

Long-Run Inflation Expