synthetic expectations. For all 15 variables, we find a statistically significant relationship between the survey bias and the synthetic bias. One may be concerned that the R^2s for these regressions are driven by the fact that the slope coefficient is sometimes significantly different from 1, meaning that the survey bias may be an exaggerated or dampened version of the synthetic bias. To address this possibility, the final column of the table shows the R^2 of the regression when b is constrained to be 1. On average, this constraint only slightly reduces the average R^2, from 0.474 to 0.437. Section III.C.1 provides more details on what determines how well synthetic expectations perform for each variable.

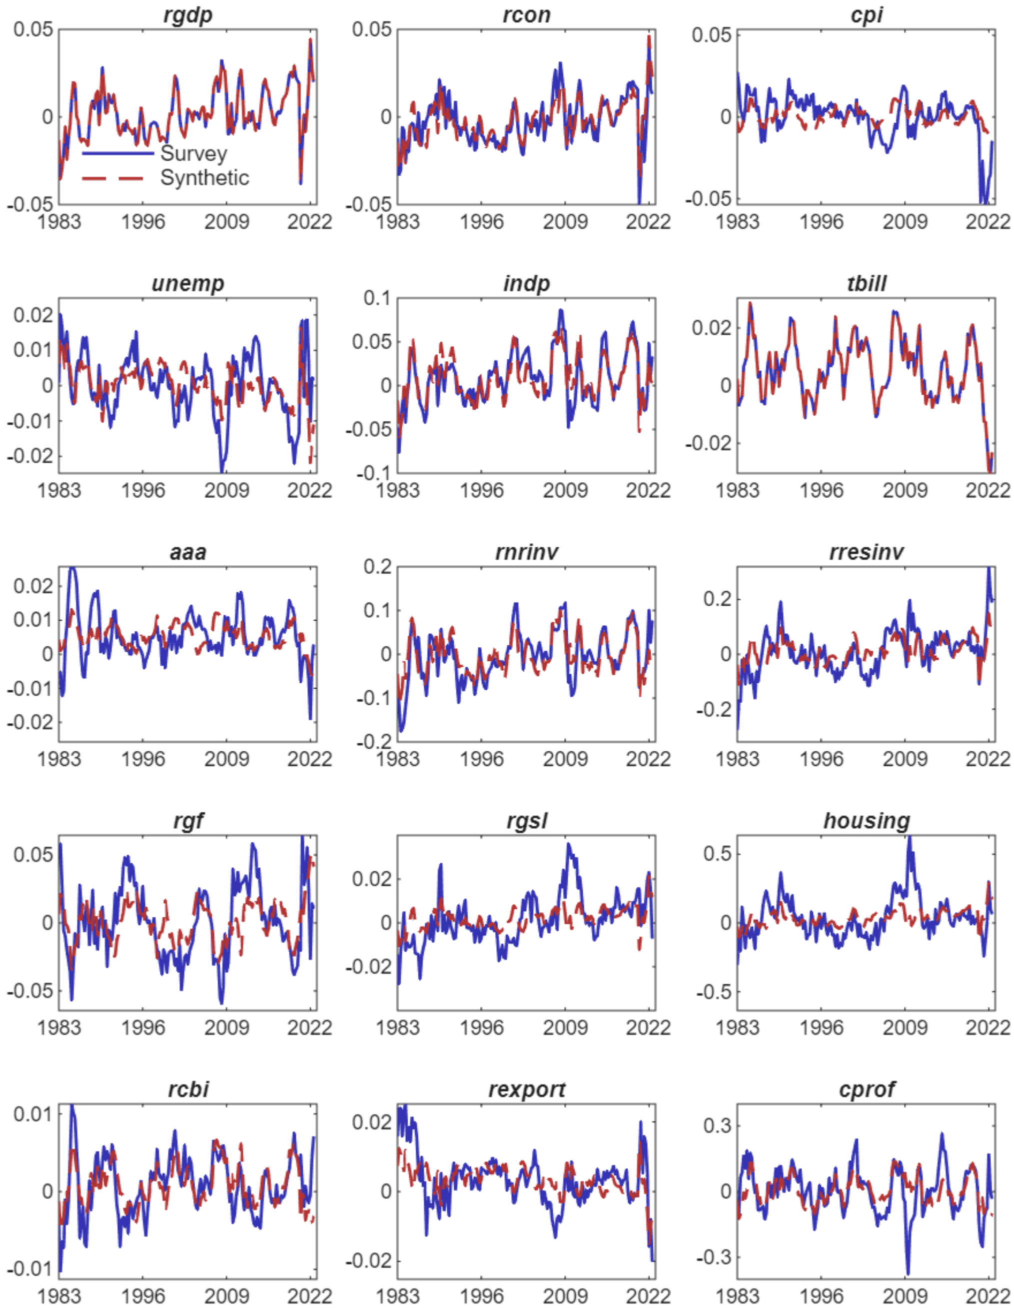

Figure 2. Comparison between synthetic biases and survey biases. This figure shows how well the synthetic bias $\hat{E}_t^*[X_{j,t+1}] - E_t[X_{j,t+1}]$, constructed using the SBF $\hat{s}_{t+1}$, matches the survey expectation bias $E_t^*[X_{j,t+1}] - E_t[X_{j,t+1}]$ for each of the 15 variables in the Survey of Professional Forecasters. The blue time series show each of the survey biases, and the red time series show the fitted synthetic biases from equation (16). The sample period is 1982Q4-2022Q2.

Importantly, the fact that R^2 does not equal 1 for all variables does not mean that these survey forecasters have non-coherent expectations, i.e., expectations that violate basic rules of addition and multiplication. Rather, it indicates that our two-factor SBF does not completely span the variation in all 15 variables. For example, incorporating a third factor into our SBF related to housing would improve our ability to match the housing starts expectations. Mechanically, a 15-factor log SBF $\hat{s}_{t+1} = -\frac{1}{2}\beta_t'\Sigma\beta_t + \beta_t'\varepsilon_{t+1}$ with $\beta_t = \Sigma^{-1}\left(E_t^*\left[X_{t+1}\right] - E_t\left[X_{t+1}\right]\right)$ would perfectly match all of the survey expectations. Just like attempts to find low-dimensional SDF's in asset pricing research, incorporating more factors into the SBF improves the explanatory power but at the cost of parsimony. Our goal is not to say that two factors is the perfect balance between parsimony and explanatory power, but rather to highlight the benefits of applying this factor approach to condense subjective expectations across multiple variables.

The final row of Table II shows a joint regression across the variables. In this joint regression, we find an insignificant intercept of (0.07) and a highly significant slope (1.03^{***}) close to unity. Note that this joint regression excludes RGDP and T-bill, so this estimate of $a \approx 0$ and $b \approx 1$ is not driven by the fact that synthetic biases perfectly match survey biases for these two variables. Additionally, for the joint regression, we scale the survey bias and synthetic bias for each variable j by the standard deviation of $E_t^*\left[X_{j,t+1}\right] - E_t\left[X_{j,t+1}\right]$ to ensure that our results are not driven by the synthetic bias simply matching the survey bias for a single variable with a highly volatile bias, i.e., all of the survey biases are standardized to have unit variance which gives them equal weight in the joint regression.

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 our Survey of Professional Forecasters sample 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 an autoregressive model for Y_{t+1},[15]

$$Y_{t+1} = a_y + B_y \begin{pmatrix} Y_t & X_t & E_t^* [X_{t+1}] \end{pmatrix} + \varepsilon_{y,t+1} \quad (17)$$

$$E_t [Y_{t+1}] = a_y + B_y \begin{pmatrix} Y_t & X_t & E_t^* [X_{t+1}] \end{pmatrix}. \quad (18)$$

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 (17). 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

$$\begin{aligned} \hat{E}_t^* [Y_{t+1}] &\equiv E_t \left[\hat{S}_{t+1} Y_{t+1} \right] \\ &= E_t [Y_{t+1}] + Cov (\varepsilon_{y,t+1}, \hat{\varepsilon}_{t+1})' \hat{\beta}_t \\ &= E_t [Y_{t+1}] + Cov (\varepsilon_{y,t+1}, \hat{\varepsilon}_{t+1})' \hat{\Sigma}^{-1} \left(E_t^* \left[\hat{X}_{t+1} \right] - E_t \left[\hat{X}_{t+1} \right] \right). \end{aligned} \quad (19)$$

Importantly, we do not use any survey expectations of Y_{t+1} in the construction of $\hat{E}_t^* [Y_{t+1}]$. Just as in equation (12), the synthetic expectations for Y_{t+1} will be based on the observed

[15]We find similar results if we estimate the same process for the log of the financial variables rather than the level.

bias for $\hat{X}_{t+1}$ and the projection of Y_{t+1} onto $\hat{X}_{t+1}$.

How well do these synthetic expectations match actual subjective expectations for these nine variables? Figure A3 compares the survey bias $E_t^*[Y_{t+1}] - E_t[Y_{t+1}]$ with the synthetic bias $\hat{E}_t^*[Y_{t+1}] - E_t[Y_{t+1}]$ for each variable. Table AIII shows the detailed regression outcomes analogous to Table II. Across the variables, we find significant slope coefficients b that are close to 1. On average, the R^2 from these regressions is 0.681. This is similar to the performance of the synthetic expectations for the Survey of Professional Forecasters variables in Table II (0.474), 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. Figure A2 shows the detailed correlations between synthetic expectations and survey expectations for each variable, as well as the correlations between the synthetic bias and the survey bias.

To push this idea of jointly explaining the Survey of Professional Forecasters and the Blue Chip forecasts further, Table III 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 the survey expectations, with an average R^2 of 75.2%. 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, 54.3%.

The last three columns of Table III compare our results to the upper bounds implied by the best linear predictor and PCA. Just as in equation (13), 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,

$$E_t^*[Z_{t+1}] - E_t[Z_{t+1}] \quad = \quad \alpha_z + \Gamma_z\left(E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]\right) + \eta_{z,t} \tag{20}$$

Table III

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 expectations $E_t^*[Z_{t+1}]$. Column 1 shows the average R^2 from regressions of the survey expectations $E_t^*[Z_{j,t+1}]$ on $\hat{E}_t^*[Z_{j,t+1}]$ for each of the 24 different variables. Column 2 shows the average R^2 from regressions of the survey bias $E_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$ on the synthetic bias $\hat{E}_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$. 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 (20) 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
$R^2(\%)$	75.2	54.3	58.6	48.1	67.4

where the coefficient matrix Γ_z is unrestricted. We find that the average R^2 produced by the synthetic expectations (54.3%) is quite close to the upper bound implied by the best linear predictor of 58.6%.

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

$$E_t^*[Z_{t+1}] - E_t[Z_{t+1}] = \alpha_z + \Gamma_z \Lambda_{z,t}. \tag{21}$$

The fourth column of Table III shows that synthetic expectations outperform the first principal component and achieve more than three fourths of the maximum possible R^2 of 67.4%.

Overall, we find that the log SBF (i) accurately predicts 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 by beliefs about a few key variables and the objective relationships between variables.

C. *Evaluating the performance of the SBF*

How does the performance of the SBF relate to how well variables are spanned? Are the results of Sections III.A and III.B simply driven by the macro variables having a low factor structure? What about other pairs of variables? In this subsection, we provide a deeper analysis of these types of questions. First, in line with Lemma 4, we show that our $\hat{S}_{t+1}$ performs best in matching the expectations of variables that are largely spanned by the chosen $\hat{X}_{t+1}$ but that the performance remains relatively high even for variables with substantial unspanned variation. Further, we show that the R^2 from Tables I and III are too high to be explained solely by the underlying 24 variables having a low-factor structure. These results are both consistent with the true SBF S_{t+1} having a low-factor structure. Finally, we provide a direct comparison of the SBF and PCA.

C.1. Performance and spanning conditions

To give a benchmark for our results, we consider a perfect foresight expectation $E_t^{pf}\left[Z_{j,t+1}\right] = Z_{j,t+1}$ for each of the 24 variables. The vertical axis of Figure 3 shows the R^2 for our synthetic expectations, $Corr\left(E_t^*\left[Z_{j,t+1}\right], \hat{E}_t^*\left[Z_{j,t+1}\right]\right)^2$, minus the R^2 for the perfect foresight benchmark, $Corr\left(E_t^*\left[Z_{j,t+1}\right], E_t^{pf}\left[Z_{j,t+1}\right]\right)^2$. In other words, we evaluate how well the synthetic expectations match variation in the survey expectations above and beyond what can be explained by the future realized outcomes.[16] The horizontal axis shows how well each variable's shock $\varepsilon_{j,t+1}$ is spanned by the RGDP and T-bill shocks $\hat{\varepsilon}_{j,t+1}$, i.e., $1 - \frac{Var(u_{j,t+1})}{Var(\varepsilon_{j,t+1})}$ where $u_{j,t+1}$ is the unspanned portion of the shock.[17]

As indicated by Proposition 4 and Lemma 4, the SBF performs particularly well for variables whose shocks are largely spanned by the RGDP and T-bill shocks. For variables such as real consumption and industrial production, the synthetic expectations produce

[16] Figure A4 shows the same exercise using the survey bias $E_t^*\left[Z_{j,t+1}\right] - E_t\left[Z_{j,t+1}\right]$ rather than the survey expectation $E_t^*\left[Z_{j,t+1}\right]$.

[17] We compute $1 - \frac{Var(u_{j,t+1})}{Var(\varepsilon_{j,t+1})}$ by regressing each variable's shock on the RGDP and T-bill shocks and calculating the regression R^2.

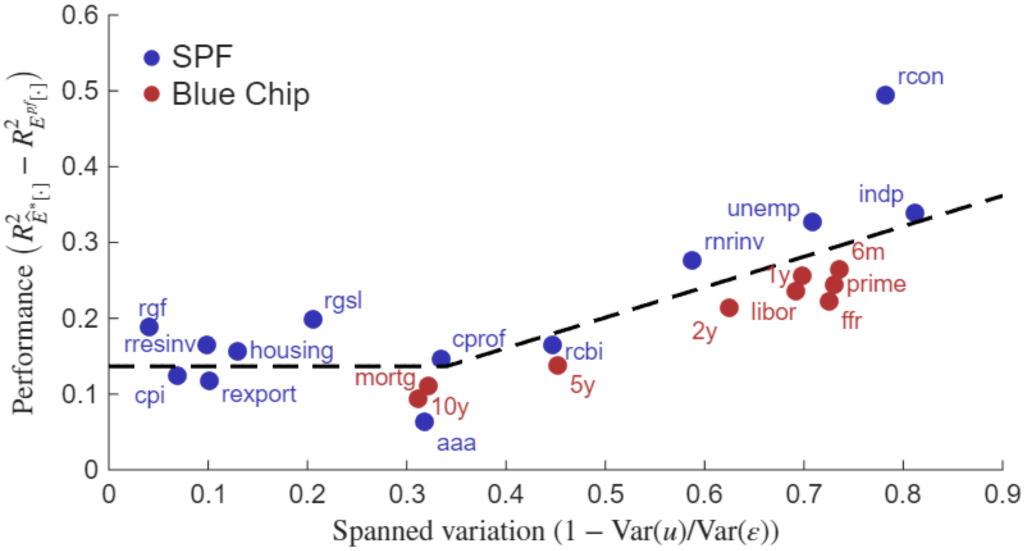

Figure 3. Performance and spanning of synthetic expectations. This figure evaluates how the performance of the synthetic expectations relates to how well each variable is spanned by the RGDP growth and T-bill shocks. The vertical axis shows the R^2 of the synthetic expectations relative to the perfect-foresight benchmark in fitting the survey expectations. The horizontal axis shows the degree to which each variable's shock $\varepsilon_{j,t+1}$ is spanned by the RGDP and T-bill shocks, $1 - Var(u_{j,t+1})/Var(\varepsilon_{j,t+1})$, where $u_{j,t+1}$ is the unspanned portion of each shock. Each point is one of the 22 Survey of Professional Forecasters and Blue Chip variables, excluding RGDP and T-bill. The sample period is 1982Q4-2022Q2.

R^2's that are substantially higher than the benchmark. The performance decreases as we consider variables whose shocks are less spanned by the RGDP and T-bill shocks. However, the performance does not go to 0 even when we consider variables such as real federal government spending where 96% of the variation is unspanned (i.e., $1 - \frac{Var(u_{j,t+1})}{Var(\varepsilon_{j,t+1})}$ is 0.04). Instead, the performance levels off as spanning decreases.

This hockey stick shape in Figure 3 has an intuitive economic interpretation. Proposition 4 and Lemma 4 emphasize that synthetic expectations differ from survey expectations only when two conditions are both present: the variable has variation that is not spanned by $\hat{\varepsilon}_{t+1}$, and the true S_{t+1} also has variation that is not spanned by $\hat{\varepsilon}_{t+1}$. If the true SBF S_{t+1} has a low-factor structure and is relatively well-spanned by $\hat{\varepsilon}_{t+1}$, then synthetic expectations will be accurate even for variables whose shocks are not well-spanned by $\hat{\varepsilon}_{t+1}$.

As a concrete example, if the true SBF was perfectly spanned by $\hat{\varepsilon}_{t+1}$, then the survey expectations and synthetic expectations would perfectly match, even for variables with a large amount of variation that is unrelated to $\hat{\varepsilon}_{t+1}$. This is because variation in the variable that is not related to the true SBF is irrelevant for biases (i.e., $E_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}] = Cov_t(S_{t+1}, Z_{j,t+1})$). Taking the example from Section I.B.1, an agent that extrapolates RGDP will also have biased government spending expectations, but only to the extent that government spending depends on RGDP; the agent will be unbiased about the variation in government spending that is unrelated to RGDP.

To summarize, the performance of the synthetic expectations depends on two forms of spanning: how well RGDP and T-bill shocks span the true SBF, and how well RGDP and T-bill shocks span the shocks to the other variables. The first component allows synthetic expectations to match survey expectations moderately well even for variables whose shocks are poorly spanned by the RGDP and T-bill shocks. The second component allows the synthetic expectations to perform particularly well for variables whose shocks are largely spanned by the RGDP and T-bill shocks. To show this visually, Figure 3 includes the best fit piecewise function $\max\left\{a, b\left(1 - \frac{Var(u_{j,t+1})}{Var(\varepsilon_{j,t+1})}\right)\right\}$. If the true SBF was poorly spanned by $\hat{\varepsilon}_{t+1}$ then a would be 0 and the performance of the synthetic expectations would go to 0 as $1 - \frac{Var(u_{j,t+1})}{Var(\varepsilon_{j,t+1})}$ approaches 0.

C.2. Performance relative to univariate models

As a second, more general benchmark, we consider whether the observed performance could arise in univariate models of expectation biases. Specifically, consider a model where the bias for each variable j is driven solely by its own shocks:

$$E_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}] = \sum_{\tau=0}^{\infty} \delta_{j,\tau}\varepsilon_{j,t-\tau} \tag{22}$$

where the shocks $\varepsilon_{j,t}$ are serially uncorrelated but can be contemporaneously correlated across variables. For example, forecasters may overreact to export shocks when forecasting exports, underreact to inflation shocks when forecasting inflation, etc. Even though the bias for each variable only depends on its own shocks, if the shocks are correlated across variables then the biases may also be correlated across variables and one may still find a low factor structure for biases. To make this benchmark as general as possible, we allow the bias for each variable to depend on its entire history of shocks and we place no restrictions on the lag coefficients $\{\delta_{j,\tau}\}_{\tau=0}^{\infty}$.

Proposition 5. *Let $\hat{\varepsilon}_t$ denote the subset of shocks used to construct the synthetic biases. If biases follow equation (22), then for each variable j, the R_j^2 from regressing the actual bias $E_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$ on the synthetic bias $\hat{E}_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$ has the following upper bound*

$$R_j^2 \leq \frac{\mathrm{Cov}(\varepsilon_{j,t+1}, \hat{\varepsilon}_{t+1})' \hat{\Sigma}^{-1} \mathrm{Cov}(\varepsilon_{j,t+1}, \hat{\varepsilon}_{t+1})}{\mathrm{Var}(\varepsilon_{j,t+1})}. \tag{23}$$

This upper bound is precisely $1 - \frac{Var(u_{j,t+1})}{Var(\varepsilon_{j,t+1})}$, i.e., how well spanned the shock $\varepsilon_{j,t+1}$ is by $\hat{\varepsilon}_{t+1}$. Appendix B contains the proof of this proposition and Appendix A.5 provides sufficient conditions to maintain the same bound in a setting with serially correlated shocks.

We compare this upper bound to the actual R^2 from regressing each variable's survey bias on its synthetic bias. Among the 22 Survey of Professional Forecasters and Blue Chip variables (excluding RGDP and T-bill themselves), 17 of them exceed their respective bounds.[18] This pattern shows that the factor structure in Sections III.A and III.B is consistently too strong for even this generalized benchmark to match. For example, although shocks to the interest rate variables in Section III.B are highly correlated with the T-bill shock, the upper bound is still exceeded for all nine variables, indicating that forecasters have biased beliefs about a common interest-rate component that goes beyond the correlation of their shocks.

C.3. Alternative pairs of variables

A common theme among these results is that it is not the specific variables chosen in $\hat{\varepsilon}_{t+1}$ that matter but what $\hat{\varepsilon}_{t+1}$ spans. Alternative variables that are highly correlated with RGDP and T-bill should deliver similar results. To confirm this, we consider all possible pairs of

[18] Table AV shows a comparison of each variable's R^2 against its upper bound.

Table IV

Top-performing alternative variable pairs for the SBF.

This table evaluates the ability of synthetic expectations to explain forecasts from the Survey of Professional Forecasters and Blue Chip. For each column, 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 $\hat{E}_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$, shown in the first row. The second and third rows show the pair of variables used for each alternative log SBF $\hat{s}_{t+1}$.

Rank	1	2	3	4	5	6	7	8	9	10
Average R^2 (%)	59.0	58.8	58.3	58.3	58.2	58.0	57.9	57.6	57.3	57.2
Variables	*rgdp* *1y*	*rgdp* *2y*	*rgdp* *6m*	*rcon* *1y*	*rcon* *2y*	*rnrinv* *1y*	*rnrinv* *2y*	*indp* *2y*	*rnrinv* *6m*	*indp* *1y*

variables for the Survey of Professional Forecasters and Blue Chip. Table IV shows the ten pairs that generate the highest average R^2 for the survey biases for the 24 combined variables. The general pattern across all of these pairs is that it is important to have one variable related to economic activity and one variable related to yields. While the top pairs use RGDP and 6-month to 2-year yields, rather than T-bill, we use RGDP and T-bill for our main SBF because (i) T-bill expectations are available for the longer Survey of Professional Forecasters sample, (ii) the choice was made ex-ante based on economic interpretation and the clustering evidence in Figure 1 rather than selected from the 276 possible pairs ex-post, and (iii) the average R^2 of 0.543 from Table III is quite close to the results for these top ex-post pairs.

C.4. Information content relative to PCA

The final point worth evaluating is how the SBF compares to PCA. As shown in Table III, both the two-factor SBF and the first two principal components capture a large portion of the variation in survey biases for the combined 24 variables. By construction, the performance of the principal components must exceed the performance of the SBF, as PCA utilizes information from all 24 biases whereas the SBF is based solely on information on RGDP and T-bill biases. Does this mean that the SBF is a noisy proxy for the first two principal components or does the SBF contain information that is not captured by them?

Table V

Informational overlap between Subjective Belief Factor and PCA.

This table evaluates whether the SBF contains information beyond the first two principal components of the survey biases, and vice versa. For each of the 24 Survey of Professional Forecasters and Blue Chip variables, we regress the survey bias on the synthetic bias and on the first two principal components of the 24 survey biases, and use F-tests based on differences in R^2 to assess incremental explanatory power. The column "SBF over PCA" reports whether the synthetic bias adds explanatory power beyond the two principal components; the column "PCA over SBF" reports whether the two principal components add explanatory power beyond the synthetic bias. The row "% variables (5% sig.)" shows the share of the 24 variables for which the corresponding null (i.e., no extra explanatory power) is rejected at the 5% level. The row "p-value on Average R^2" shows the p-value of the aggregate test based on whether the average R^2 is significantly different.

Metric	SBF over PCA	PCA over SBF
% variables (5% sig.)	75.0%	87.5%
p-value on Average R^2	3×10^{-5}	6×10^{-14}

Table V addresses this question. For each of the 24 variables $Z_{j,t+1}$, we consider the regression

$$E_t^*\left[Z_{j,t+1}\right] - E_t\left[Z_{j,t+1}\right] \quad = \quad a_j + b_j\left(\hat{E}_t^*\left[Z_{j,t+1}\right] - E_t\left[Z_{j,t+1}\right]\right) + \gamma_{1,j}\Lambda_{t,1} + \gamma_{2,j}\Lambda_{t,2} + \eta_{j,t}$$

where $\Lambda_{t,1}$ and $\Lambda_{t,2}$ are the first two principal components of the 24 survey biases from equation (21). For each variable, we estimate the incremental explanatory power of the synthetic expectations and the two principal components using standard F-tests based on differences in R^2. Specifically, we test whether allowing $b_j \neq 0$ significantly increases explanatory power, i.e., whether the SBF contributes information beyond what is reflected in the first two principal components.

The first row and first column in Table V report the share of variables for which the SBF significantly increases the R^2 relative to the model with only the first two principal components. This null is rejected at the 5% level for 75% of the variables. Analogously, the second column shows that for 87.5% of the variables, including the principal components into the regression (i.e., allowing $\gamma_{1,j} \neq 0$ and $\gamma_{2,j} \neq 0$) significantly increases the R^2 compared to only using the synthetic expectations. The second row of Table V provides the average analogue of these comparisons. Averaging the R^2 across the 24 variables, the first column shows the p-value for the null that including synthetic expectations does not increase the R^2. The second column shows the same test but for the first two PCs. Both nulls are clearly rejected.

These results highlight that the SBF captures information that is not included in the first two principal components. As discussed in Section III.A, PCA is the best in sample method for summarizing the 24 survey biases. The SBF is more restrictive, but it has two advantages. First, it is economically interpretable: it depends on biases in RGDP growth and the T-bill rate and on the objective covariance of shocks. Second, it can be applied to variables for which no survey expectations are available. This was the key exercise in Section III.B, where we formed synthetic expectations for the Blue Chip variables without using any survey data from Blue Chip. Further, in Section V, we can calculate synthetic expectations for model objects such as $E_t^*\left[\tilde{M}_{t+1}R^e_{j,t+1}\right]$ and $E_t^*\left[R^e_{j,t+1}\right]$ by calculating the objective expectations and incorporating the SBF, $E_t\left[\hat{S}_{t+1}\tilde{M}_{t+1}R^e_{j,t+1}\right]$ and $E_t\left[\hat{S}_{t+1}R^e_{j,t+1}\right]$. Once the SBF has been estimated, it can be incorporated into any objective expectation to create a synthetic subjective expectation.

D. Testing the coherence assumption

The results of the previous sections already provide indirect support for coherence, in the sense that synthetic expectations (which are coherent by construction) match the survey expectations relatively well. In this subsection, we provide a more direct test of coherence.

Testing coherence requires observing expectations for variables linked by a linear identity. We use the national income accounting identity for real GDP, which is included in the Survey of Professional Forecasters expectations since 1992Q1.[19] Specifically, for each quarter in our sample, we compare the Survey of Professional Forecasters forecast for real GDP against the sum of forecasts for its components: real personal consumption, real nonresidential fixed investment, real residential investment, real federal government expenditures, real state and local government expenditures, real change in private inventories, and real net exports. We define the coherence residual as the difference between the reported real GDP forecast and

[19] Previous quarters use real GNP rather than real GDP.

the sum of the forecasts of the components, expressed as a fraction of the real GDP forecast, and we report statistics on its absolute value and range. We distinguish between two subperiods because a structural change in the BEA's methodology breaks the identity starting in 1996.

Prior to 1996, the BEA computed real GDP using a fixed-weight price index, under which the identity holds exactly: realized real GDP is the precise sum of the realized components. Using the subjective expectations, the average absolute residual is only 5.2 bps. In other words, the forecasts for the seven components sum almost exactly to the forecast for real GDP, with only a small residual which may simply be due to rounding across the seven components since forecasters cannot report an unlimited number of digits in their forecasts. The range for the residual is also fairly small, with the standard deviation of the residual being 6.6 bps and the maximum absolute residual being 14.4 bps.

Starting in 1996, the BEA adopted chain-weighting. Under chain weighting, the relative prices used to aggregate components vary over time, so real GDP is no longer the exact sum of its real components. To give a benchmark, for realized real GDP and the realized components, the mean absolute residual in the post 1996 period is 20.2 bps, with a standard deviation of 27.7 bps and a maximum residual of 102 bps. Importantly, the subjective expectations of the components match the subjective expectations for real GDP to a similar degree of accuracy as the realized data over the same period, with the residuals having an average absolute value of 14.8 bps, a standard deviation of 21.2 bps, and a maximum absolute value of 90.5 bps. Thus, in a setting where we have an approximate identity for the realized data, we find a similarly approximate identity for the subjective expectations.

It is important to emphasize that coherence tests should rely on objective identities rather than economic relationships that models may assume for simplicity. Adding assumed economic relationships transforms the exercise into a joint test of coherence and the accuracy of such relationships. For example, one might consider evaluating whether subjective expectations satisfy a Taylor rule linking yields, inflation, and output (e.g., Bauer et al. 2024b).

However, because the Taylor rule is not an identity, deviations between these expectations would not necessarily imply incoherent beliefs. Instead, they might simply imply a rejection of the Taylor rule being an exact relation.

As a detailed example, consider the analysis in Shue et al. (2024). They document that expected changes in long-term interest rate forecasts (specifically the 10-year Treasury yield and the mortgage rate) comove with expected changes in the short-term rate with OLS coefficients of approximately 0.38 and 0.29, respectively. One natural interpretation is that this comovement reflects a lack of coherence: if beliefs were internally consistent, the long rate, being approximately the sum of expected future short rates, should have already incorporated any expected change in short rates and the OLS coefficients should be very small values of $1/n$, where n is the horizon of the long rate. However, this relationship only holds if term premia in yields are constant. Interpreting their regressions as a pure coherence test therefore requires the expectations hypothesis to hold exactly, which it does not (Shiller, 1979). Thus, their results could represent either violations of coherence or violations of constant term premia.

Interestingly, once we consider time-varying term premia, their results are quite consistent with our empirical estimates of comovements under coherence. Using the Blue Chip data and our two factors of RGDP and T-bill of Section III.B, 10-year yield expectations are predicted to comove with the short-term T-bill expectations with a coefficient $Cov\left(\varepsilon_{y,t+1}, \hat{\varepsilon}_{t+1}\right)' \hat{\Sigma}^{-1}$ of 0.36. Mortgage rate expectations are predicted to have a coefficient of 0.35. These two-factor implied values are close to Shue et al. (2024)'s empirical estimates of 0.38 and 0.29, suggesting that their findings of coefficients noticeably larger than $1/n$ could be explained by coherent beliefs under a low-factor structure rather than a failure of coherence.

IV. Predictable forecast errors and expectation models

This section evaluates how well our SBF replicates a number of untargeted moments of forecast error predictability, namely (i) the wide range of serial autocorrelation of forecast errors across variables, (ii) the wide range of Coibion-Gorodnichenko (CG) coefficients across variables, (iii) decreasing CG coefficients across maturities, and (iv) cross-variable CG coefficients. While our results are non-causal, our findings highlight that a two-factor SBF provides a parsimonious explanation for multiple empirical puzzles regarding forecast error predictability.

We focus on the predictability of forecast errors using lagged forecast errors (i.e., the serial correlation of forecast errors), and the predictability of forecast errors using changes in expectations (i.e., the CG coefficient). Let $F_{j,t+1} \equiv Z_{j,t+1} - E_t^*[Z_{j,t+1}]$ be the forecast error for each variable j. We measure the coefficients $b_{j,1}$ and $b_{j,2}$ from regressions

$$F_{j,t+1} = a_{j,1} + b_{j,1} F_{j,t} + \eta_{j,1,t+1} \tag{24}$$

$$F_{j,t+1} = a_{j,2} + b_{j,2}\left(E_t^*[Z_{j,t+1}] - E_{t-\Delta}^*[Z_{j,t+1-\Delta}]\right) + \eta_{j,2,t+1} \tag{25}$$

where $\Delta = \frac{1}{4}$ denotes that we are using the three-month change in the one-year expectations.[20] Following the framework of Wu (2023), we also estimate cross-variable CG coefficients from

$$\begin{aligned} F_{j,t+1} = \; & a_{j,k}^{cross} + b_{j,k}^{cross}\left(E_t^*[Z_{k,t+1}] - E_{t-\Delta}^*[Z_{k,t+1-\Delta}]\right) \\ & + c_{j,k}^{cross}\left(E_t^*[Z_{j,t+1}] - E_{t-\Delta}^*[Z_{j,t+1-\Delta}]\right) + \eta_{j,k,t+1}^{cross} \end{aligned} \tag{26}$$

which measure how changes in expectations for variable k predict future forecast errors for variable j, controlling for changes in expectations of variable j.

The magnitudes of the regression coefficients $(b_{j,1}, b_{j,2})$ provide valuable information about

[20] We focus on changes in one-year expectations, rather than revisions $E_t[Z_{t+1}] - E_{t-\Delta}[Z_{t+1}]$ which utilize a one-year expectation and a 15-month expectation. This ensures that the independent variable in equation (25) can be replicated by a single time series for the SBF. We set Δ to the smallest value possible in our data, which is one quarter, to provide the maximum overlap between the two forecasted periods (i.e., $t-\Delta$ to $t+1-\Delta$ and t to $t+1$).

model parameters, such as the speed of information acquisition in learning models or the degree of extrapolation in behavioral models. In particular, positive serial correlation in forecast errors and positive CG coefficients indicate underreaction in subjective expectations, while negative serial correlation and negative CG coefficients indicate overreaction, and the magnitudes of the coefficients indicate the degree of under/overreaction (Coibion and Gorodnichenko, 2015; Bordalo et al., 2020). Across the 24 Survey of Professional Forecasters and Blue Chip variables, we find a wide range of serial correlations, from -0.18 to 0.51, and a wide range of CG coefficients, from -0.29 to 1.26. Further, we find cross-variable CG coefficients that are positively associated with the objective correlation of variables j and k, in line with the central finding of Wu (2023).

Our goal is to evaluate whether the synthetic versions of lagged forecast errors and changes in expectations can predict future forecast errors with similar coefficients. We repeat regressions (24) and (25), but regress future forecast errors $F_{j,t+1}$ onto synthetic versions of the predictors, $\hat{F}_{j,t} \equiv Z_{j,t} - \hat{E}^*_{t-1}[Z_{j,t}]$ and $\hat{E}^*_t[Z_{j,t+1}] - \hat{E}^*_{t-\Delta}[Z_{j,t+1-\Delta}]$. Figure 4 compares $(b_{j,1}, b_{j,2})$ from both sets of regressions. The x-axis shows the results using the synthetic predictors, $\hat{F}_{j,t}$ and $\hat{E}^*_t[Z_{j,t+1}] - \hat{E}^*_{t-\Delta}[Z_{j,t+1-\Delta}]$, and the y-axis shows the results using the survey version of the predictors $F_{j,t}$ and $E^*_t[Z_{j,t+1}] - E^*_{t-\Delta}[Z_{j,t+1-\Delta}]$. Each point in the scatterplot represents one of the 22 Survey of Professional Forecasters and Blue Chip variables, excluding RGDP and T-bill for which we perfectly replicate the forecast error predictability results by construction.

The left panel demonstrates that the predictability of forecast errors using lagged synthetic forecast errors closely matches the predictability using lagged survey forecast errors, with most observations falling near the 45-degree line. The correlation of the synthetic $b_{j,1}$ with the survey $b_{j,1}$ across all variables is 0.81. To ensure that the result is not driven simply by one set of expectations, we show the correlation for $b_{j,1}$ for the 13 Survey of Professional Forecasters variables (excluding RGDP and T-bill), the 9 Blue Chip variables, and the combined 22 variables. The right panel confirms that the CG coefficients $(b_{j,2})$ derived

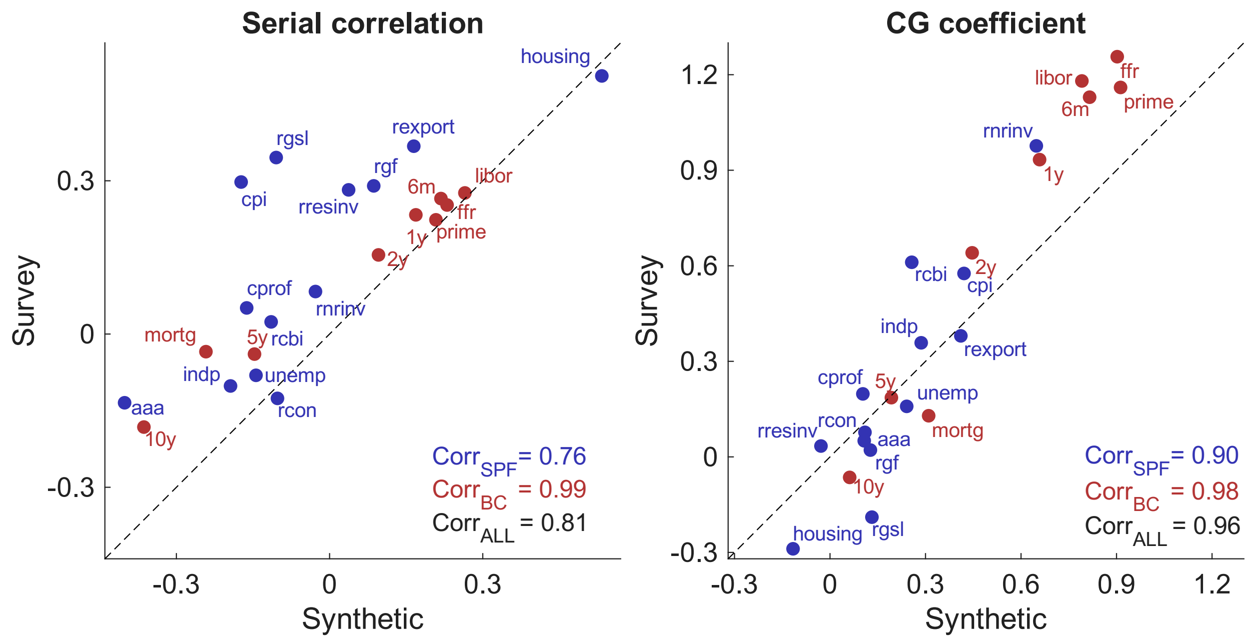

Figure 4. Forecast error predictability tests. The figure evaluates the ability of the synthetic expectations to replicate forecast error predictability patterns in the survey expectations for each of the variables in the Survey of Professional Forecasters and Blue Chip data. The left panel shows the results of the serial correlation test in (24), obtained by regressing forecast errors on lagged forecast errors. The scatter plot compares the coefficients $b_{j,1}$ using survey lagged forecast errors and synthetic lagged forecast errors as the predictor. The right panel shows the Coibion-Gorodnichenko coefficients from (25), obtained by regressing forecast errors on changes in expectations. The scatter plot compares the coefficients $b_{j,2}$ using changes in survey expectations and changes in synthetic expectations as the predictor. In both panels, the dashed line shows the 45-degree line.

from synthetic expectations also align closely with those from survey expectations, with the observations falling near the 45-degree line and a high correlation across all variables of 0.96.

How do these results connect back to RGDP and T-bill? Empirically, RGDP and T-bill differ in their predictability of forecast errors. While T-bill has a positive serial correlation of forecast errors of 0.16 and a positive Coibion-Gorodnichenko coefficient of 0.94, RGDP has a slightly negative serial correlation of forecast errors of -0.13 and a small positive Coibion-Gorodnichenko coefficient of 0.12. Overall, equations (12) and (19) and the results for the synthetic predictors demonstrate that we can replicate the wide range of serial correlations and Coibion-Gorodnichenko coefficients for the other 22 variables simply using the bias in RGDP and T-bill survey expectations and how each variable objectively relates to RGDP and T-bill. For example, underreaction to innovations in T-bill expectations can generate underreaction, no bias, or even overreaction in expectations of $Z_{j,t+1}$ depending on how innovations in $Z_{j,t+1}$ objectively relate to innovations in T-bill.

A useful illustration comes from the Blue Chip expectations for yields at different maturities, shown in red in Figure 4. We see a clear pattern in which underreaction is decreasing with maturity, in line with the findings of d'Arienzo (2020). We find underreaction (i.e., positive serial correlation and positive Coibion-Gorodnichenko coefficients) for short maturities such as LIBOR and 6-month yields. The amount of underreaction monotonically decreases and eventually flips to overreaction as we increase the maturity to 1-year, 2-years, 5-years, and finally 10-years. While synthetic expectations are constructed without using any information from the Blue Chip surveys and only involve Survey of Professional Forecasters expectations for the 3-month yield, we find that synthetic expectations closely match this pattern in forecast error predictability across different maturities.

Does the SBF match the findings of Wu (2023) on the cross-variable CG coefficients? Below, we replicate his central result that variables with higher objective correlation tend to have higher cross-variable CG coefficients.[21] The left panel of Figure 5 shows that the cross-variable CG coefficients tend to be positively related to $Corr\left(Z_{j,t}, Z_{k,t}\right)$ with a significant slope coefficient of 0.59 (t-stat 2.61).[22] In other words, for variables that are positively correlated, changes in expectations of k tend to positively predict forecast errors for j, even after controlling for changes in expectations of j.

The right panel of Figure 5 shows the results using synthetic expectations. We repeat regression (26) but regress forecast errors onto synthetic versions of the predictors $\hat{E}_t^*\left[Z_{k,t+1}\right] - \hat{E}_{t-\Delta}^*\left[Z_{k,t+1-\Delta}\right]$ and $\hat{E}_t^*\left[Z_{j,t+1}\right] - \hat{E}_{t-\Delta}^*\left[Z_{j,t+1-\Delta}\right]$. We find that the synthetic

[21] To fully align with the methodology in Wu (2023), we normalize the cross terms to $b_{j,k}^{sum} = \sqrt{\frac{Var(Z_{j,t})}{Var(Z_{k,t})}} b_{j,k}^{cross} + \sqrt{\frac{Var(Z_{k,t})}{Var(Z_{j,t})}} b_{k,j}^{cross}$ which captures both how expectations for k impact forecast errors for j and vice versa and scales the coefficients to account for differences in volatility across variables.

[22] To account for the fact that the same variable appears in many points in the scatter plot of Figure 5 (e.g., real consumption appears in 23 pairs), we use bootstrap standard errors for the slope coefficients. Specifically, we select random subsets of the 24 variables to ensure that the result is not driven by a single outlier variable.

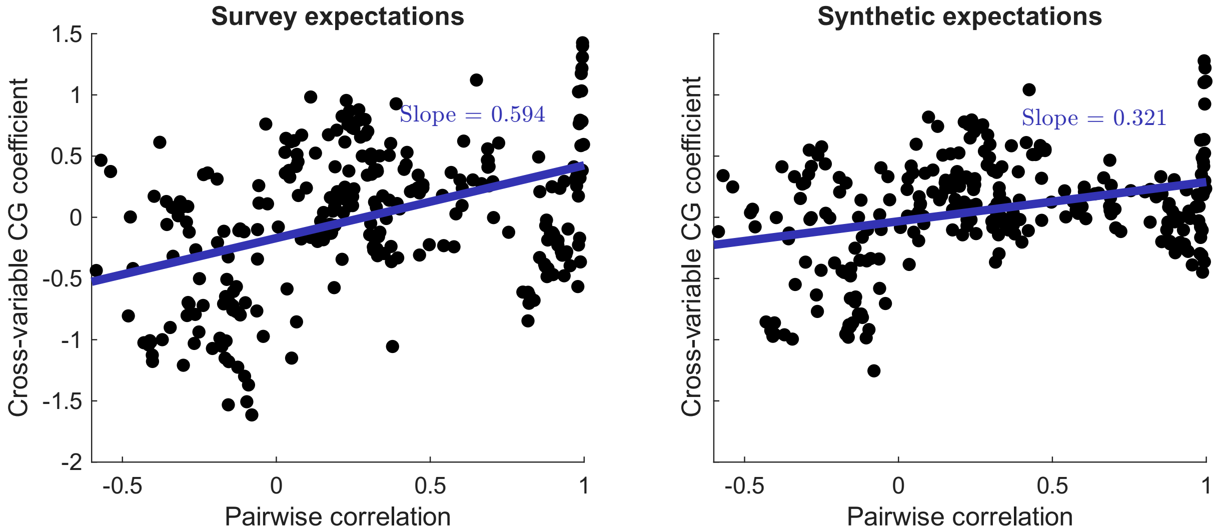

Figure 5. Cross-variable forecast error predictability tests. This figure evaluates the ability of the synthetic expectations to match the cross-variable forecast error predictability in the survey data. Each panel plots, for every pair of variables Z_j and Z_k, the cross-variable Coibion-Gorodnichenko coefficient against their pairwise correlation $Corr(Z_j, Z_k)$. The cross-variable CG coefficient is obtained from equation (26). In the left panel, changes in survey expectations are used as the predictors. In the right panel, changes in synthetic expectations are used as the predictors. The reported slope in blue is from regressing the cross-variable CG coefficients on the pairwise correlations. The sample period is 1982Q4-2022Q2.

cross-variable CG coefficients are also positively associated with $Corr\left(Z_{j,t}, Z_{k,t}\right)$ with a significant slope coefficient of 0.32 (t-stat 2.05). Further, we cannot reject that the slope coefficients in the left and right panels are the same.

In Wu (2023), this result is interpreted as evidence that forecasters misperceive the relationship between $Z_{j,t+1}$ and $Z_{k,t+1}$. Our results offer a different perspective: by construction, synthetic expectations have no misperceptions of the comovements of variables, following equation (12), yet they produce a similar pattern. Instead, we emphasize that this pattern arises naturally when expectations are formed based on a small set of key variables $\hat{X}_{t+1}$.

To give intuition, suppose the agent has biased beliefs about future $\hat{X}_{t+1}$ which then impact her expectations of all other variables,

$$E_t^*\left[Z_{j,t+1}\right]-E_t\left[Z_{j,t+1}\right] \;=\; Cov\left(\varepsilon_{j,t+1},\hat{\varepsilon}_{t+1}\right)\hat{\Sigma}^{-1}\left(E_t^*\left[\hat{X}_{t+1}\right]-E_t\left[\hat{X}_{t+1}\right]\right).$$

The key insight is that the underlying bias depends on $E_t^*\left[\hat{X}_{t+1}\right]-E_t\left[\hat{X}_{t+1}\right]$. Variables that predict future forecast errors, such as changes in expectations $E_t^*\left[Z_{j,t+1}\right]-E_{t-\Delta}^*\left[Z_{j,t+1-\Delta}\right]$ and $E_t^*\left[Z_{k,t+1}\right]-E_{t-\Delta}^*\left[Z_{k,t+1-\Delta}\right]$, will be noisy proxies of $E_t^*\left[\hat{X}_{t+1}\right]-E_t\left[\hat{X}_{t+1}\right]$, which means including both as predictors will help to predict future forecast errors. How the change in expectations for variable k help to predict forecast errors for variable j will depend on how k covaries with $\hat{X}_{t+1}$ compared to how j covaries with $\hat{X}_{t+1}$, which is related to the correlation of variables j and k. Thus, the cross-variable CG pattern can reflect the low-factor structure of the SBF, rather than a separate behavioral bias about cross-variable misperceptions.

A. *Comparison with Univariate Benchmarks*

To give a benchmark for our joint approach, we evaluate how well common univariate models perform in matching these four sets of results (i.e., the wide range of forecast error autocorrelation, the wide range of CG coefficients, CG coefficients decreasing in maturity, and cross-variable CG coefficients increasing in objective correlation).

In the simplest univariate models, agents underreact/overreact for all variables by the same amount. This prediction underlies the idea that the underreaction/overreaction is a "deep" parameter, meaning that it can be estimated from expectations of one variable and assumed to apply to the agent's expectations of other variables. By definition, these types of models cannot explain the large range of forecast error autocorrelations and CG coefficients observed across variables, including the differences across maturities. This highlights that it is not trivial to take information from expectations of a few variables and accurately predict the behavior of expectations of other variables. Further, because these are univariate models, they imply cross-variable CG coefficients will be zero, as the bias for each variable only depends on its own revision.

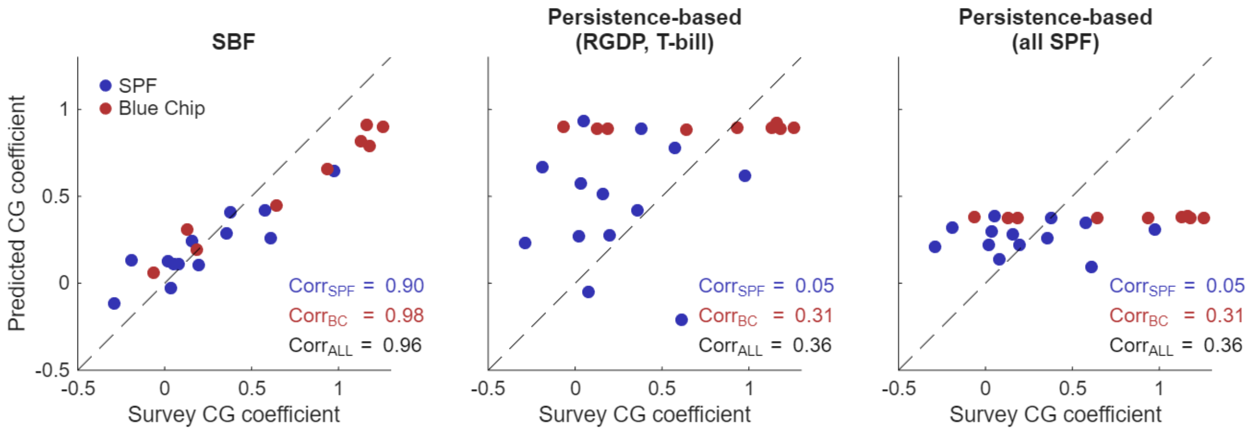

Figure 6. Comparison of the SBF and univariate persistence-based benchmarks. This figure illustrates how synthetic expectations and univariate persistence-based models match the Coibion-Gorodnichenko coefficients across the 22 Survey of Professional Forecasters and Blue Chip variables (excluding RGDP and T-bill). In all three panels, the horizontal axis is the actual survey CG coefficient for each variable and the vertical axis is the predicted CG coefficient. The left panel shows the prediction from the SBF (as in Figure 4). The middle and right panels show predictions from a persistence-based model, $CG_j = a + b\rho_j + \eta_j$, where ρ_j is the objective persistence of each variable j. In the middle panel, the model is estimated using only RGDP and T-bill, and in the right panel, the estimation uses all 15 Survey of Professional Forecasters variables. Blue points denote Survey of Professional Forecasters variables and red points denote Blue Chip variables. The sample period is 1982Q4-2022Q2.

A more interesting benchmark is one that directly attempts to explain differences in CG coefficients across variables. Both Afrouzi et al. (2023) and Bastianello and Imas (2026) propose models in which CG coefficients increase with the objective persistence of the variable. Figure 6 evaluates how well this benchmark matches empirical CG coefficients across the 22 variables. To facilitate comparisons, the horizontal axis for all three panels is the actual survey CG coefficients for the 22 variables. For reference, the left panel shows the predictions of the SBF from Figure 4.

The middle and right panels show the prediction based on the objective persistence of each variable. In line with Bastianello and Imas (2026), we estimate a linear relationship,

$$CG_j = a + b\rho_j + \eta_j$$

where ρ_j is the objective persistence of the variable. The predicted CG coefficient is then

$a + b\rho_j$. In the middle panel, we estimate a and b solely using expectations of RGDP and T-bill. Like the synthetic expectations, we want to test how well a model trained to match RGDP and T-bill performs in predicting the CG coefficients for other variables. Consistent with the model, RGDP has lower persistence and a lower CG coefficient than T-bill. However, the estimated a and b perform poorly in predicting the CG coefficients for other variables.

Perhaps training on two variables is simply not enough. In the right panel, a and b are estimated using all 15 Survey of Professional Forecasters variables. This compresses the range of predicted CG coefficients but does not noticeably improve the fit. Further, this rescaling does not change the correlation between the survey CG coefficients and the predicted CG coefficients.

Focusing just on the red points, which represent the Blue Chip expectations, we see the first issue. All of the yield variables have comparable persistences, near one. Because of this, the persistence-based model predicts all nine variables should have nearly identical CG coefficients. In reality, the CG coefficients vary widely across the different yield forecasts. Focusing on the Survey of Professional Forecasters variables, shown in blue, we see the second issue, which is that the variables differ substantially in their persistences but that this does not translate into clear differences in CG coefficients. Instead, the results for the Survey of Professional Forecasters variables are essentially a cloud with an unadjusted R^2 of 0.003.

These results highlight the difficulty in taking information from expectations of some variables and correctly predicting our four sets of forecast error predictability moments. Figure 6 shows the test of predicting the wide range of CG coefficients. Figure A6 shows similar results when trying to predict forecast error autocorrelation. As mentioned above, the persistence-based models struggle to match the d'Arienzo (2020) result on differences in CG coefficients across maturities, as all of the yields have similar persistences. Finally, while these models are more complex than the simplest underreaction/overreaction models, they are still univariate models, so they imply that only changes in expectations of $Z_{j,t+1}$ should predict forecast errors for variable j and that the cross-variable CG coefficients should be zero.

Importantly, we are not claiming that the persistence-based models lack merit nor that our two-factor SBF explains all patterns in forecast error predictability. Instead, our goal is to emphasize that a low-factor SBF provides a parsimonious way to match a number of untargeted moments of forecast error predictability, both qualitatively and quantitatively.

B. Implications for models of expectation formation

The previous results show that a two-factor SBF is fairly successful in matching both the time series of observed survey expectations, as well as a number of documented moments of biases in these expectations. What does the success of a low-factor SBF tell us about expectation formation?

In general, models of expectations partition variables into two groups. The first group is the set of variables for which the particular expectation-formation mechanism is applied directly. The second group is the set of variables whose expectations are inferred from beliefs about the first group of variables and the objective relationships between the variables. As a concrete example, in Bordalo et al. (2019), agents form diagnostic expectations about a latent component of earnings $f_{i,t}$ (Group 1), which then impact expectations of earnings and earnings growth (Group 2) through the objective relationship between $f_{i,t+1}$ and next period earnings. Similarly, in Nagel and Xu (2022), agents learn about consumption growth (Group 1) with fading memory, and these beliefs then affect expectations of future dividends and their future utility (Group 2) based on the objective relationships between Δc_{t+1}, Δd_{t+1}, and V_{t+1}.

A common approach is to assume that all variables are in Group 1. As an example, one might conjecture that people form diagnostic expectations for each individual variable with noisy information (Bordalo et al., 2020). Or, as in Afrouzi et al. (2023) and Bastianello and Imas (2026), expectations for each variable might be formed with a bias that depends on that variable's own persistence.

In contrast, the empirical success of a low-factor SBF is consistent with models in which only a small number of variables belong to Group 1, with the vast majority falling into Group 2. As discussed in Section III.C, the synthetic R^2 values are difficult to reconcile with univariate models of biases. Instead, the results are more consistent with a framework in which beliefs about one or two key variables spill over into expectations about many other outcomes through objective covariance relationships. As highlighted in Table IV, these variables do not need to correspond exactly to RGDP and T-bill themselves, but rather to broader underlying state variables correlated with them, such as TFP or the stance of monetary policy.

More broadly, our results highlight that many puzzles surrounding empirical biases in expectations become *easier* to understand once we think about expectations as being formed based on a small set of Group 1 variables rather than univariately. When the baseline model is that agents underreact/overreact univariately, the wide range of empirical CG coefficients and forecast error autocorrelations requires that the model must be modified so that agents are nearly rational in some settings (to generate near zero CG coefficients and autocorrelation for some variables) and highly biased in other settings.[23] Similarly, cross-variable CG coefficients must be explained by agents univariately under/overreacting to each variable while also making mistakes about cross-variable relationships (Wu, 2023). Our results demonstrate that these patterns (the wide range of CG coefficients, the variation in forecast error autocorrelation, and the cross-variable CG coefficients) arise naturally in a setting where agents have biased beliefs about a small set of $\hat{X}_{t+1}$ which then impact their expectations of all other variables.

It is important to note that our results are non-causal and do not rule out alternative explanations. Perhaps the true model is that agents underreact/overreact to each individual variable and also make mistakes about cross-variable relationships. Or perhaps agents form

[23] As an example, in Bastianello and Imas (2026), the agent relies on a default persistence $\bar{\rho}$ which makes her nearly rational when forecasting an outcome whose persistence is close to $\bar{\rho}$ and highly biased when forecasting an outcome with significantly higher or lower persistence.

expectations univariately using a default persistence $\bar{\rho}$ but were also impacted by another offsetting bias leading to the results of Figure 6. Perhaps there is a further bias which generates differences in forecast errors across maturities. We are simply highlighting that a low-factor SBF provides a parsimonious alternative that replicates these observed moments of biases reasonably well despite all of the results of Section IV being untargeted.

Overall, these results are good news for researchers seeking to incorporate subjective expectations into economic analyses, as they underscore the possibility 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 the same underlying SBF. Rather than requiring a separate expectation process for a large number of variables, it may be sufficient to model beliefs about a small number of underlying state variables. In Section V, we demonstrate how the SBF can be incorporated into economic models and show an application to a cross-sectional asset pricing model.

V. Application to Economic Models

A key feature of many economic models is that decisions and equilibrium outcomes depend on expectations of multiple variables. This can make it difficult to apply general mechanisms learned from lab experiments or surveys. For example, while experiments and surveys may present evidence of extrapolation, how do we apply this to a model where outcomes depend on expectations of output and inflation? Should the agent be modeled as extrapolating output, extrapolating inflation, extrapolating both, or extrapolating some combination of output and inflation? This only becomes more complicated as we consider economic models with more variables $X_{j,t+1}$.

To address this challenge, there are two standard approaches. The first is to assume a full model of expectation formation that specifies the agent's joint beliefs about all of the relevant variables and insert this into the economic model. The second is to collect survey data on expectations of all of the variables needed in the model and to insert this data into the model.

The SBF offers a third, middle-ground option. By imposing a mild coherence assumption, we can take survey data on a small number of variables and infer what the agent believes about correlated variables. This requires fewer assumptions than assuming a full model of expectation formation. At the same time, the assumption of coherence removes the need to have survey data on every relevant variable. Specifically, we can estimate $\hat{S}_{t+1}$ from a few variables $\hat{X}_{t+1}$ and evaluate equilibrium conditions under distorted beliefs $E_t\left[\hat{S}_{t+1}X_{t+1}\right]$. As discussed in Section I.C, these distorted beliefs $E_t\left[\hat{S}_{t+1}X_{t+1}\right]$ will capture all biases in expectations of X_{t+1} that are spanned by $\hat{X}_{t+1}$.

In this section, we apply the SBF approach to a classic model involving many expectations: a representative investor that simultaneously prices all assets and has differing preferences for payoffs in different states. By summarizing subjective expectations for all of the various assets with a single SBF, we can easily integrate subjective expectations into the equilibrium asset pricing equations, enabling us to distinguish the roles of beliefs and preferences. We apply this methodology specifically 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 176 anomalies compiled in Chen and Zimmermann (2022).[24] In each case, we provide evidence that the belief-based component explains a substantial amount of the excess returns of these anomalies.

A. Representative investor with multiple assets

Consider a model with a risk-free asset with return R_t^f and risky assets j with log-normal excess returns $R_{j,t+1}^e \equiv R_{j,t+1}/R_t^f$. Assets are priced by a representative investor who has conditionally log-normal preferences $\tilde{M}_{t+1}$. The equilibrium conditions in this model are

[24] Details on the construction of annual returns are shown in Appendix F.

simply

$$E_t^*\left[\tilde{M}_{t+1}\right]R_t^f = 1 \tag{27}$$

$$E_t^*\left[\tilde{M}_{t+1}R_{j,t+1}\right] = 1\ \forall j. \tag{28}$$

The goal of the model is to explain average excess returns on many assets.

Using our estimated conditionally log-normal SBF $\hat{S}_{t+1}$, the subjective expectations are $E_t^*\left[\tilde{M}_{t+1}\right] = E_t\left[\hat{S}_{t+1}\tilde{M}_{t+1}\right]$ and $E_t^*\left[\tilde{M}_{t+1}R_{j,t+1}\right] = E_t\left[\hat{S}_{t+1}\tilde{M}_{t+1}R_{j,t+1}\right]$. Given equations (27)-(28), Lemma 5 shows that average excess return can be easily decomposed into the belief-based component and the preference-based or risk-based component.

Lemma 5. *The average excess return for an asset* $E\left[R_{j,t+1}^e\right]$ *is*

$$\log\left(E\left[R_{j,t+1}^e\right]\right) = Cov\left(-\hat{s}_{t+1}, r_{j,t+1}^e\right) + Cov\left(-\tilde{m}_{t+1}, r_{j,t+1}^e\right). \tag{29}$$

Intuitively, Lemma 5 says that if the agent prices the risk-free bond and each asset, then a high average $R_{j,t+1}^e$ must be due to (i) the risky asset paying off in states the agent thinks are unlikely, $Cov\left(-\hat{s}_{t+1}, r_{j,t+1}^e\right)$, and/or (ii) the risky asset paying off in states the agent does not prefer, $Cov\left(-\tilde{m}_{t+1}, r_{j,t+1}^e\right)$.[25] Appendix A.6 provides a slightly more complicated relationship that holds in a more general model where we do not impose assumptions about the distributions of $\hat{S}_{t+1}$, $\tilde{M}_{t+1}$, and $R_{j,t+1}^e$.

Conveniently, equation (29) provides a clean separation of beliefs and preferences. Even if we cannot observe $\tilde{m}_{t+1}$, we can still estimate the unconditional covariance $Cov\left(-\hat{s}_{t+1}, r_{j,t+1}^e\right)$ using the observed $R_{j,t+1}^e$ and our estimated $\hat{S}_{t+1}$ and then compare it to $\log\left(E\left[R_{j,t+1}^e\right]\right)$ to assess whether the distorted beliefs captured by the estimated SBF can explain the average excess return. Importantly, we can do this without needing survey data for each individual $R_{j,t+1}$ or survey data on the expected product $\tilde{M}_{t+1}R_{j,t+1}$. This relaxes the need in previous papers (e.g., Engelberg, McLean, and Pontiff, 2020; Bordalo et al., 2024; Jensen, 2024; Décaire and Graham, 2026; Delao, Han, and Myers, 2026) to have survey data for each

[25] For S_{t+1}, $\tilde{M}_{t+1}$, and $R_{j,t+1}^e$, lower case letters denote log values.

individual asset j and to either impose a constant $\tilde{M}_{t+1}$ or proxy for $\tilde{M}_{t+1}$ using survey data on expected cash flows for all horizons.[26]

Further, decomposition (29) holds even if $\hat{s}_{t+1}$ 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 $Cov\left(-\hat{s}_{t+1}, r^e_{t+1}\right)$ will capture the magnitude of the agent's biased beliefs and $Cov\left(-\tilde{m}_{t+1}, r^e_{t+1}\right)$ 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 $Cov\left(-\hat{s}_{t+1}, r^e_{t+1}\right)$ will be much smaller than $Cov\left(-\tilde{m}_{t+1}, r^e_{t+1}\right)$. Even if the researcher does not observe $\tilde{m}_{t+1}$, the researcher will still correctly conclude that $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ is small compared to $\log\left(E\left[R^e_{j,t+1}\right]\right)$ and that $Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right)$ must account for most of the excess return.

To give an analogy for this decomposition, this exercise is similar to asset pricing research that estimates an SDF to match one set of assets and then measures whether that SDF accounts for the excess returns of some other set of test assets. In this case, we are estimating a distortion $\hat{s}_{t+1}$ 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. As in earlier exercises, the results are non-causal. It could be that the marginal investor's beliefs differ from those of the professional forecasters. However, we find that a framework in which the marginal investor shares the same beliefs as the professional forecasters goes a long way towards explaining a number of return anomalies.

By relating excess returns for many assets to a single distortion $\hat{s}_{t+1}$, our approach is reminiscent of research arguing that future returns are correlated with qualitative measures

[26] Because survey data on the product $\tilde{M}_{t+1}R_{j,t+1}$ is not available, these papers decompose average returns using an approximate Campbell-Shiller present-value identity $E\left[r_{j,t+1}\right] = E\left[E^*_t\left[r_{j,t+1}\right]\right] + E\left[\left(E^*_{t+1} - E^*_t\right)\left[\sum_{k=1}^{\infty} \Delta d_{j,t+k}\right]\right] - E\left[\left(E^*_{t+1} - E^*_t\right)\left[\sum_{k=2}^{\infty} r_{j,t+k}\right]\right]$ which requires survey data on expected returns and expected cash flows for all horizons.

of "sentiment."[27] While our approach is related to the 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. As discussed above, by using a quantitative measure of belief distortions rather than a qualitative measure of sentiment, equation (29) can still separate the belief component and the preference component even when $\hat{s}_{t+1}$ and $\tilde{m}_{t+1}$ are correlated.

Phrased differently, our test is restrictive in the sense advocated by Lewellen, Nagel, and Shanken (2010). They emphasize that asset pricing tests are most persuasive when the cross-sectional price of risk is pinned down by theory rather than estimated as a free parameter because, otherwise, candidate factors that are merely correlated with the true pricing factors will appear to explain returns at the wrong magnitude. In the decomposition in Lemma 5, the coefficient on $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ is unity by construction: $\hat{s}_{t+1}$ is calibrated entirely to subjective expectations of RGDP growth and the T-bill rate and is never rescaled to fit anomaly returns. This unit-coefficient discipline carries through every application below.

Continuing with the approach advocated in Lewellen, Nagel, and Shanken (2010), we assess how well the SBF performs for a large number of cross-sectional excess returns. Tests that focus only on a narrow menu of assets, such as the typical size-B/M portfolios, may not be sufficiently stringent because almost any factor weakly correlated with SMB and HML could capture these differences in mean returns. In the subsections below, we evaluate the SBF using Fama-French factors, the behavioral factors of Daniel, Hirshleifer, and Sun (2020), and a large set of 176 anomalies grouped into 22 conceptually distinct categories (profitability, leverage, momentum, issuance, accruals, etc.).

[27]See Baker and Wurgler (2006) and Kumar and Lee (2006) for examples in this literature.

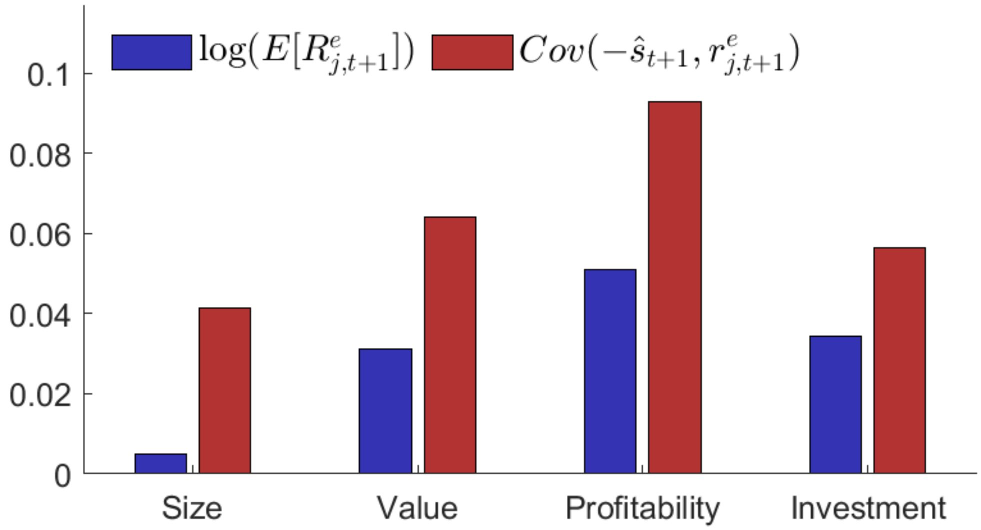

Figure 7. Fama-French factors and their covariance with the SBF. Using the log SBF $\hat{s}_{t+1}$, this figure shows the decomposition (29) of excess returns for each of the Fama-French cross-sectional factors. The blue bars show the average excess return $\log\left(E\left[R^e_{j,t+1}\right]\right)$ 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, $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$.

B. Fama-French factors

We first evaluate the four Fama-French anomalies.[28] 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 1982Q4 to 2022Q2.

Figure 7 shows the average excess returns $\log\left(E\left[R^e_{j,t+1}\right]\right)$ and the belief component denoted as $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ for each of 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 $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ gives reasonable values for annual anomaly returns of 4.2 pp to 9.3 pp. To give a comparison, the asset pricing literature has shown that SDFs that are estimated without using anomaly returns data often predict average excess returns that are orders of

[28] All four long-short return portfolios are taken from Ken French's data library and compounded into annual returns for the same sample as the distortion, 1982-2022. Appendix F contains details on the portfolio construction.

magnitude smaller or larger than what we observe in the data.[29]

For each anomaly, we find a positive $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ 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. In fact, for each factor, the belief component is larger than the average excess return, which indicates a negative risk component $Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right)$. For most of the factors, this has a fairly intuitive implication. This would translate to investors viewing growth firms as riskier than value firms, unprofitable firms as riskier than profitable firms, and aggressive firms as riskier than conservative firms. For the size factor, this translates to investors viewing large firms as riskier than small firms, which could be consistent with larger firms being harder to diversify, as their "granular" idiosyncratic shocks are large enough to impact the aggregate economy (Gabaix, 2011).

C. Behavioral factors

Since the above results suggest that some of the Fama-French factors can be explained with survey expectations, 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, constructed as the return spread between recent issuers and repurchasers, together with the PEAD factor, constructed as the return 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 long-horizon mispricing and one related to the four-day response to earnings surprises.

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

[29]See Cochrane and Hansen (1992) for a deeper discussion.

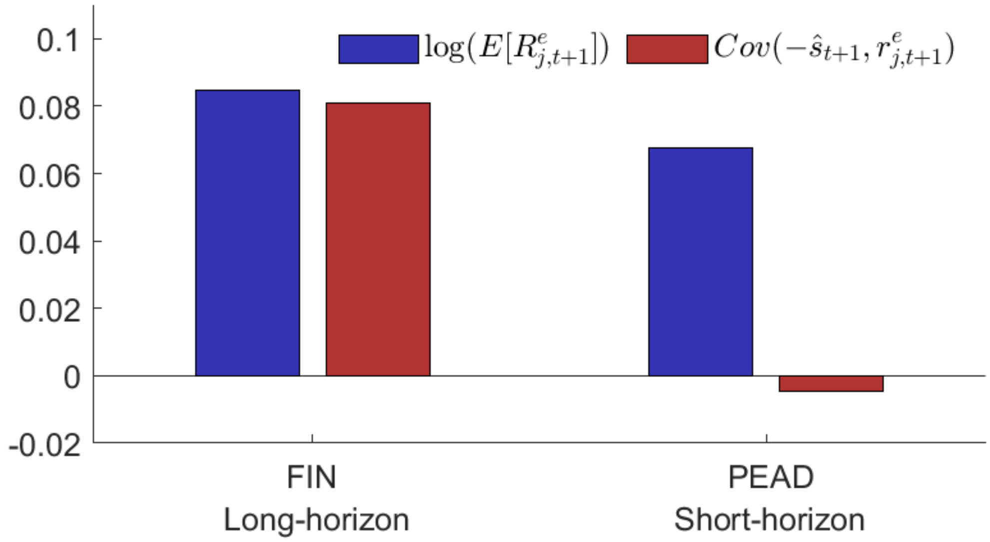

Figure 8. Behavioral factors and their covariance with the SBF. Using the log SBF $\hat{s}_{t+1}$, this figure shows the decomposition (29) of excess returns for each of the behavioral factors from Daniel, Hirshleifer, and Sun (2020). The blue bars show the average excess return $\log\left(E\left[R^e_{j,t+1}\right]\right)$ of the FIN (Financing) and PEAD (Post-Earnings Announcement Drift) factors. The red bars show the portion of the average return attributable to the covariance of the excess return with the log SBF, $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$.

short-horizon factor (PEAD) is intended to capture "high frequency" biases.

Figure 8 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 return for the financing factor, accounting for 96% of the average excess return generated by the factor. This supports a behavioral explanation for the financing factor and provides evidence against the idea that the excess return on FIN is mainly due to a rational risk premium. On the other hand, Figure 8 shows that our SBF has nearly no comovement with the short-horizon PEAD factor, with $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ being -0.48 pp compared to the observed 6.8 pp 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 the risk component $Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right)$.

D. *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 1982Q4 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 $\log\left(E\left[R^e_{j,t+1}\right]\right)$ and the average $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ across all anomalies j in each category.[30] Figure 9 shows the results.[31]

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

$$Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right) = \alpha_s + \gamma_s \log\left(E\left[R^e_{j,t+1}\right]\right) + \eta_{j,s} \tag{30}$$

$$Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right) = \alpha_m + \gamma_{\tilde{m}} \log\left(E\left[R^e_{j,t+1}\right]\right) + \eta_{j,\tilde{m}} \tag{31}$$

From equation (29), we know that

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

In other words, a one unit increase in $\log\left(E\left[R^e_{j,t+1}\right]\right)$ must correspond to a one unit increase in $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right) + Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right)$, and γ_s and $\gamma_{\tilde{m}}$ tell us whether it is primarily $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ that increases or primarily $Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right)$ that increases. Using our

[30] To 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 over our sample. The only exception is "Size," which only contains a single anomaly but is not combined with any other category given its historical importance.

[31] Note 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) provide multiple ways to define profitability, the results for categories like "profitability" and "investment" in Figure 9 can differ from the values for the individual anomalies in Figure 7. 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 $-\hat{s}_{t+1}$.

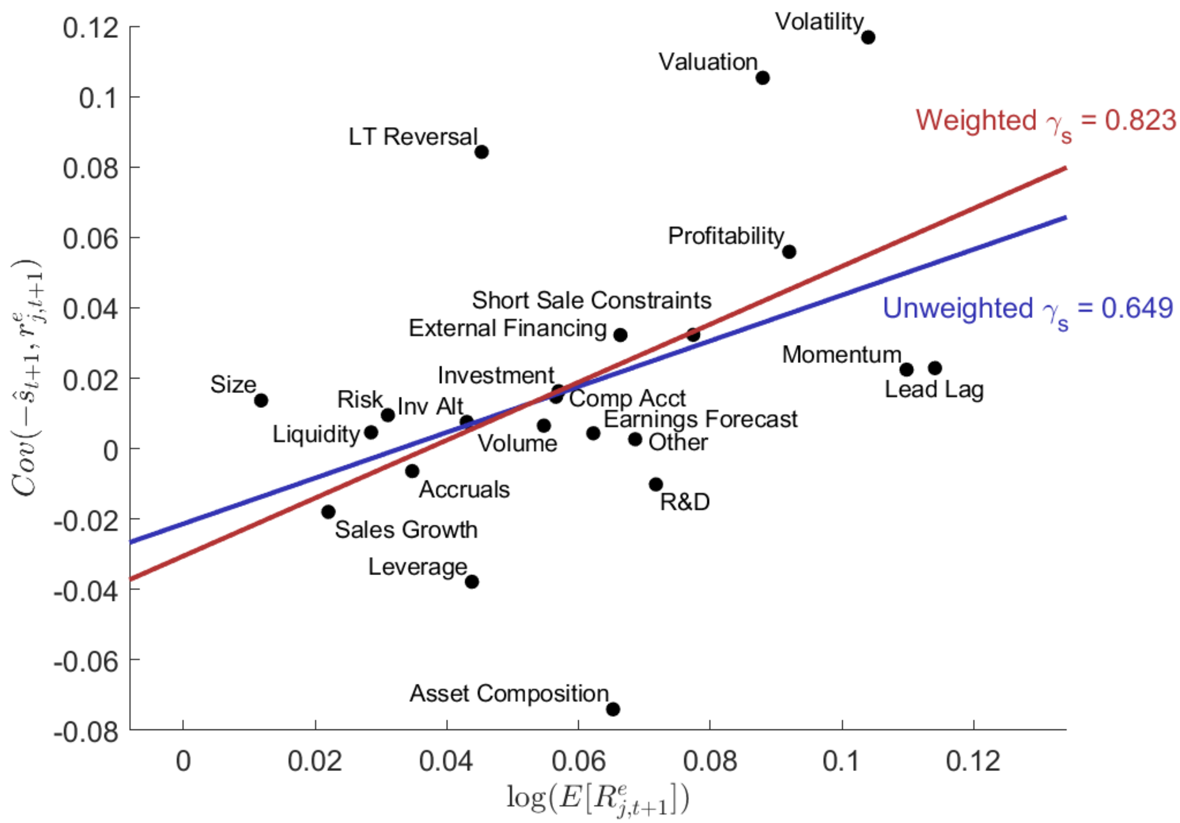

Figure 9. Anomalies and their covariance with the SBF. Using the log SBF $\hat{s}_{t+1}$, this figure shows the decomposition of excess return for the Chen and Zimmermann (2022) anomalies. The x-axis shows the average excess return $\log\left(E\left[R^e_{j,t+1}\right]\right)$ of the anomaly portfolios. The y-axis shows $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$, 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 $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ on $\log\left(E\left[R^e_{j,t+1}\right]\right)$ as in equation (30). The line in red shows the slope of the same regression weighted by the number of anomalies in each category.

estimated $-\hat{s}_{t+1}$, Figure 9 shows that regression (30) gives a coefficient γ_s equal to 0.649 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 0.823. Both results point to the belief component and the risk component playing a non-trivial role, with either a 60/40 split or an 80/20 split between the two.

While γ_s does quantify the relationship between $\log\left(E\left[R^e_{j,t+1}\right]\right)$ and $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$, it is important to note that γ_s potentially incorporates interaction between $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ and $Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right)$ since $\log\left(E\left[R^e_{j,t+1}\right]\right)$ is the sum of both components. From equations (29) and (30), the coefficient γ_s can be decomposed as

$$\gamma_s = \frac{Var^{CX}\left[Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)\right]}{Var^{CX}\left[\log\left(E\left[R^e_{j,t+1}\right]\right)\right]} + \frac{Cov^{CX}\left[Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right), Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right)\right]}{Var^{CX}\left[\log\left(E\left[R^e_{j,t+1}\right]\right)\right]} \quad (32)$$

where $Var^{CX}[\cdot]$ and $Cov^{CX}[\cdot,\cdot]$ are the cross-sectional variance and covariance across assets j. The first term captures a direct effect measuring the pure cross-sectional variation in the belief component, and the second term captures the indirect effect, measuring whether risk and distorted beliefs are correlated in the cross-section, i.e., whether anomalies with high SBF exposure also tend to have high $\tilde{m}_{t+1}$ exposure. In the equally-weighted case, the direct effect is 2.32 and the indirect effect is -1.67; in the weighted case, these are 2.99 and -2.16. The indirect effect is negative in both cases, indicating that anomalies concentrated in states forecasters *underweight* tend also to be concentrated in states investors *prefer*. As a result, γ_s actually understates the direct contribution of the SBF: the belief component alone more than accounts for cross-sectional dispersion in excess returns, while its interaction with preferences offsets some of this variation. This pattern mirrors the Fama-French results, where the belief component exceeded each factor's full average return, implying that there would need to be an offsetting $Cov\left(-\tilde{m}_{t+1}, r_{j,t+1}\right)$. Rephrased, the decomposition in equation (32) rules out the possibility that γ_s is high because of positive comovement between the belief and preference components.

These results show that our single SBF $\hat{s}_{t+1}$ based on beliefs for just two macroeconomic variables accounts for a quantitatively large share of cross-sectional variation in excess returns, but there is still a similarly large amount of variation remaining that can be attributed to risk or preferences.

E. Interpretation under misspecification

Equation (29) shows how to incorporate the estimated log SBF $\hat{s}_{t+1}$ into the representative investor model. A relevant question is how potential misspecification would enter into this equation. If the true log SBF for the model agent is s_{t+1} which differs from $\hat{s}_{t+1}$, then the decomposition becomes

$$\begin{aligned}\log\left(E\left[R^e_{j,t+1}\right]\right) &= Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right) + Cov\left(-s_{t+1}+\hat{s}_{t+1}, r^e_{j,t+1}\right) \\ &\quad + Cov\left(-\tilde{m}_{t+1}, r^e_{j,t+1}\right). \end{aligned} \tag{33}$$

For our empirical exercises, we directly measure $\log\left(E\left[R^e_{j,t+1}\right]\right)$ and $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$. Because of this, our results only reflect how well biased beliefs about one-year RGDP and T-bill shocks (i.e., the biased beliefs incorporated in $\hat{s}_{t+1}$) explain various returns. Any gap between $\log\left(E\left[R^e_{j,t+1}\right]\right)$ and $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ could be due to either preferences or biased beliefs about shocks not spanned by $\hat{\varepsilon}_{t+1}$. For example, in Section V.C, we discuss that the failure of our estimated $\hat{s}_{t+1}$ to explain the PEAD return is unsurprising given that PEAD is meant to capture short-horizon behavioral biases whereas our $\hat{s}_{t+1}$ only captures biases in long-horizon (i.e., one-year) expectations. Additional biases in expectations, $s_{t+1} - \hat{s}_{t+1}$, that negatively comove with realized returns will lessen the need for preferences to explain returns, and conversely additional biases that positively comove with realized returns will increase the need for preferences to explain returns. Thus, as with the earlier discussion about spanning, these estimates do not speak to all possible biased beliefs, but rather only show how well biased beliefs about shocks spanned by $\hat{\varepsilon}_{t+1}$ can account for the observed phenomena.

F. *Returns using alternative SBFs*

The results in Sections V.B-V.D regarding the role of the belief component are based on an SBF estimated from two variables in $\hat{X}_{t+1}$, RGDP growth and the T-bill rate. To assess robustness to this choice, we replicate the exercises in Sections V.B–V.D using SBFs estimated from the ten alternative pairs of variables shown in Table IV. For each pair, the corresponding figures are reported in Appendix I.

Given the emphasis on spanning in the previous results, a natural question is why shouldn't one simply use a 24-factor SBF that captures all 24 available subjective expectations? Using the SBF to explain returns for which we lack survey data is effectively an out-of-sample exercise. Because of this, it is important to avoid overfitting in sample. As a simple example, the loadings $\hat{\beta}_t = \hat{\Sigma}^{-1}\left(E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]\right)$ depend on a stable estimate for $\hat{\Sigma}^{-1}$ which has 100 times as many entries for a 24-factor $\hat{X}_{t+1}$ compared to a two-factor $\hat{X}_{t+1}$. Thus, rather than using a 24-factor SBF, we focus on two-factor SBFs that perform well in summarizing the combined 24 subjective expectations, similar in spirit to Stock and Watson (2002).

Appendix Figures A7-A9 show the contribution of the belief component to the Fama-French factors, behavioral factors, and Chen-Zimmermann anomalies across the ten alternative pairs. The main conclusions of Sections V.B-V.C are preserved across all pairs: the belief component remains positive and quantitatively large for the Fama-French factors (Figure A7), the SBF continues to account for the entirety of the FIN factor while leaving the short-horizon PEAD factor largely unexplained (Figure A8).

Appendix Figure A9 shows the Chen-Zimmermann scatter plots across the ten alternative pairs. The results are also broadly consistent with the baseline. Nine of the ten pairs produce positive γ_s coefficients. One pair produces a slightly negative slope driven by the momentum anomaly, which has a relatively high average excess return but a negative covariance with that specific $-\hat{s}_{t+1}$. Taken together, the average unweighted γ_s across the ten pairs is 0.37 and the average weighted γ_s is 0.47, suggesting a robust relationship between excess returns and $Cov\left(-\hat{s}_{t+1}, r_{j,t+1}^e\right)$.

VI. 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. We demonstrate that an SBF based on subjective expectations for RGDP growth and the T-bill rate goes a substantial way toward 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 and Hansen-Jagannathan bounds, can potentially provide important insights for subjective expectations. Conversely, models of expectation formation, tests for biases, 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, *Handbook of Economic Expectations* 477–506.

Afrouzi, Hassan, Spencer Kwon, Augustin Landier, Yueran Ma, and David Thesmar, 2023, Overreaction in expectations: Evidence and theory, *The Quarterly Journal of Economics* 138, 1713–1764.

Bachmann, Rüdiger, Giorgio Topa, and Wilbert van der Klaauw, 2023, *Handbook of Economic Expectations* (Academic Press).

Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *The Journal of Finance* 61, 1645–1680.

Bastianello, Francesca, and Alex Imas, 2026, Biases in belief updating within and across domains, *Working Paper,* University of Chicago.

Bauer, Michael D, Carolin E Pflueger, and Adi Sunderam, 2024a, Changing perceptions and post-pandemic monetary policy, *Proceedings of 2024 Jackson Hole Economic Symposium: Reassessing the Effectiveness and Transmission of Monetary Policy* 285–336.

Bauer, Michael D, Carolin E Pflueger, and Adi Sunderam, 2024b, Perceptions about monetary policy, *The Quarterly Journal of Economics* 139, 2227–2278.

Bhandari, Anmol, Jaroslav Borovička, and Paul Ho, 2025, Survey data and subjective beliefs in business cycle models, *The Review of Economic Studies* 92, 1375–1437.

Bianchi, Francesco, Sydney Ludvigson, and Sai Ma, 2026, 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, 2019, Diagnostic expectations and stock returns, *Journal of Finance* 74, 2839–2874.

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

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2025, Finance without exotic risk, *Journal of Financial Economics* 173, 104145.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2018, Diagnostic expectations and credit cycles, *Journal of Finance* 73, 199–227.

Borovička, Jaroslav, Lars Peter Hansen, and José A Scheinkman, 2016, Misspecified recovery, *The Journal of Finance* 71, 2493–2544.

Brav, Alon, and J B Heaton, 2002, Competing theories of financial anomalies, *The Review of Financial Studies* 15, 575–606.

Bundick, Brent, 2015, Estimating the monetary policy rule perceived by forecasters, *Federal Reserve Bank of Kansas City, Economic Review* 4, 33–49.

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.

Chen, Xiaohong, Lars Peter Hansen, and Peter G. Hansen, 2024, Robust inference for moment condition models without rational expectations, *Journal of Econometrics* 243.

Chernov, Mikhail, and Philippe Mueller, 2012, The term structure of inflation expectations, *Journal of Financial Economics* 106, 367–394.

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 (Princeton University Press).

Cochrane, John H., and Lars P. Hansen, 1992, Asset pricing explorations for macroeconomics, *NBER Macroeconomics Annual* 7.

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.

Couts, Spencer J, Andrei S Gonçalves, and Johnathan A Loudis, 2025, The subjective risk and return expectations of institutional investors, *Working Paper,* Fisher College of Business.

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

Daniel, Kent, David Hirshleifer, and Lin Sun, 2020, Short- and long-horizon behavioral factors, *Review of Financial Studies* 33, 1673–1736.

d'Arienzo, Daniele, 2020, Maturity increasing overreaction and bond market puzzles, *Working Paper,* Nova School of Business and Economics.

Delao, Ricardo, Xiao Han, and Sean Myers, 2026, The cross-section of subjective expectations: Understanding prices and anomalies, *Working Paper,* SSRN.

Décaire, Paul H, and John R Graham, 2026, Valuation fundamentals, *Working Paper,* Arizona State University.

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.

Farmer, Leland E, Emi Nakamura, and Jón Steinsson, 2024, Learning about the long run, *Journal of Political Economy* 132, 3334–3377.

Gabaix, Xavier, 2011, The granular origins of aggregate fluctuations, *Econometrica* 79, 733–772.

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

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

Hansen, Lars Peter, and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.

Hansen, Lars Peter, and Scott F Richard, 1987, The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587–613.

Hansen, Lars Peter, and Thomas J Sargent, 2001, Robust control and model uncertainty, *American Economic Review* 91, 60–66.

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, 2024, Subjective risk and return, *Working Paper,* Yale University.

Jia, Pengfei, Haopeng Shen, and Shikun Zheng, 2023, Monetary policy rules and opinionated markets, *Economics Letters* 223, 110995.

Kim, Jinill, and Seth Pruitt, 2017, Estimating monetary policy rules when nominal interest rates are stuck at zero, *Journal of Money, Credit and Banking* 49, 585–602.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2018, Interpreting factor models, *The Journal of Finance* 73, 1183–1223.

Kremens, Lukas, Ian W R Martin, and Liliana Varela, 2025, Long-horizon exchange rate expectations, *The Journal of Finance* 80, 3695–3724.

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.

Lewellen, Jonathan, Stefan Nagel, and Jay Shanken, 2010, A skeptical appraisal of asset pricing tests, *Journal of Financial Economics* 96, 175–194.

Li, Zigang, Stijn Van Nieuwerburgh, and Wang Renxuan, 2025, Understanding rationality and disagreement in house price expectations, *The Review of Financial Studies* 39, 297–342.

Link, Sebastian, Andreas Peichl, Christopher Roth, and Johannes Wohlfart, 2023, Information frictions among firms and households, *Journal of Monetary Economics* 135, 99–115.

Maenhout, Pascal J, Andrea Vedolin, and Hao Xing, 2025, Robustness and dynamic sentiment, *Journal of Financial Economics* 163, 103953.

Malmendier, Ulrike, and Stefan Nagel, 2016, Learning from inflation experiences, *Quarterly Journal of Economics* 131, 53–87.

Maxted, Peter, 2024, A macro-finance model with sentiment, *The Review of Economic Studies* 91, 438–475.

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.

Nagel, Stefan, and Zhengyang Xu, 2022, Asset pricing with fading memory, *Review of Financial Studies* 35, 2190–2245.

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

Shiller, Robert J, 1979, The volatility of long-term interest rates and expectations models of the term structure, *Journal of Political Economy* 87, 1190–1219.

Shue, Kelly, Richard Townsend, and Chen Wang, 2024, Categorical thinking about interest rates, *Working Paper,* Yale University.

Stambaugh, Robert F, Jianfeng Yu, and Yu Yuan, 2012, The short of it: Investor sentiment and anomalies, *Journal of Financial Economics* 104, 288–302.

Stock, James H, and Mark W Watson, 2002, Macroeconomic forecasting using diffusion indexes, *Journal of Business and Economic Statistics* 20, 147–162.

Wu, Lingxuan, 2023, Mental macro-finance models: Evidence and theory, *Working Paper,* New York University.

Appendix

A. *Additional analytical results*

A.1. Estimating a range for synthetic expectations

Consider the environment of Section III.B in which $Y_{t+1} = E_t[Y_{t+1}] + \varepsilon_{y,t+1}$, $\varepsilon_{y,t+1}$ is a multivariate Gaussian shock, and $\hat{S}_{t+1}$ is defined by equation (11). We focus on the environment of Section III.B to highlight that this process can be done without using any survey data on expectations of Y_{t+1}. All of these results also apply to the environment of Section III.A. Then the unspanned component for Y_{t+1} is

$$u_{t+1} = \varepsilon_{y,t+1} - Cov\left(\varepsilon_{y,t+1}, \hat{\varepsilon}_{t+1}\right)' \hat{\Sigma}^{-1} \hat{\varepsilon}_{t+1}.$$

Let $\sigma_{u,j}$ be the standard deviation of $u_{j,t+1}$ for each variable $Y_{j,t+1}$. From Lemma 4, we have that

$$\begin{aligned} E_t^*[Y_{j,t+1}] &= E_t\left[\hat{S}_{t+1} Y_{j,t+1}\right] + Cov_t\left(u_{j,t+1}, u_{t+1}^S\right) \\ &= E_t\left[\hat{S}_{t+1} Y_{j,t+1}\right] + \sigma_{u,j} \rho_{j,t} \sigma_{u^S,t} \end{aligned} \tag{A1}$$

where $\rho_{j,t}$ is the conditional correlation between u_{t+1}^S and $u_{j,t+1}$ and $\sigma_{u^S,t}$ is the conditional standard deviation of u_{t+1}^S. This means we have the following bound

$$E_t^*[Y_{j,t+1}] \in \left[E_t\left[\hat{S}_{t+1} Y_{j,t+1}\right] - \sigma_{u,j}\sigma_{u^S,t}, E_t\left[\hat{S}_{t+1} Y_{j,t+1}\right] + \sigma_{u,j}\sigma_{u^S,t}\right]$$

Note that the width of this range is proportional to $\sigma_{u,j}$, meaning that the range is tighter for variables with less unspanned variation.

Importantly, we do not need any survey data on Y_{t+1} to measure $\sigma_{u,j}$ or $E_t\left[\hat{S}_{t+1} Y_{j,t+1}\right]$. Thus, given an estimate for $\sigma_{u^S,t}$, which we will denote as $\hat{\sigma}_{u^S,t}$, we can then estimate a synthetic range $\left[E_t\left[\hat{S}_{t+1} Y_{j,t+1}\right] - \sigma_{u,j}\hat{\sigma}_{u^S,t}, E_t\left[\hat{S}_{t+1} Y_{j,t+1}\right] + \sigma_{u,j}\hat{\sigma}_{u^S,t}\right]$ for the subjective expectation of $Y_{j,t+1}$. The width of this synthetic range will also be proportional to $\sigma_{u,j}$.

A.2. Incorporating subjective variances and covariances

Given a multivariate X_{t+1}, suppose we have data on the subjective expected mean $E_t^*[X_{t+1}]$ and the subjective expected covariance matrix Σ_t^*. Our goal is to find a variable S_{t+1} such that

$$E_t^*[X_{t+1}] = E_t[S_{t+1}X_{t+1}] \tag{A2}$$

$$\Sigma_t^* = E_t\left[S_{t+1}X_{t+1}X_{t+1}'\right] - E_t[S_{t+1}X_{t+1}]\,E_t[S_{t+1}X_{t+1}]'. \tag{A3}$$

Note that if we are only trying to match subjective expectations about a piece of the covariance matrix, then this task is even easier as there are fewer moments that S_{t+1} needs to satisfy. For example, suppose we only have data on the subjective expected mean and the subjective expected variance (i.e., the diagonal of Σ_t^*) and we want to find an S_{t+1} that matches those data. Given an objective conditional correlation matrix C_t, we can simply create a hypothetical subjective covariance matrix $D_tC_tD_t$ where D_t is a diagonal matrix whose elements are the subjective expected standard deviations. We can then apply the methodology below to this hypothetical subjective covariance matrix, which will guarantee that we match the data on the subjective expected variance.

Similar to Proposition 2, if the variables are objectively normally distributed, then we have a straightforward representation for S_{t+1}.

Proposition 6. *For a multivariate X_{t+1} and subjective expectations $E_t^*[X_{t+1}]$ and Σ_t^*, if the objective conditional distribution is $X_{t+1} \sim N(E_t[X_{t+1}], \Sigma_t)$, then equations (A2) and (A3) are satisfied for*

$$\begin{aligned}
s_{t+1} &= -\frac{1}{2}\beta_t'\Sigma_t^*\beta_t + \beta_t'\varepsilon_{t+1} + \frac{1}{2}\left[\log\left(\frac{\det(\Sigma_t)}{\det(\Sigma_t^*)}\right) + \varepsilon_{t+1}'\left(\Sigma_t^{-1} - \Sigma_t^{*-1}\right)\varepsilon_{t+1}\right] \\
\beta_t &= \Sigma_t^{*-1}\left(E_t^*[X_{t+1}] - E_t[X_{t+1}]\right) \\
\varepsilon_{t+1} &= X_{t+1} - E_t[X_{t+1}].
\end{aligned}$$

Compared to Proposition 2, the log SBF s_{t+1} now depends not just on the shocks ε_{t+1} but also on the squared shocks $\varepsilon_{t+1}'\left(\Sigma_t^{-1} - \Sigma_t^{*-1}\right)\varepsilon_{t+1}$ to account for higher-order subjective expectations.

A.3. The SBF in various expectation formation models

Consider the environment of Section I.B.1 for a simple case of AR(1) real GDP growth,

$$g_{t+1} = \rho g_t + \varepsilon_{t+1}, \qquad \rho > 0, \qquad \varepsilon_{t+1} \sim N\left(0, \sigma^2\right). \tag{A4}$$

Section I.B shows that the subjective expectations of g_{t+1} can be represented by the log SBF $s_{t+1} = -\frac{1}{2}\beta_t^2\sigma^2 + \beta_t\varepsilon_{t+1}$. Different models of expectation formation imply different values of β_t. In that section, we show the case of an extrapolative agent who believes the persistence is ρ^* rather than ρ, which implies the loading

$$\beta_t^{EX} = \frac{\rho^* - \rho}{\sigma^2} g_t. \tag{A5}$$

We now extend this exercise to three other expectation formation models: diagnostic expectations, Bayesian learning about the mean, and robust control.

Diagnostic Expectations Suppose the agent has diagnostic expectations of future growth, $E_t^*\left[g_{t+1}\right] = E_t\left[g_{t+1}\right] + \theta\left(E_t\left[g_{t+1}\right] - E_{t-1}\left[g_{t+1}\right]\right)$. This simplifies to $\rho g_t + \rho\theta\varepsilon_t$, where $\theta > 0$ reflects how much the agent overreacts to the most recent shock ε_t. In this case, the loading is

$$\beta_t^{DE} = \frac{\theta\rho}{\sigma^2}\varepsilon_t. \tag{A6}$$

The loading no longer depends on current growth. It only depends on the most recent shock. When the recent shock ε_t is positive, it is as if the agent exaggerates the probability of positive future ε_{t+1} and understates the probability of negative ε_{t+1}.

Bayesian Learning Suppose we have a rational Bayesian agent who is learning about the true mean of the process. Learning starts at $t = 0$ with a prior belief of μ_0^* and initial uncertainty h_0. The agent updates her estimate μ_t^* every period, applying a

Kalman gain $K_t = \frac{h_0}{\sigma^2 + h_0 t}$ to every innovation.[32] The agent expects next period growth to be $\mu_t^* + \rho(g_t - \mu_t^*)$. This gives a loading of

$$\beta_t^{BL} = \frac{1}{\sigma^2}\left[(1 - tK_t)(1-\rho)\mu_0^* + tK_t\left(\frac{1}{t}\sum_{j=0}^{t-1}\varepsilon_{t-j}\right)\right]. \tag{A7}$$

The loading depends on a weighted average of the initial bias $(1-\rho)\mu_0^*$ and the average of all observed shocks. Over time, the weight placed on the initial bias deterministically shrinks towards zero.

Robust Control Suppose we have a rational agent who is uncertain about the true underlying model à la Hansen and Sargent (2001). As an example, we will focus on the robust control specification of Bhandari, Borovička, and Ho (2025). Given a parameter θ which determines the amount of concern about model misspecification and a value function V_t, the agent's expectation of next period growth is $\rho g_t - \theta\sigma_{t,v,g}$ where $\sigma_{t,v,g}$ is the conditional covariance of V_{t+1} and g_{t+1}. Rephrased, the agent adjusts her expectations depending on how much her future value function depends on g_{t+1}. This gives a loading of

$$\beta_t^{RC} = -\frac{\theta}{\sigma^2}\sigma_{t,v,g}. \tag{A8}$$

When V_{t+1} is positively related to g_{t+1} and $\theta > 0$, the loading is negative which means that the agent is exaggerating the probability of negative future ε_{t+1} and understating the probability of positive future ε_{t+1}. Note that the second term in the SBF (i.e., $\beta_t^{RC}\varepsilon_{t+1}$) is just $-\theta$ multiplied by the projection of V_{t+1} onto g_{t+1}, i.e., $\frac{\sigma_{t,v,g}}{\sigma^2}\varepsilon_{t+1}$.

A.4. Breaking coherence in expectation formation models

Appendix A.3 shows examples of the SBF in various models of expectation formation. This raises a natural question of which types of models do not have an SBF representation. We know from Proposition 1 that coherence is all that is needed for the SBF to exist, so we can reframe this question as asking which types of models violate coherence.

[32] The solution to the agent's learning problem is identical to learning the mean of an iid process $z_t = g_t - \rho g_{t-1}$.

Conceptually, satisfying or breaking coherence is less about the specific belief formation mechanism (e.g., diagnostic expectations, extrapolation, sticky expectations) and more about the set of variables to which the mechanism is applied. As discussed in Section IV.B, models of expectations partition variables into two groups. Group 1 is the set of variables to which we apply the specific belief formation mechanism, such as extrapolation. Group 2 is the set of variables for which the agent forms her expectations based on her beliefs about the Group 1 variables and the objective relationships between variables. For example, the agent might extrapolate inflation (Group 1) and then form her expectations about future wages based on the objective relationship between wages and inflation (Group 2). Coherence can be broken when Group 1 contains variables with degenerate shocks.

To provide a concrete example of the same belief formation mechanism breaking coherence and satisfying coherence, we can consider two models of diagnostic expectations: Bordalo et al. (2018) and Bordalo et al. (2019). Since both papers use the same notation $E_t^{\theta}[\cdot]$ to denote diagnostic expectations, we will introduce a more detailed notation $E_t^{\theta,Z_{t+1}}[\cdot]$ to denote diagnostic expectations when the agent is applying the diagnostic mechanism to beliefs about Z_{t+1}. This will help to clarify the difference between the two models.

In Bordalo et al. (2018), the agent is assumed to apply the diagnostic mechanism separately to each individual variable. In other words, all variables are Group 1 and the paper is using $E_t^{\theta}[Z_{t+1}] \equiv E_t^{\theta,Z_{t+1}}[Z_{t+1}]$ for all Z_{t+1}. As detailed in their internet appendix, this means that $E_t^{\theta}[X_{t+1} - X_t] \neq E_t^{\theta}[X_{t+1}] - E_t^{\theta}[X_t] = E_t^{\theta}[X_{t+1}] - X_t$. Thus, coherence is violated in this model of diagnostic expectations.

In Bordalo et al. (2019), the agent is assumed to apply the diagnostic mechanism to expectations of a latent variable f_t and forms her beliefs about all other variables based on her beliefs about f_t and how the variables objectively relate to f_t. In other words, f_t is a Group 1 variable and all other variables are Group 2, so the paper is using $E_t^{\theta}[Z_{t+1}] \equiv E_t^{\theta,f_t}[Z_{t+1}]$ for all Z_{t+1}. In this model, coherence is satisfied and the paper repeatedly uses the fact that $E_t^{\theta}[X_{t+1} - X_t] = E_t^{\theta}[X_{t+1}] - X_t$ in its derivations.

How can it be that $E_t^\theta[X_{t+1} - X_t] \neq E_t^\theta[X_{t+1}] - X_t$ in one model of diagnostic expectations and $E_t^\theta[X_{t+1} - X_t] = E_t^\theta[X_{t+1}] - X_t$ in another? It is because the more precise statement is $E_t^{\theta, X_{t+1}-X_t}[X_{t+1} - X_t] \neq E_t^{\theta, X_{t+1}}[X_{t+1}] - X_t$ but $E_t^{\theta, f_t}[X_{t+1} - X_t] = E_t^{\theta, f_t}[X_{t+1}] - X_t$. If the agent is diagnostic about both the *change* $X_{t+1} - X_t$ and the *level* X_{t+1}, then coherence is violated. This is because the change and the level depend on the same objective shock, $\varepsilon_{t+1} \equiv X_{t+1} - E_t[X_{t+1}] = \Delta X_{t+1} - E_t[\Delta X_{t+1}]$. Applying the diagnostic mechanism to beliefs about the change and the level implies conflicting beliefs about the future shock ε_{t+1}. In that case, there is no single distortion that can explain both beliefs.

This example illustrates the general point that coherence can be violated when a behavioral mechanism is applied separately to variables whose shocks are degenerate (i.e., variables that are linked to the same underlying shock). Because of this, our results are more in line with models of diagnostic expectations like Bordalo et al. (2019), in which the agent's beliefs across different variables are characterized by a single distortion (i.e., the agent is always diagnostic about f_t), rather than models of diagnostic expectations like Bordalo et al. (2018) in which beliefs about each variable utilize a different distortion. This also highlights why our results fit nicely with robust control, as the mechanism of robust control always distorts probabilities for a single variable, specifically utility,[33] regardless of how it is applied. Following Hansen and Sargent (2001), both the multiplier robust-control problem and the constraint robust-control problem use a probability distortion that exaggerates the states with the lowest utility and the agent's beliefs about inflation, unemployment, or any other variables are distorted based on how they comove with utility. Rephrased, the agent exaggerates the probability of the "worst-case" state and the definition of the "worst-case" state is always the state with the lowest utility.

[33] Bhandari, Borovička, and Ho (2025) provide a particularly straightforward setting, where the distortion simplifies to $\frac{\exp(\theta_t V_{t+1})}{E_t[\exp(\theta_t V_{t+1})]}$, with V_{t+1} denoting the utility, and θ_t denoting the degree of belief distortion.

A.5. Upper bound from shocks correlated across variables and across time

In Section III.C.2, we provide an upper bound of the results in a model where each variable's bias is driven solely by its own shocks and the shocks can be correlated across variables but are not serially autocorrelated. In this section, we show sufficient conditions to achieve the same bound as Proposition 5 in the presence of serial autocorrelation.

We consider a model in which the bias for each variable can depend on its entire history of shocks:

$$E_t^*\left[Z_{j,t+1}\right] - E_t\left[Z_{j,t+1}\right] \;=\; \sum_{\tau=0}^{\infty} \delta_{j,\tau}\varepsilon_{j,t-\tau}. \tag{A9}$$

where the shocks $\varepsilon_{j,t}$ can be contemporaneously correlated across variables. The lag coefficients $\{\delta_{j,\tau}\}$ can take any *non-negative* number. The shocks can be serially correlated with their lag-h covariances $\Omega_{jk}(h) \equiv Cov(\varepsilon_{j,t}, \varepsilon_{k,t-h})$ satisfying the following conditions:

$$\begin{aligned} \Upsilon'\hat{\Sigma}^{-1/2}\Omega_{jj}(h)^{-1}\ \Omega_{j\hat{s}}(h) \quad &\leq \quad \Upsilon'\hat{\Sigma}^{-1/2}\Omega_{jj}(0)^{-1}\ \Omega_{j\hat{s}}(h) \\ diag(\Upsilon)\hat{\Sigma}^{-1/2}\Omega_{j\hat{s}}(h) \quad &\leq \quad diag(\Upsilon)\hat{\Sigma}^{-1/2}\Omega_{\hat{s}\hat{s}}(h)\hat{\Sigma}^{-1}\Omega_{j\hat{s}}(0), \text{ for } h \geq 1, \end{aligned} \tag{A10}$$

where $\Omega_{j\hat{s}}(h) = Cov(\varepsilon_{j,t}, \hat{\varepsilon}_{t-h})$, $\Omega_{\hat{s}\hat{s}}(h) = Cov(\hat{\varepsilon}_t, \hat{\varepsilon}_{t-h})$, and $\Upsilon = \hat{\Sigma}^{-1/2}\Omega_{j\hat{s}}(0)$.

Intuitively, both conditions ensure that the serial cross-covariances at extended lags $\Omega_{j\hat{s}}(h)$ do not grow too much relative to the serial autocovariance of the shocks $\varepsilon_{j,t}$ and $\hat{\varepsilon}_t$ (i.e., $\Omega_{jj}(h)$ and $\Omega_{\hat{s}\hat{s}}(h)$) as the lags extend. This is easy to see in the special case where Υ, the contemporaneous projection of $\varepsilon_{j,t}$ on the shocks $\hat{\varepsilon}_t$, and $\hat{\Sigma}^{-1/2}$ have all non-negative elements. In this case, the conditions simplify to the intuitive relations:

$$\begin{aligned} \Omega_{jj}(h)^{-1}\ \Omega_{j\hat{s}}(h) \quad &\leq \quad \Omega_{jj}(0)^{-1}\Omega_{j\hat{s}}(0) \\ \Omega_{\hat{s}\hat{s}}(h)^{-1}\Omega_{j\hat{s}}(h) \quad &\leq \quad \hat{\Sigma}^{-1}\Omega_{j\hat{s}}(0), \qquad \text{for } h \geq 1. \end{aligned}$$

Proposition 7. *Let biases follow equation (A9) and $\hat{\varepsilon}_t$ denote the subset of shocks used to construct the synthetic bias satisfying equation (A10). Then for each variable j, the R_j^2 from regressing the actual bias $E_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$ on the synthetic bias $\hat{E}_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$*

has the following upper bound

$$R_j^2 \;\leq\; \frac{\text{Cov}(\varepsilon_{j,t+1}, \hat{\varepsilon}_{t+1})' \hat{\Sigma}^{-1} \text{Cov}(\varepsilon_{j,t+1}, \hat{\varepsilon}_{t+1})}{\text{Var}(\varepsilon_{j,t})}. \tag{A11}$$

This is the same bound as the one derived in Proposition 5.

A.6. Excess returns without assuming log-normality

Lemma 5 shows that we can cleanly decompose excess returns into two comovements, one associated with beliefs and one associated with risk/preferences, if variables are log-normally distributed. If we make no assumptions about the distributions of the variables, then we have the following relationship.

Lemma 6. *Let* $R^e_{j,t+1} \equiv R_{j,t+1}/R^f_t$. *If* $E^*_t\left[\tilde{M}_{t+1}R_{j,t+1}\right] = E^*_t\left[\tilde{M}_{t+1}\right]R^f_t = 1$, *then*

$$\begin{aligned} E\left[R^e_{j,t+1}\right] \;&=\; 1 - Cov\left(S_{t+1}, R^e_{j,t+1}\right) - Cov\left(\tilde{M}_{t+1}R^f_t, R^e_{j,t+1}\right) \\ &\quad - Cov\left(\left[S_{t+1}-1\right]\left[\tilde{M}_{t+1}R^f_t - 1\right], R^e_{j,t+1}\right). \end{aligned} \tag{A12}$$

An asset's expected return can be impacted by the agent understating certain states of the world, which is reflected by S_{t+1}. An asset can also be affected by the agent having a low preference for payoffs in certain states of the world, which is reflected by $\tilde{M}_{t+1}R^f_t$. Note that the inclusion of R^f_t in $\tilde{M}_{t+1}R^f_t$ is simply to offset the fact that $\tilde{M}_{t+1}$ includes time discounting as well as preferences for different states of the world. The combined $\tilde{M}_{t+1}R^f_t$ solely reflects preferences for different states of the world.

The first RHS covariance in equation (A12) tells us the excess return that is explained by distorted beliefs about the probability of different states (S_{t+1}), assuming equal preferences for payoffs in all states of the world, i.e., $\tilde{M}_{t+1}R^f_t = 1$. Similarly, the second RHS covariance tells us the excess return that is explained by preferences for different states of the world, assuming the SBF is equal for all states of the world, i.e., $S_{t+1} = 1$. The third RHS covariance captures the interaction between beliefs and preferences. For example, is the agent overstating/understating the probability of states of the world that she particularly cares about (i.e., states that have a high preference)?

Combining Lemma 3 and Lemma 6 provides a method to estimate the SBF S_{t+1} and gauge its ability to explain excess returns without imposing any assumptions other than subjective expectations being coherent. We can estimate the SBF from expectations data following Lemma 3 and then measure the first RHS covariance in equation (A12) even if we do not know $\tilde{M}_{t+1}$. As mentioned in Section I.B, we choose to make assumptions about the objective distribution so that we can ensure the estimated SBF is always non-negative, as this allows us to easily relate the estimated SBF to models of expectation-formation. However, for research where a negative estimated SBF would not be problematic, Lemmas 3 and 6 provide a useful minimum-assumptions method to connect subjective expectations to excess returns.

B. 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 Riesz 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},

$$E^*_{i,t}[X_{t+1}] \quad = \quad E_t[S_{i,t+1}X_{t+1}].$$

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

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

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

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

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

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

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

Proof of Lemma 3: Equation (4) states that

$$E_t\left[S_{t+1}X_{t+1}\right] \;=\; E_t\left[X_{t+1}\right] + Cov_t\left(S_{t+1}, X_{t+1}\right)$$

Then, we simply note that $Cov_t\left(S_{t+1}, X_{t+1}\right) = Cov_t\left(\beta_t'\varepsilon_{t+1}, \varepsilon_{t+1}\right) = E_t^{*}\left[X_{t+1}\right] - E_t\left[X_{t+1}\right]$ given the definition of β_t.

Proof of Proposition 3: We show the proof of Proposition 3 before the proof of Proposition 2 because Proposition 3 is more general. Let $\phi(\cdot)$ denote the standard normal distribution. We have that

$$\begin{aligned}
s_{t+1} &= -\frac{1}{2}\beta_t'\beta_t + \beta_t'\varepsilon_{t+1} \\
&= -\frac{1}{2}\left(\varepsilon_{t+1} - \beta_t\right)'\left(\varepsilon_{t+1} - \beta_t\right) + \frac{1}{2}\varepsilon_{t+1}'\varepsilon_{t+1} \qquad \text{(A13)}
\end{aligned}$$

which implies that

$$S_{t+1} \;=\; \frac{\phi\left(\left(\varepsilon_{t+1} - \beta_t\right)'\left(\varepsilon_{t+1} - \beta_t\right)\right)}{\phi\left(\varepsilon_{t+1}'\varepsilon_{t+1}\right)}. \qquad \text{(A14)}$$

Equation (A14) shows that S_{t+1} is equal to the ratio of two normal pdf's, one centered at β_t and the other centered at 0. Thus, when calculating expectations $E_t\left[S_{t+1}X_{t+1}\right] = \int_{\varepsilon_{t+1}} \phi\left(\varepsilon_{t+1}'\varepsilon_{t+1}\right) S_{t+1} f_t\left(\varepsilon_{t+1}\right) d\varepsilon_{t+1}$, the inclusion of S_{t+1} means that we change from using the objective pdf $\phi\left(\varepsilon_{t+1}'\varepsilon_{t+1}\right)$ to the distorted pdf centered at β_t. Specifically,

$$\begin{aligned}
E_t\left[S_{t+1}X_{t+1}\right] &= \int_{\varepsilon_{t+1}} \phi\left(\varepsilon_{t+1}'\varepsilon_{t+1}\right) S_{t+1} f_t\left(\varepsilon_{t+1}\right) d\varepsilon_{t+1} \\
&= \int_{\varepsilon_{t+1}} \phi\left(\left(\varepsilon_{t+1}-\beta_t\right)'\left(\varepsilon_{t+1}-\beta_t\right)\right) f_t\left(\varepsilon_{t+1}\right) d\varepsilon_{t+1} \\
&= E_t\left[f_t\left(\varepsilon_{t+1}+\beta_t\right)\right].
\end{aligned}$$

Given that $\beta_t = h_t^{-1}\left(E_t^*\left[X_{t+1}\right]\right)$, we know that $E_t\left[f_t\left(\varepsilon_{t+1}+\beta_t\right)\right] = E_t^*\left[X_{t+1}\right]$.

Proof of Proposition 2: Proposition 2 is a special case of Proposition 3. Let $\eta_{t+1} \equiv \Sigma^{-1/2}\varepsilon_{t+1}$ denote the standard normal shocks studied in Proposition 3. We have $X_{t+1} = f_t\left(\eta_{t+1}\right) = E_t\left[X_{t+1}\right] + \Sigma^{1/2}\eta_{t+1}$ and $h_t\left(\beta\right) = E_t\left[f_t\left(\beta+\eta_{t+1}\right)\right] = E_t\left[X_{t+1}\right] + \Sigma^{1/2}\beta$. This means $h_t^{-1}\left(E_t^*\left[X_{t+1}\right]\right)$ equals $\Sigma^{-1/2}\left(E_t^*\left[X_{t+1}\right] - E_t\left[X_{t+1}\right]\right)$. Thus, $E_t^*\left[X_{t+1}\right] = E_t\left[S_{t+1}X_{t+1}\right]$ for $s_{t+1} \equiv -\frac{1}{2}\beta_t'\beta_t + \beta_t'\eta_{t+1}$ and $\beta_t = \Sigma^{-1/2}\left(E_t^*\left[X_{t+1}\right] - E_t\left[X_{t+1}\right]\right)$. The final step is noting that this s_{t+1} is identical to $s_{t+1} \equiv -\frac{1}{2}\beta_t'\Sigma\beta_t + \beta_t'\varepsilon_{t+1}$ with $\beta_t = \Sigma^{-1}\left(E_t^*\left[X_{t+1}\right] - E_t\left[X_{t+1}\right]\right)$.

Proof of Proposition 4: Given equation (4) and the fact that $E_t\left[S_{t+1}\hat{X}_{t+1}\right] = E_t^*\left[\hat{X}_{t+1}\right] = E_t\left[\hat{S}_{t+1}\hat{X}_{t+1}\right]$, we know that

$$Cov_t\left(S_{t+1}, \hat{X}_{t+1}\right) = Cov_t\left(\hat{S}_{t+1}, \hat{X}_{t+1}\right).$$

Thus, $Cov_t\left(S_{t+1}-\hat{S}_{t+1}, \hat{X}_{t+1}\right) = 0$. From here, it follows that

$$\begin{aligned}
E_t^*\left[X_{t+1}\right] - E_t\left[\hat{S}_{t+1}X_{t+1}\right] &= E_t\left[\left(S_{t+1}-\hat{S}_{t+1}\right)X_{t+1}\right] \\
&= E_t\left[\left(S_{t+1}-\hat{S}_{t+1}\right)u_{t+1}\right] \\
&= Cov_t\left(S_{t+1}-\hat{S}_{t+1}, u_{t+1}\right) \\
&= Cov_t\left(u_{t+1}^{S_{t+1}-\hat{S}_{t+1}}, u_{t+1}\right).
\end{aligned}$$

Proof of Lemma 4: By Proposition 2 applied to $\hat{X}_{t+1}$, we know that $\hat{S}_{t+1}$ is a deterministic function of $\hat{X}_{t+1}$. With $(X_{t+1}, \hat{X}_{t+1})$ conditionally normal, u_{t+1} satisfies $E_t\left[u_{t+1}|\hat{X}_{t+1}\right] = 0$.

Hence,

$$\begin{aligned} E_t\left[\hat{S}_{t+1}u_{t+1}\right] &= E_t\left[E_t\left[\hat{S}_{t+1}u_{t+1}|\hat{X}_{t+1}\right]\right] \\ &= E_t\left[E_t\left[\hat{S}_{t+1}|\hat{X}_{t+1}\right]E_t\left[u_{t+1}|\hat{X}_{t+1}\right]\right] \\ &= 0 \end{aligned}$$

where the second line uses the fact that $Cov_t\left(\hat{S}_{t+1}, u_{t+1}|\hat{X}_{t+1}\right)$ is 0 since $\hat{S}_{t+1}$ is known given $\hat{X}_{t+1}$ and the third line uses the fact that $E_t\left[u_{t+1}|\hat{X}_{t+1}\right]$ is 0. Thus, $Cov_t(\hat{S}_{t+1}, u_{t+1}) = 0$, and the rest of the proof follows from Proposition 4.

Proof of Proposition 5: For clarity, the entire proof focuses on a fixed variable j. For exposition, we assume that $\hat{X}_{t+1}$ is the first two elements of the larger vector of variables Z_{t+1}. The proof can then easily be extended to an arbitrary number of elements.

We can define the standardized shocks as $\begin{pmatrix} \tilde{\varepsilon}_{1,t-\tau} \\ \tilde{\varepsilon}_{2,t-\tau} \end{pmatrix} = \hat{\Sigma}^{-1/2}\begin{pmatrix} \varepsilon_{1,t-\tau} \\ \varepsilon_{2,t-\tau} \end{pmatrix}$. Let $\Upsilon \equiv \begin{pmatrix} \tilde{\Sigma}_{j,1} \\ \tilde{\Sigma}_{j,2} \end{pmatrix}$ where $\tilde{\Sigma}_{j,k} = Cov(\varepsilon_j, \tilde{\varepsilon}_k)$ for $k = 1, 2$. Denote the projected synthetic lags as $w_\tau \equiv P\begin{pmatrix} \delta_{1,\tau} \\ \delta_{2,\tau} \end{pmatrix}$, where $P = \begin{pmatrix} \Upsilon_1 & 0 \\ 0 & \Upsilon_2 \end{pmatrix}$, and $|w_\tau| = w_\tau' w_\tau$.

It is convenient to collect the lag weights of the variable into an infinite vector $\overrightarrow{\boldsymbol{\delta}} \equiv (\delta_{j,0}, \delta_{j,1}, ...)$ with norm $||\cdot||$. Similarly, we collect $\overrightarrow{\boldsymbol{w}} \equiv (w_0, w_1, ...)$ and $\overrightarrow{|\boldsymbol{w}|} \equiv (|w_0|, |w_1|, ...)$. Then, the R_j^2 from regressing the actual bias $E_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$ on the synthetic bias $\hat{E}_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$ is the squared correlation,

$$
\begin{aligned}
R_j^2 &= \frac{Cov\left(\sum_{\tau=0}^{\infty}\delta_{j,\tau}\varepsilon_{j,t-\tau}, \sum_{\tau=0}^{\infty}\begin{pmatrix}\tilde{\varepsilon}_{1,t-\tau}\\ \tilde{\varepsilon}_{2,t-\tau}\end{pmatrix}' w_\tau\right)^2}{Var\left(\sum_{\tau=0}^{\infty}\delta_{j,\tau}\varepsilon_{j,t-\tau}\right) Var\left(\sum_{\tau=0}^{\infty}\begin{pmatrix}\tilde{\varepsilon}_{1,t-\tau}\\ \tilde{\varepsilon}_{2,t-\tau}\end{pmatrix}' w_\tau\right)} \\
&= \frac{\left\langle \overrightarrow{\boldsymbol{\delta}}, \Upsilon' \overrightarrow{\boldsymbol{w}} \right\rangle^2}{Var(\varepsilon_{j,t}) \left\| \overrightarrow{\boldsymbol{\delta}} \right\|^2 \left\| \overrightarrow{|\boldsymbol{w}|} \right\|^2} \\
&\leq \frac{\left\| \overrightarrow{\boldsymbol{\delta}} \right\|^2 \left\| \Upsilon' \overrightarrow{\boldsymbol{w}} \right\|^2}{Var(\varepsilon_{j,t}) \left\| \overrightarrow{\boldsymbol{\delta}} \right\|^2 \left\| \overrightarrow{|\boldsymbol{w}|} \right\|^2} \\
&\leq \frac{\left\| (\Upsilon'\Upsilon) \overrightarrow{|\boldsymbol{w}|} \right\|^2}{Var(\varepsilon_{j,t}) \left\| \overrightarrow{|\boldsymbol{w}|} \right\|^2} = \frac{\Upsilon'\Upsilon}{Var(\varepsilon_{j,t})}.
\end{aligned}
$$

The second line comes after applying the fact that errors are serially uncorrelated. The third and fourth lines are Cauchy-Schwarz inequalities. Section A.5 extends this result to the case where the errors are serially correlated by arriving to the same upper bound under additional sufficient conditions.

Proof of Lemma 5: For concision of notation, we drop the j subscript since this Lemma can be applied to each asset individually. 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 $\left(\varepsilon_{t+1}^{r}, \varepsilon_{t+1}^{s}, \varepsilon_{t+1}^{m}\right)$ are potentially correlated Gaussian shocks. We have that

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

where the 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

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

In a model where we assume that the belief distortion is given by the measured $\hat{s}_{t+1}$, we have equation (29). In a more general setting where s_{t+1} can differ from the measured $\hat{s}_{t+1}$, we have equation (33).

Proof of Proposition 6: Again, let $\phi(\cdot)$ denote the standard normal pdf. The logic of the proof is virtually identical to the proof of Proposition 3. We have that

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

which implies that

$$
S_{t+1} = \frac{\phi\left(\left(\varepsilon_{t+1} - \Sigma_t^* \beta_t \right)' \Sigma_t^{*-1} \left(\varepsilon_{t+1} - \Sigma_t^* \beta_t \right) \right)}{\phi\left(\varepsilon_{t+1}' \Sigma_t^{-1} \varepsilon_{t+1} \right)}. \tag{A15}
$$

Equation (A15) shows that S_{t+1} is equal to the ratio of two normal pdf's, one centered at $\Sigma_t^* \beta_t$ with covariance matrix Σ_t^* and the other centered at 0 with covariance matrix Σ_t.

Thus, when calculating expectations $E_t\left[S_{t+1}X_{t+1}\right]=\int_{\varepsilon_{t+1}}\phi\left(\varepsilon_{t+1}'\Sigma_t^{-1}\varepsilon_{t+1}\right)S_{t+1}f_t\left(\varepsilon_{t+1}\right)d\varepsilon_{t+1}$, the inclusion of S_{t+1} means that we change from using the objective pdf $\phi\left(\varepsilon_{t+1}'\Sigma_t^{-1}\varepsilon_{t+1}\right)$ to the distorted pdf. Specifically,

$$\begin{aligned}
E_t\left[S_{t+1}X_{t+1}\right] &= E_t\left[X_{t+1}\right]+\int_{\varepsilon_{t+1}}\phi\left(\varepsilon_{t+1}'\Sigma_t^{-1}\varepsilon_{t+1}\right)S_{t+1}\varepsilon_{t+1}d\varepsilon_{t+1} \\
&= E_t\left[X_{t+1}\right]+\int_{\varepsilon_{t+1}}\phi\left(\left(\varepsilon_{t+1}-\Sigma_t^*\beta_t\right)'\Sigma_t^{*-1}\left(\varepsilon_{t+1}-\Sigma_t^*\beta_t\right)\right)\varepsilon_{t+1}d\varepsilon_{t+1} \\
&= E_t\left[X_{t+1}\right]+\Sigma_t^*\beta_t.
\end{aligned}$$

Given that $\beta_t=\Sigma_t^{*-1}\left(E_t^*\left[X_{t+1}\right]-E_t\left[X_{t+1}\right]\right)$, we know that $E_t\left[S_{t+1}X_{t+1}\right]=E_t^*\left[X_{t+1}\right]$. Similarly, we have that the inclusion of S_{t+1} means that we change from the objective pdf to the distorted pdf when calculating the covariance,

$$\begin{aligned}
E_t\left[S_{t+1}X_{t+1}X_{t+1}'\right] &= \int_{\varepsilon_{t+1}}\phi\left(\varepsilon_{t+1}'\Sigma_t^{-1}\varepsilon_{t+1}\right)S_{t+1}X_{t+1}X_{t+1}'d\varepsilon_{t+1} \\
&= \int_{\varepsilon_{t+1}}\phi\left(\left(\varepsilon_{t+1}-\Sigma_t^*\beta_t\right)'\Sigma_t^{*-1}\left(\varepsilon_{t+1}-\Sigma_t^*\beta_t\right)\right)X_{t+1}X_{t+1}'d\varepsilon_{t+1} \\
&= \left(E_t\left[X_{t+1}\right]+\Sigma_t^*\beta_t\right)\left(E_t\left[X_{t+1}\right]+\Sigma_t^*\beta_t\right)' \\
&\quad +\int_{\varepsilon_{t+1}}\phi\left(\left(\varepsilon_{t+1}-\Sigma_t^*\beta_t\right)'\Sigma_t^{*-1}\left(\varepsilon_{t+1}-\Sigma_t^*\beta_t\right)\right)\left(\varepsilon_{t+1}-\Sigma_t^*\beta_t\right)\left(\varepsilon_{t+1}-\Sigma_t^*\beta_t\right)'d\varepsilon_{t+1} \\
&= \left(E_t\left[X_{t+1}\right]+\Sigma_t^*\beta_t\right)\left(E_t\left[X_{t+1}\right]+\Sigma_t^*\beta_t\right)'+\Sigma_t^* \\
&= E_t\left[S_{t+1}X_{t+1}\right]E_t\left[S_{t+1}X_{t+1}\right]'+\Sigma_t^*.
\end{aligned}$$

Proof of Lemma 6: For concision of notation, we drop the j subscript since this Lemma can be applied to each asset individually. Given $E_t^*\left[\tilde{M}_{t+1}R_{t+1}\right]=E_t^*\left[\tilde{M}_{t+1}\right]R_t^f=1$ and $R_{t+1}^e\equiv R_{t+1}/R_t^f$, we know that

$$E_t\left[S_{t+1}\tilde{M}_{t+1}R_t^f\left(R_{t+1}^e-1\right)\right] = 0$$

which implies $E\left[S_{t+1}\tilde{M}_{t+1}R_t^f\left(R_{t+1}^e-1\right)\right]=0$. Using the definition of covariance and the fact that $E\left[S_{t+1}\tilde{M}_{t+1}R_t^f\right]=E\left[E_t\left[S_{t+1}\tilde{M}_{t+1}R_t^f\right]\right]=1$, we have

$$E\left[R_{t+1}^e\right] = 1-Cov\left(S_{t+1}\tilde{M}_{t+1}R_t^f,R_{t+1}^e\right).$$

The final step is simply to expand $Cov\left(S_{t+1}\tilde{M}_{t+1}R_t^f,R_{t+1}^e\right)$ into the three covariance terms in the RHS of equation (A12).

Proof of Proposition 7: For exposition, we assume that $\hat{X}_{t+1}$ is the first two elements of the larger vector of variables Z_{t+1}. The proof can then easily be extended to an arbitrary number of elements. We define the same objects as in the proof of Proposition 5. We also define the standardized construction shocks as $\begin{pmatrix} \tilde{\varepsilon}_{1,t} \\ \tilde{\varepsilon}_{2,t} \end{pmatrix} = \hat{\Sigma}^{-1/2} \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{pmatrix}$, the lag-h covariances $\tilde{\Omega}_{jk}(h) \equiv Cov(\varepsilon_{j,t}, \tilde{\varepsilon}_{k,t-h})$, the vector of shocks for variable j $\overrightarrow{\boldsymbol{\varepsilon}}_{\boldsymbol{j}} \equiv (\varepsilon_{j,0}, \varepsilon_{j,1}, ...)$, the stacked vectors $\overrightarrow{\tilde{\boldsymbol{\varepsilon}}} = (\tilde{\varepsilon}_{1,0}, \tilde{\varepsilon}_{1,1}, ..., \tilde{\varepsilon}_{1,\tau}, ..., \tilde{\varepsilon}_{2,0}, \tilde{\varepsilon}_{2,1}, ..., \tilde{\varepsilon}_{2,\tau}, ..)$, $\overrightarrow{\tilde{\boldsymbol{\delta}}} = (\delta_{1,0}, \delta_{1,1}, ..., \delta_{1,\tau}, ..., \delta_{2,0}, \delta_{2,1}, ..., \delta_{2,\tau}, ..)$, and $\overrightarrow{\boldsymbol{\zeta}} = (w_{1,0}, w_{1,1}..., w_{1,\tau}, ..., w_{2,0}, w_{2,1}, ...w_{2,t}..) = (P \bigotimes \boldsymbol{I})\, \overrightarrow{\tilde{\boldsymbol{\delta}}}$ where $w_{1,\tau}$ and $w_{2,\tau}$ are the two elements of w_τ.

We now solve an optimization problem to find the lag weights that maximize R_j^2. This problem is equivalent to:

$$\begin{aligned} \max_{\overrightarrow{\boldsymbol{\delta}}, \overrightarrow{\boldsymbol{\zeta}}} \quad & Corr\left(\sum_{\tau=0} \delta_{j,\tau}\varepsilon_{j,t-\tau}, \sum_{\tau=0} w'_\tau \begin{pmatrix} \tilde{\varepsilon}_{1,t} \\ \tilde{\varepsilon}_{2,t} \end{pmatrix}\right) \\ s.t. \quad & \overrightarrow{\boldsymbol{\delta}} \geq 0,\ \overrightarrow{\tilde{\boldsymbol{\delta}}} \geq 0 \end{aligned}$$

$$\begin{aligned} = \max_{\overrightarrow{\boldsymbol{\delta}}, \overrightarrow{\boldsymbol{\zeta}}} \quad & Cov\left(\overrightarrow{\boldsymbol{\delta}}'\overrightarrow{\boldsymbol{\varepsilon}}_{\boldsymbol{j}}, \overrightarrow{\boldsymbol{\zeta}}'\overrightarrow{\tilde{\boldsymbol{\varepsilon}}}\right) \\ s.t. \quad & Var\left(\overrightarrow{\boldsymbol{\delta}}'\overrightarrow{\boldsymbol{\varepsilon}}_{\boldsymbol{j}}\right) = 1,\ Var\left(\overrightarrow{\boldsymbol{\zeta}}'\overrightarrow{\tilde{\boldsymbol{\varepsilon}}}\right) = 1,\ \overrightarrow{\boldsymbol{\delta}} \geq 0,\ \overrightarrow{\tilde{\boldsymbol{\delta}}} \geq 0 \end{aligned}$$

$$\begin{aligned} = \max_{\overrightarrow{\boldsymbol{\delta}}, \overrightarrow{\boldsymbol{\zeta}}} \quad & \overrightarrow{\boldsymbol{\delta}}'\begin{pmatrix} \boldsymbol{\Gamma_{j1}} & \boldsymbol{\Gamma_{j2}} \end{pmatrix}\overrightarrow{\boldsymbol{\zeta}} \\ s.t. \quad & \overrightarrow{\boldsymbol{\delta}}'\boldsymbol{\Gamma_{jj}}\overrightarrow{\boldsymbol{\delta}} = 1,\ \overrightarrow{\boldsymbol{\zeta}}'\begin{pmatrix} \boldsymbol{\Gamma_{11}} & \boldsymbol{\Gamma_{12}} \\ \boldsymbol{\Gamma_{21}} & \boldsymbol{\Gamma_{22}} \end{pmatrix}\overrightarrow{\boldsymbol{\zeta}} = 1,\ \overrightarrow{\boldsymbol{\delta}} \geq 0,\ \overrightarrow{\tilde{\boldsymbol{\delta}}} \geq 0, \end{aligned}$$

where $\Gamma_{jk} = \left[\tilde{\Omega}_{jk}\left(|\tau - \tau'|\right)\right]_{\tau,\tau' \geq 0}$ is the Toeplitz covariance matrix of the shocks.
The FOC for $\overrightarrow{\boldsymbol{\delta}}$ and $\overrightarrow{\boldsymbol{\zeta}}$ are:

$$\begin{pmatrix} \mathbf{\Gamma}_{j1} & \mathbf{\Gamma}_{j2} \end{pmatrix} \overrightarrow{\boldsymbol{\zeta}} \leq \lambda \mathbf{\Gamma}_{jj} \overrightarrow{\boldsymbol{\delta}}$$

$$\left(\boldsymbol{I} \bigotimes P\right) \begin{pmatrix} \mathbf{\Gamma}_{j1} & \mathbf{\Gamma}_{j2} \end{pmatrix}' \overrightarrow{\boldsymbol{\delta}} \leq \mu \left(\boldsymbol{I} \bigotimes P\right) \begin{pmatrix} \mathbf{\Gamma}_{11} & \mathbf{\Gamma}_{12} \\ \mathbf{\Gamma}_{21} & \mathbf{\Gamma}_{22} \end{pmatrix} \overrightarrow{\boldsymbol{\zeta}},$$

where λ, μ are the Lagrange multipliers. Given the assumptions on $\Omega_{jk}(h)$ in Appendix A.5, the solution for this system of equations is $\delta^*_{j,0} = \tilde{\Omega}_{jj}(0)^{-1/2}$, $w^*_{1,0} = \tilde{\Omega}_{j1}(0) \left(\tilde{\Omega}_{j1}(0)^2 + \tilde{\Omega}_{j2}(0)^2\right)^{-1/2}$, $w^*_{2,0} = \tilde{\Omega}_{j2}(0) \left(\tilde{\Omega}_{j1}(0)^2 + \tilde{\Omega}_{j2}(0)^2\right)^{-1/2}$, and $\delta^*_{j,\tau} = w^*_{1,\tau} = w^*_{2,\tau} = 0$ for all $\tau > 0$. One can check that the unit variance constraints are straightforwardly satisfied and that the evaluated FOC strictly bind on $\delta^*_{j,0}$, $w^*_{1,0}$, $w^*_{2,0}$, and are satisfied for all other $\delta^*_{j,\tau}, w^*_{1,\tau}, w^*_{2,\tau}$ by condition (A10).

Plugging these optimal choices for $\overrightarrow{\boldsymbol{\delta}}$ and $\overrightarrow{\boldsymbol{\zeta}}$ into the calculation for the R^2_j, we have that the maximum possible value of R^2_j is

$$\begin{aligned} R^2_j &\leq Cov\left(\delta^*_{j,0}\varepsilon_{j,t}, \begin{pmatrix} \tilde{\varepsilon}_{1,t} \\ \tilde{\varepsilon}_{2,t} \end{pmatrix}' w^*_\tau\right)^2 = \left(\delta^*_{j,0}\, \Upsilon' w^*_0\right)^2 \\ &\leq \delta^{*2}_{j,0} \left(\Upsilon'\Upsilon\right) |w^*_0|^2 = \frac{\Upsilon'\Upsilon}{Var(\varepsilon_{j,t})}. \end{aligned}$$

C. *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 1968Q4, which included forecasts for *ngdp*, *pgdp*, *cprof*, *unemp*, *indp*, *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

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
rgdp	Real GDP
rcon	Real personal consumption expenditures
cpi	CPI inflation rate
unemp	Unemployment rate
indp	Index of industrial production
tbill	Three-month Treasury bill rate
aaa	Moody's Aaa corporate bond yield
rnrinv	Real nonresidential fixed investment
rresinv	Real residential fixed investment
rgf	Real federal government consumption and gross investment
rgsl	Real state and local government consumption and gross investment
housing	Level of housing starts
rcbi	Real net change in private inventories
rexport	Real net exports
cprof	Level of nominal corporate profits

survey variables included since the 1981Q3 survey are *tbill*, *aaa*, *rcon*, *rnrinv*, *rresinv*, *rgf*, *rgsl*, *rcbi*, *rexport*, and *cpi*. In our analysis, we use all variables for which we have survey data since 1981Q3. Note that there is some redundancy in the variables due to the fact that 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 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 Chip survey has variables added, and sometimes removed, particularly in the earlier years of the survey. A round of significant changes occurred in 1988Q1 when a host of interest rates were added

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 Forecasters dataset shown in Table AI.

Variable Name	Variable Description
ffr	Federal funds rate
libor	3-month LIBOR (SOFR from 2022:Q1)
6m	6-month Treasury bill rate
1y	1-year Treasury bill rate
2y	2-year Treasury note rate
5y	5-year Treasury note rate
10y	10-year Treasury note rate
prime	Prime bank rate
mortg	Home mortgage rate

to the survey: *1y, 6m, libor* (3-month)*, 2y,* and *10y.* Earlier variables available include the prime rate (*prime)*, the fed funds rate (*ffr)*, *munis*, *aaa*, mortgage rate (*mortg*), among others. In our analysis, we use all variables for which we have survey data since 1988Q1. 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).[34] 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 over our sample. Table AII shows the resulting 9 Blue Chip variables.

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*, *rnrinv*, *rresinv*, *rgf*, *rgsl*, *housing*, *cprof*, we convert point forecasts into implied growth

[34] We 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.

rates. For example, for $rgdp$, we compute the forecaster's point forecast for $rgdp_{t+4}$ divided by the initial release of $rgdp_{t-1}$ 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 (*rexport*) 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 and the Federal Reserve Board. 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) to match the fact that survey participants are asked to forecast the average of the interest rate over that 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 2001Q2.

D. Additional comparisons of synthetic and subjective expectations

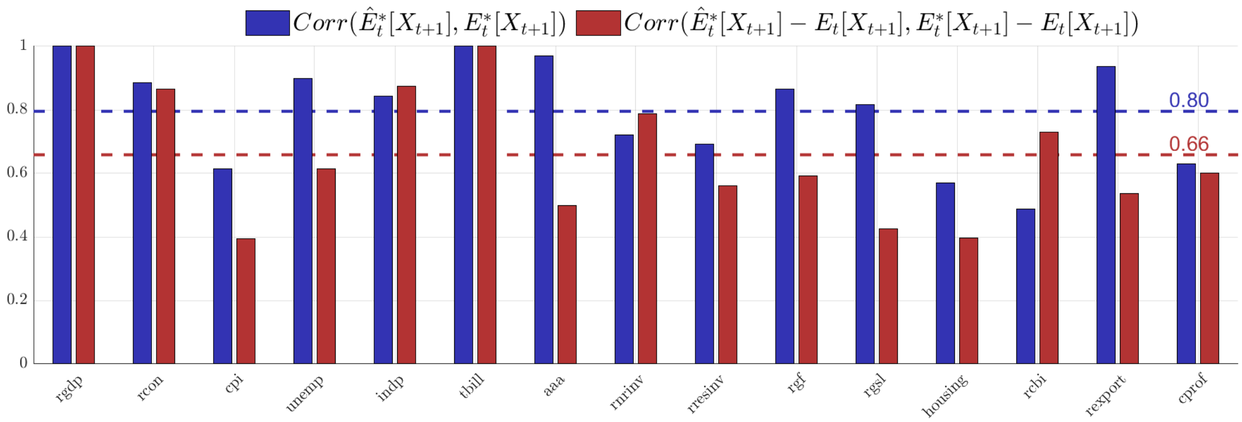

Figure A1. Correlations of synthetic expectations and the Survey of Professional Forecasters. The figure evaluates the accuracy of the log SBF $\hat{s}_{t+1}$ in generating synthetic expectations for the survey variables. The log SBF $\hat{s}_{t+1}$ is constructed using only RGDP growth and the T-bill rate biases. The blue bars show the correlation of survey expectations with synthetic expectations for the 15 variables. The red bars show the correlation of the survey biases with the synthetic biases. The dashed lines show the average values across the 15 variables.

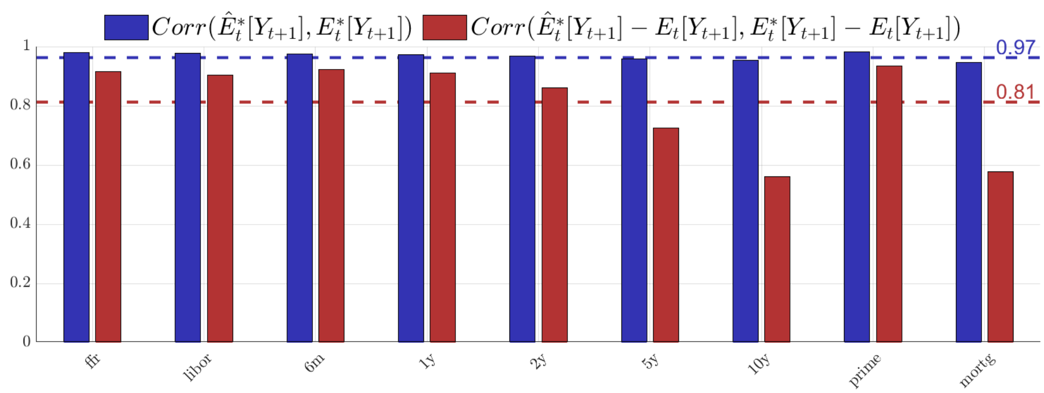

Figure A2. Correlations of synthetic expectations and Blue Chip expectations. The figure evaluates the accuracy of the log SBF $\hat{s}_{t+1}$ in generating synthetic expectations for Blue Chip variables. The log SBF is constructed using only the RGDP growth and T-bill rate biases from the Survey of Professional Forecasters. The blue bars show the correlation of survey expectations with synthetic expectations for the 9 variables. The red bars show the correlation of the survey biases with the synthetic biases. The dashed lines show the average values across the 9 variables.

Table AIII

Matching the Blue Chip Forecasts

This table evaluates the ability of each individual synthetic bias to match the survey bias of the 9 variables in Blue Chip. The table shows the intercept and slope coefficients from regressing $E_t^*[Y_{j,t+1}] - E_t[Y_{j,t+1}]$ on $\hat{E}_t^*[Y_{j,t+1}] - E_t[Y_{j,t+1}]$, where $\hat{E}_t^*[Y_{j,t+1}]$ is the synthetic expectation constructed using the log SBF $\hat{s}_{t+1}$. The last row shows a joint regression across all variables. In the joint regression, the survey and synthetic bias for each variable are scaled by the standard deviation of the survey bias to ensure comparability across variables. The fifth column shows the R^2 of the regression, and the last column shows the R^2 when the slope of the regression b is constrained to be 1. The row labeled "Average" shows the average R^2 and constrained R^2 across the first 9 rows. Superscripts indicate significance at the 1% (∗∗∗), 5% (∗∗), and 10% (∗) level.

$Y_{j,t+1}$	a	se_a	b	se_b	R^2	constrained R^2
ffr	-0.000	(0.000)	1.209***	(0.054)	0.840	0.815
libor	-0.000	(0.001)	1.190***	(0.087)	0.816	0.795
6m	-0.000	(0.001)	1.250***	(0.064)	0.852	0.818
1y	0.000	(0.001)	1.229***	(0.074)	0.827	0.798
2y	0.001	(0.001)	1.201***	(0.089)	0.742	0.721
5y	0.002**	(0.001)	1.131***	(0.140)	0.527	0.520
10y	0.003***	(0.001)	1.003***	(0.201)	0.314	0.314
prime	-0.001*	(0.000)	1.212***	(0.044)	0.873	0.846
mortg	0.002*	(0.001)	1.131***	(0.321)	0.334	0.330
Average					0.681	0.662
Joint	0.119***	(0.039)	1.174***	(0.039)	0.645	0.614

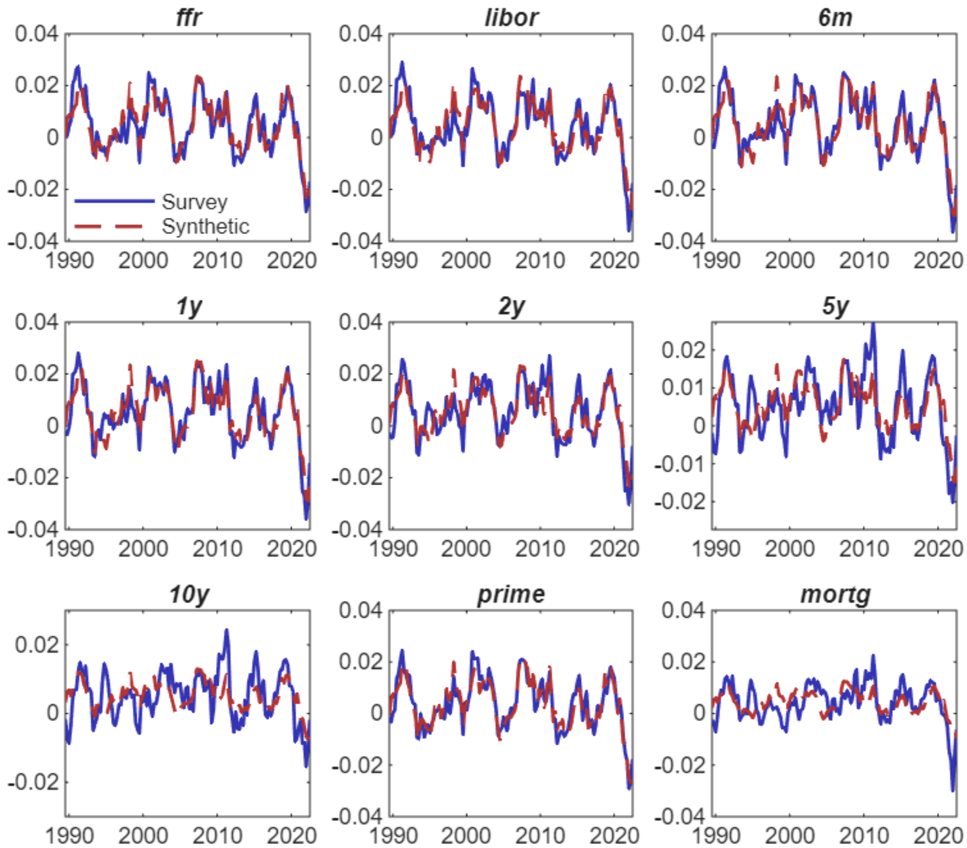

Figure A3. Comparison between synthetic biases and Blue Chip biases. This figure shows how well the synthetic bias $\hat{E}_t^* [Y_{j,t+1}] - E_t [Y_{j,t+1}]$, constructed using the SBF $\hat{s}_{t+1}$, matches the survey expectation bias $E_t^* [Y_{j,t+1}] - E_t [Y_{j,t+1}]$ for each of the 9 variables in Blue Chip. The blue time series show each of the survey biases, and the red time series show the fitted synthetic biases. The sample period is 1989Q2-2022Q2.

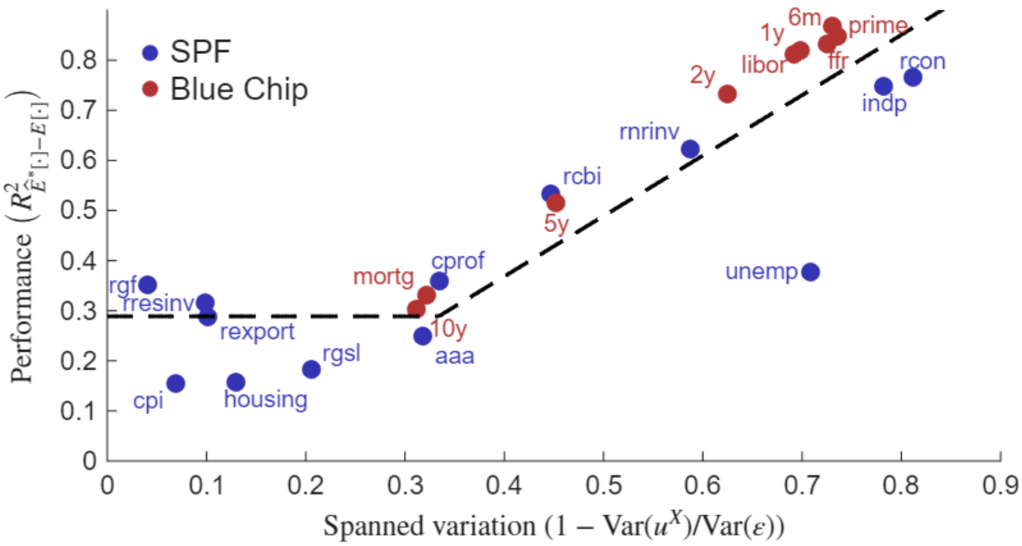

Figure A4. Performance and spanning of synthetic biases. This figure repeats the spanning exercise of Figure 3 using the synthetic biases rather than the level of the synthetic expectations. The horizontal axis shows the degree to which each variable's shock $\varepsilon_{j,t+1}$ is spanned by the RGDP and T-bill shocks, $1 - Var(u_{j,t+1})/Var(\varepsilon_{j,t+1})$, where $u_{j,t+1}$ is the unspanned portion of each shock. The vertical axis shows the R^2 of the synthetic bias $Corr\left(\hat{E}^*[\cdot] - E[\cdot], E^*[\cdot] - E[\cdot]\right)^2$ relative to the R^2 of the perfect-foresight benchmark in fitting the survey biases, which is equal to zero. Each point is one of the 22 Survey of Professional Forecasters and Blue Chip variables, excluding RGDP and T-bill. The sample period is 1982Q4-2022Q2.

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

The top panel of Figure A5 shows the time series for the gap between subjective and statistical expectations for RGDP growth and the T-bill rate ($E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]$). In general, the gap between subjective and statistical expectations for RGDP growth and the T-bill rate are weakly correlated (0.27). 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. Unconditionally, we find a significantly positive average bias for T-bill (0.0054 log points with a p-value of 0.003) indicating that forecasters on average did not completely predict the large decline in interest rates over the past 40 years. In comparison, we find that average value for the RGDP growth bias is smaller and insignificant, at 0.0023 log points with a p-value of 0.255. Compared to statistical expectations, subjective expectations

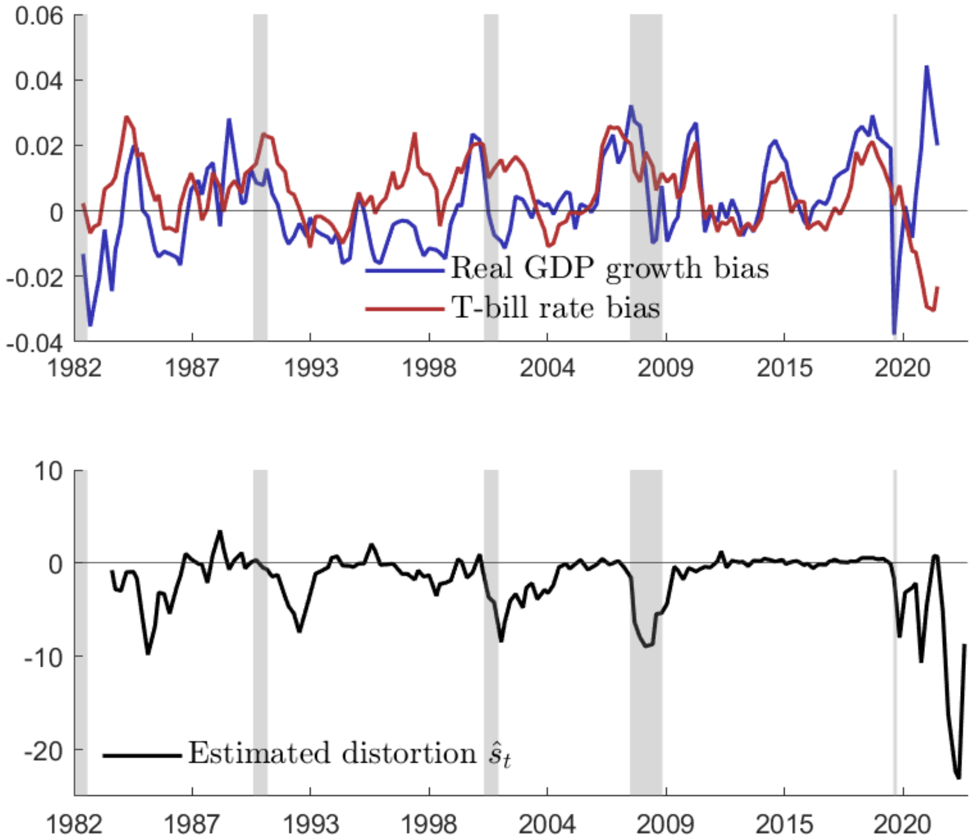

Figure A5. Biases and estimated log SBF over time. The top figure shows the RGDP growth and T-bill rate biases $E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]$ over the main sample. The bottom figure shows the log SBF $\hat{s}_t$ estimated from these two biases using equation (11). Note that the sample for $\hat{s}_t$ starts one year after the sample for $E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]$, as $E_t^*\left[\hat{X}_{t+1}\right] - E_t\left[\hat{X}_{t+1}\right]$ are combined with the future shocks $\hat{\varepsilon}_{t+1}$ to calculate the realized value for the log SBF. Shaded bars indicate NBER recessions.

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.

F. 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 V, we obtain the monthly returns from the data libraries provided publicly by each of the authors. Using geometric averages, we annualize the risk-free rate provided by French and the excess returns for each anomaly.

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 $R_{t+1} = R_t^f + R^e_{anom,t+1}$, where $R^e_{anom,t+1}$ is the excess level return directly taken from the authors, and ensures that our measured $R^e_{t+1} \equiv R_{t+1}/R_t^f$ is simply $1 + R^e_{anom,t+1}/R_t^f$.

For the DHS behavioral anomalies, we only have returns data up to 2018 and therefore evaluate the performance of our SBF over this sample (1982-2018). For the Chen-Zimmermann anomalies, we consider all anomalies with returns over the sample of our SBF (1982-2022). This results in 176 anomalies assigned to 32 broader anomaly categories by the authors' original classification. Given that our goal is to calculate the average $\log\left(E\left[R^e_{j,t+1}\right]\right)$ and the average $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ across all anomalies j in each category, we reassign anomalies which belong to groups with three or fewer anomalies.[35] 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.

[35] We set the threshold at 3 or fewer anomalies as 5 of the 32 categories have exactly 4 anomalies and we want to limit the number of categories that are being reassigned.

Table AIV

Chen-Zimmermann Anomalies and Categories.

This table shows the broader set of Chen and Zimmermann (2022) anomalies which are available over the same sample as our log SBF $\hat{s}_{t+1}$, i.e., 1982 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

G. *Upper bound from correlated shocks across variables*

Table AV

Test of joint biases

This table compares the fit of the synthetic bias with the upper bound that could be derived from correlated shocks. Columns 1 and 5 show each of the 22 variables in the Survey of Professional Forecasters and Blue Chip, excluding RGDP and T-bill which are used to construct the synthetic biases. Columns 2 and 6 show the R_j^2 obtained from regressing each bias $E_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$ on its synthetic bias $\hat{E}_t^*[Z_{j,t+1}] - E_t[Z_{j,t+1}]$. Columns 3 and 7 show the upper bound derived in Proposition 5. Columns 4 and 8 report whether the R^2 violates the upper bound (Y) or does not (N).

$Z_{j,t+1}$	R^2	Upper bound	Violates bound (Y/N)	$Z_{j,t+1}$	R^2	Upper bound	Violates bound (Y/N)
rcon	0.75	0.78	N	*rexport*	0.29	0.10	Y
cpi	0.16	0.07	Y	*cprof*	0.36	0.33	Y
unemp	0.38	0.71	N	*ffr*	0.84	0.73	Y
indp	0.76	0.81	N	*libor*	0.82	0.69	Y
aaa	0.25	0.32	N	*6m*	0.85	0.74	Y
rnrinv	0.62	0.59	Y	*1y*	0.83	0.70	Y
rresinv	0.32	0.10	Y	*2y*	0.74	0.62	Y
rgf	0.35	0.04	Y	*5y*	0.53	0.45	Y
rgsl	0.18	0.21	N	*10y*	0.31	0.31	Y
housing	0.16	0.13	Y	*prime*	0.87	0.73	Y
rcbi	0.53	0.45	Y	*mortg*	0.33	0.32	Y

H. *Univariate benchmark for autocorrelation*

Similar to Figure 6, Figure A6 shows how well our synthetic expectations match the survey autocorrelation of forecast errors compared to univariate persistence-based models. As in Section IV.A, we estimate a linear relationship between forecast error autocorrelation and persistence,

$$AC_j = a + b\rho_j + \eta_j$$

where ρ_j is the objective persistence of the variable. The predicted forecast error autocorrelation is then $a + b\rho_j$. In the middle panel, we estimate a and b solely using expectations of RGDP and T-bill. While RGDP has lower persistence and lower forecast error autocorrelation than T-bill, the estimated a and b perform relatively poorly in predicting the forecast error autocorrelation for other variables. In the right panel, we estimate a and b using all

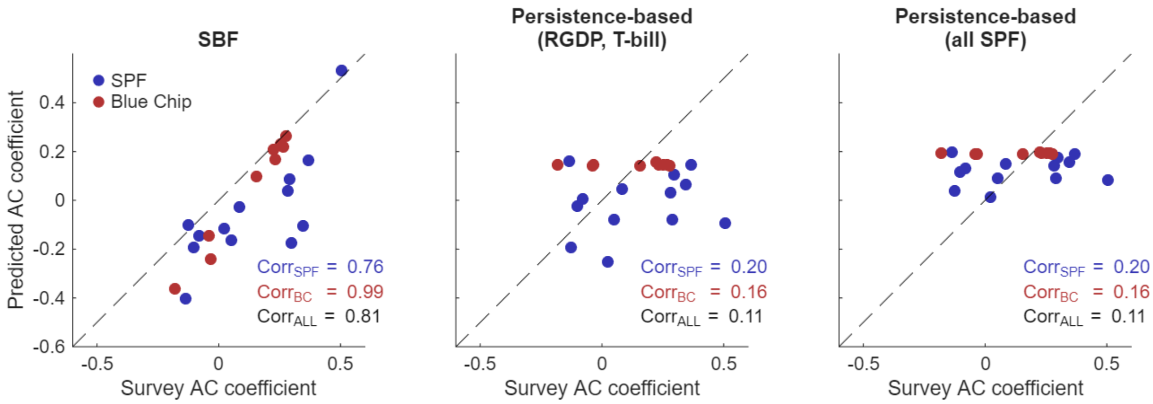

Figure A6. Comparison of the SBF and univariate persistence-based benchmarks for forecast error autocorrelation. This figure illustrates how synthetic expectations and univariate persistence-based models match the autocorrelation coefficients across the 22 Survey of Professional Forecasters and Blue Chip variables (excluding RGDP and T-bill). In all three panels, the horizontal axis is the actual forecast error autocorrelation for each variable and the vertical axis is the predicted value. The left panel shows the prediction from the SBF. The middle and right panels show predictions from a persistence-based model, $AC_j = a + b\rho_j + \eta_j$, where ρ_j is the objective persistence of each variable j. In the middle panel, the model is estimated using only RGDP and T-bill, and in the right panel, the estimation uses all 15 Survey of Professional Forecasters variables. Blue points denote Survey of Professional Forecasters variables and red points denote Blue Chip variables. The sample period is 1982Q4-2022Q2.

15 Survey of Professional Forecasters variables. Like Figure 6, this compresses the range of predicted forecast error autocorrelation, but does not noticeably improve the fit.

I. Return results using alternative SBFs

The figures below replicate the results of Figures 7-9 using ten alternative SBFs based on the top ten pairs of factors in Table IV.

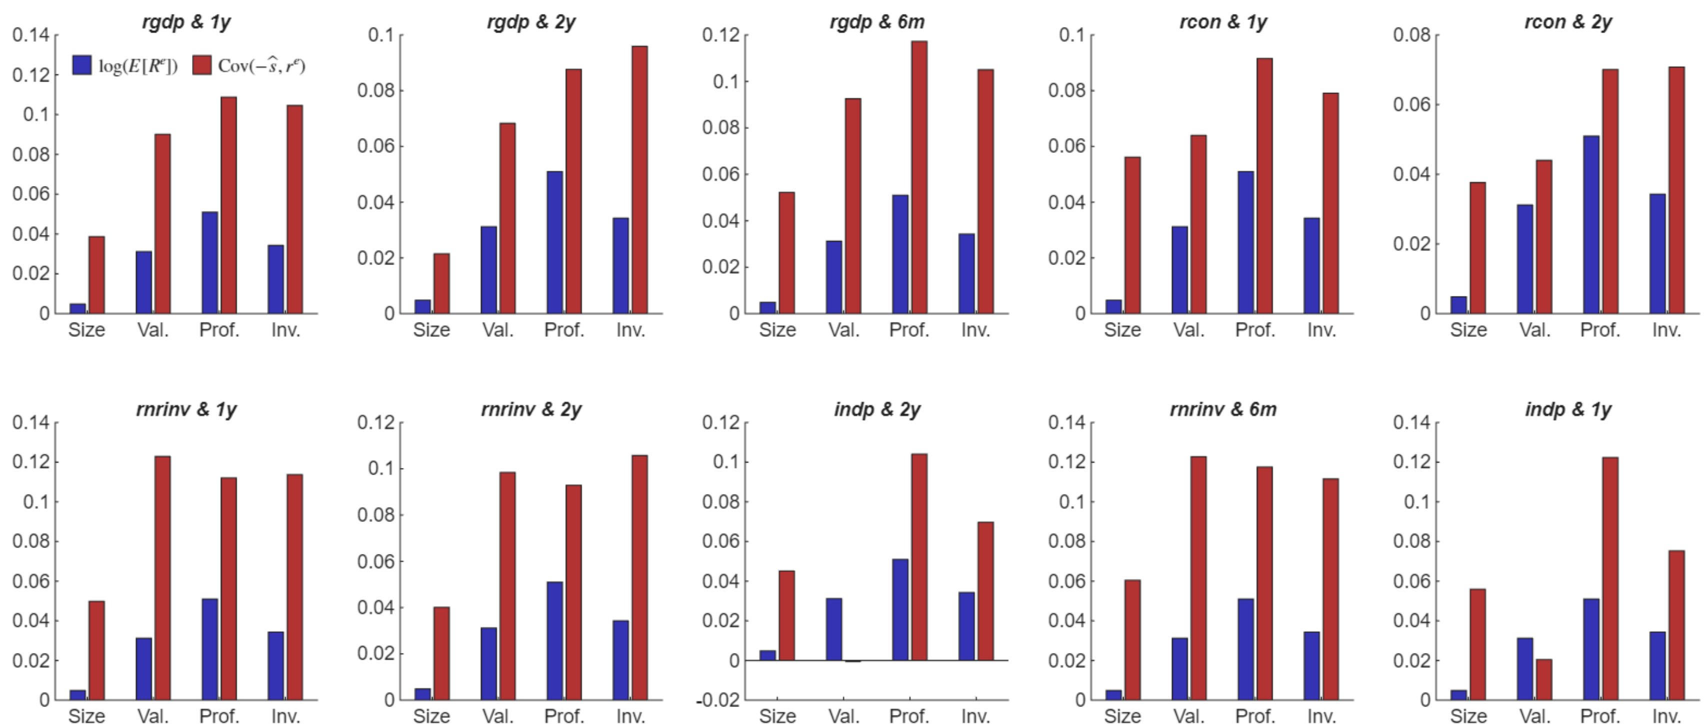

Figure A7. Fama-French factors and their covariance with the SBF using alternative pairs. This figure replicates Figure 7 using each of the ten alternative log SBFs constructed from the top ten variable pairs in Table IV. For each alternative SBF $\hat{s}_{t+1}$, this figure shows the decomposition (29) of excess returns for each of the Fama-French cross-sectional factors. The blue bars show the average excess return $\log\left(E\left[R^e_{j,t+1}\right]\right)$ 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, $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$.

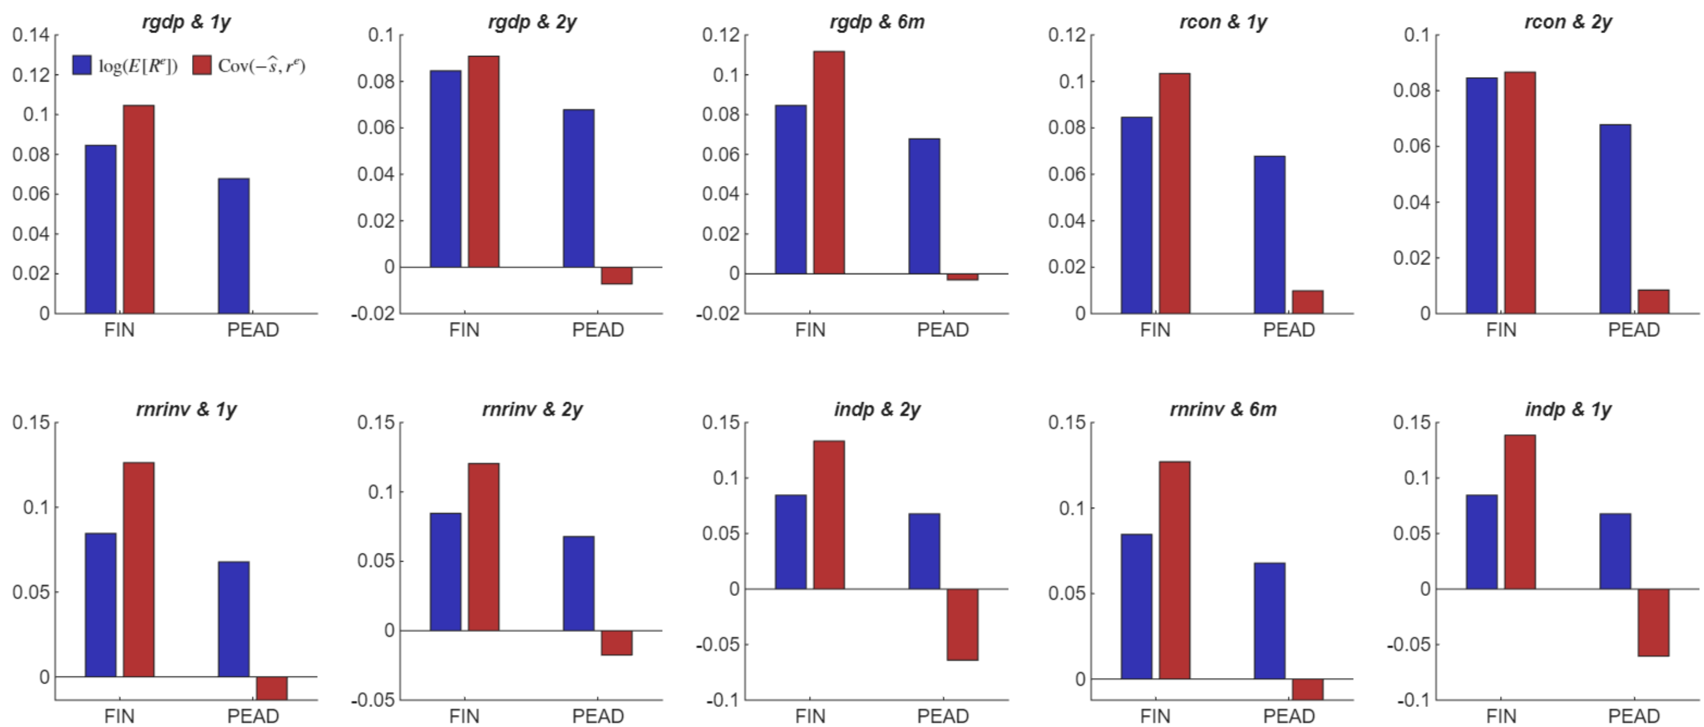

Figure A8. Behavioral factors and their covariance with the SBF using alternative pairs. This figure replicates Figure 8 using each of the ten alternative log SBFs constructed from the top ten variable pairs in Table IV. For each alternative SBF $\hat{s}_{t+1}$, this figure shows the decomposition (29) of excess returns for each of the behavioral factors from Daniel, Hirshleifer, and Sun (2020). The blue bars show the average excess return $\log\left(E\left[R^e_{j,t+1}\right]\right)$ of the FIN (Financing) and PEAD (Post-Earnings Announcement Drift) factors. The red bars show the portion of the average return attributable to the covariance of the excess return with the log SBF, $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$.

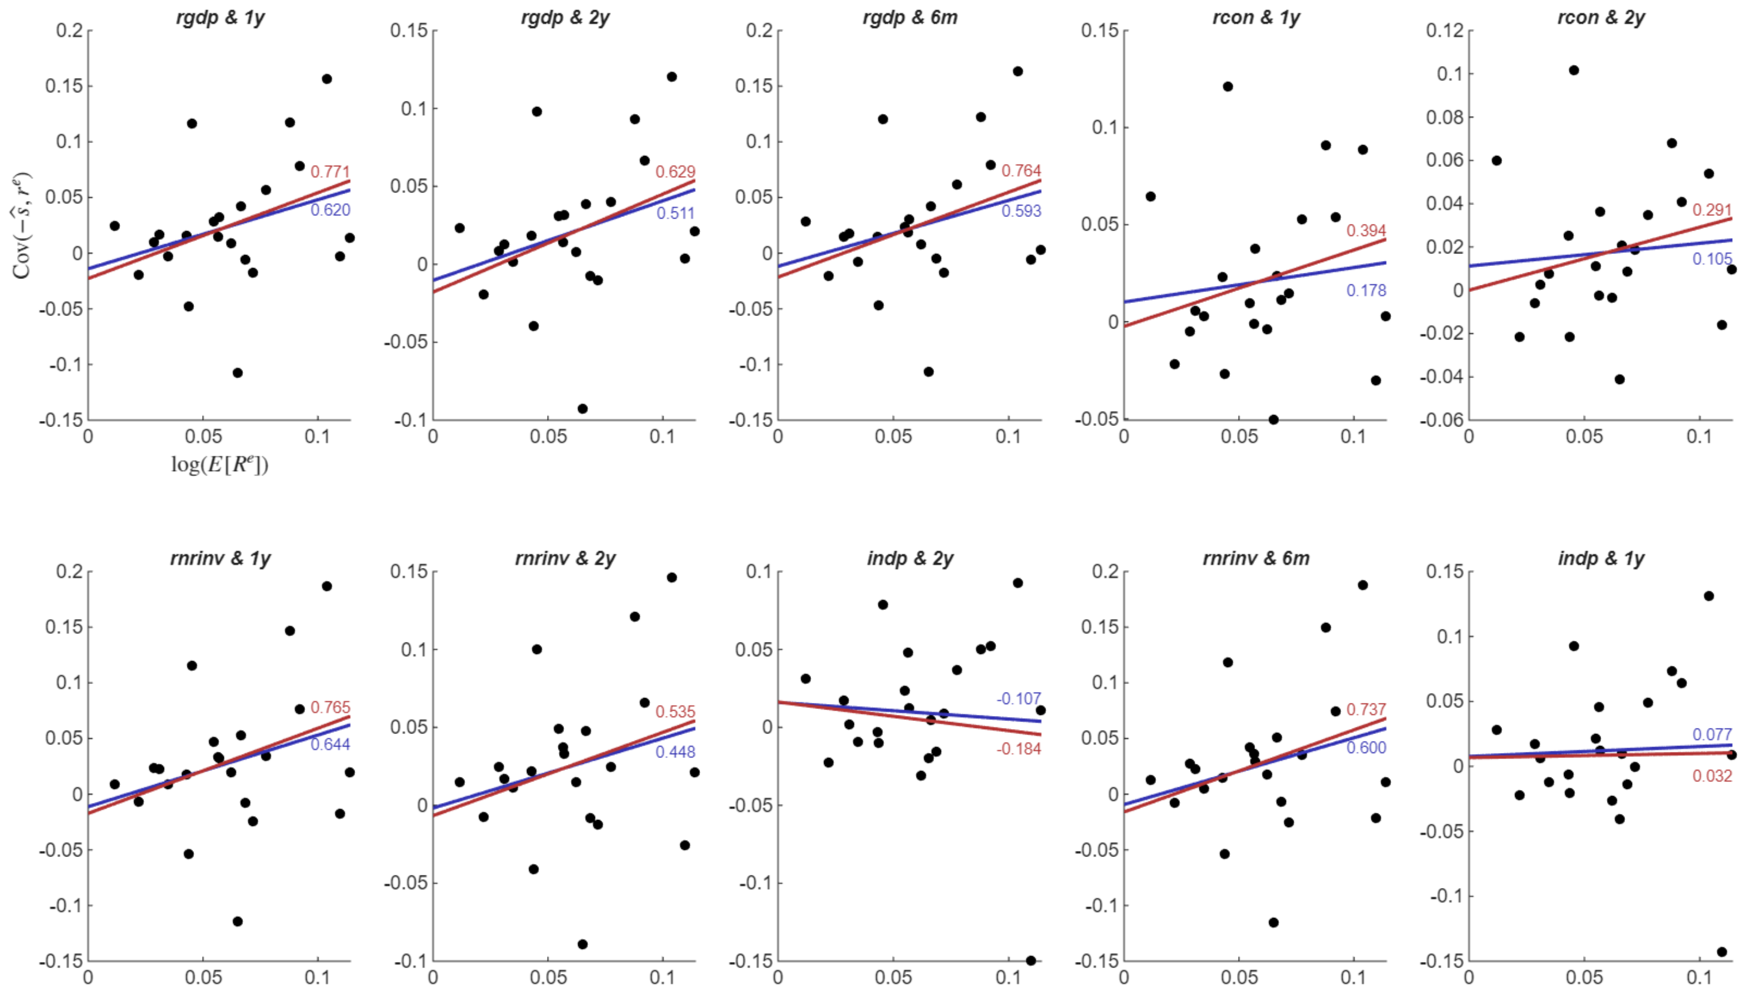

Figure A9. Anomalies and their covariance with the SBF using alternative pairs. This figure replicates Figure 9 using each of the ten alternative log SBFs constructed from the top ten variable pairs in Table IV. For each alternative SBF $\hat{s}_{t+1}$, this figure shows the decomposition of excess return for the Chen and Zimmermann (2022) anomalies. The x-axis shows the average excess return $\log\left(E\left[R^e_{j,t+1}\right]\right)$ of the anomaly portfolios. The y-axis shows $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$, 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 $Cov\left(-\hat{s}_{t+1}, r^e_{j,t+1}\right)$ on $\log\left(E\left[R^e_{j,t+1}\right]\right)$ as in equation (30). The line in red shows the slope of the same regression weighted by the number of anomalies in each category.

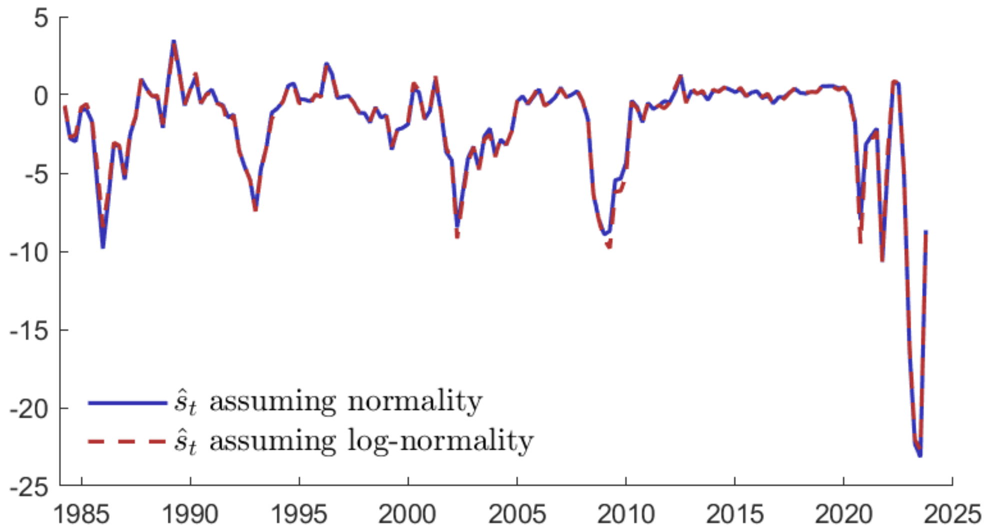

Figure A10. Estimated SBF. This figure compares the estimated SBF when we assume that the Survey of Professional Forecasters variables are objectively normally distributed and when we assume they are objectively log-normally distributed. The blue solid line shows the estimated SBF assuming normality, which is the SBF used in the paper. The red dashed line shows the estimated SBF assuming log-normality.

J. Using log-normal distributions rather than normal distributions

For our 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 A10 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). Below, we detail the process for creating synthetic expectations for log-normally distributed variables.

Of the 15 Survey of Professional Forecasters 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, 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 an autoregressive model for $x_{t+1} \equiv \log(X_{t+1})$,

$$x_{t+1} = a + B \begin{pmatrix} x_t & \log\left(E_t^*\left[X_{t+1}\right]\right) \end{pmatrix} + \varepsilon_{t+1} \tag{A16}$$

$$E_t\left[X_{t+1}\right] = \exp\left(E_t\left[x_{t+1}\right] + \frac{1}{2}\mathrm{diag}\left(\Sigma\right)\right) \tag{A17}$$

where Σ is the estimated covariance matrix of the shocks ε_{t+1}. We include the log of the survey expectations in equation (A16) to ensure that our statistical expectations contain any information known to the forecasters. 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 $E_t^*\left[\hat{X}_{t+1}\right]$,

$$\hat{s}_{t+1} \equiv -\frac{1}{2}\hat{\beta}_t'\hat{\Sigma}\hat{\beta}_t + \hat{\beta}_t'\hat{\varepsilon}_{t+1} \tag{A18}$$

where $\hat{\varepsilon}_{t+1}$ is the objective shocks to $\hat{x}_{t+1}$, $\hat{\Sigma}$ is the covariance matrix of $\hat{\varepsilon}_{t+1}$, and $\hat{\beta}_t = \hat{\Sigma}^{-1}\left(\log\left(E_t^*\left[\hat{X}_{t+1}\right]\right) - \log\left(E_t\left[\hat{X}_{t+1}\right]\right)\right)$ following Proposition 3. Then, we can estimate synthetic expectations based on $\hat{s}_{t+1}$ for the remaining variables as

$$\begin{aligned} \hat{E}_t^*\left[X_{t+1}\right] &\equiv E_t\left[\hat{S}_{t+1}X_{t+1}\right] \\ &= E_t\left[X_{t+1}\right]\exp\left(Cov\left(\varepsilon_{t+1}, \hat{\varepsilon}_{t+1}\right)\hat{\beta}_t\right). \end{aligned} \tag{A19}$$

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 $\hat{E}_t^*\left[Y_{t+1}\right]$ solely using realized data and the Survey of Professional Forecasters survey data. First, we estimate an autoregressive model for $y_{t+1} \equiv \log(Y_{t+1})$,

$$y_{t+1} = a_y + B_y \begin{pmatrix} y_t & x_t & \log\left(E_t^*\left[X_{t+1}\right]\right) \end{pmatrix} + \varepsilon_{y,t+1} \tag{A20}$$

$$E_t\left[Y_{t+1}\right] = \exp\left(E_t\left[y_{t+1}\right] + \frac{1}{2}\mathrm{diag}\left(\Sigma_y\right)\right) \tag{A21}$$

where Σ_y is the estimated covariance matrix of $\varepsilon_{y,t+1}$. To ensure our statistical expectations contain as much current information as possible, we include the current value x_t and the survey expectations $\log\left(E_t^*\left[X_{t+1}\right]\right)$ for the Survey of Professional Forecasters variables in equation (A20). Then, using our log SBF $\hat{s}_{t+1}$ estimated from the Survey of Professional Forecasters RGDP growth and T-bill rate expectations, we can calculate our synthetic expectations for Y_{t+1} as

$$\begin{aligned} \hat{E}_t^*\left[Y_{t+1}\right] &\equiv E_t\left[\hat{S}_{t+1}Y_{t+1}\right] \\ &= E_t\left[Y_{t+1}\right]\exp\left(Cov\left(\varepsilon_{y,t+1},\hat{\varepsilon}_{t+1}\right)\hat{\beta}_t\right). \end{aligned} \tag{A22}$$

Importantly, the construction of $\hat{E}_t^*\left[Y_{t+1}\right]$ does not require any survey expectations of Y_{t+1}.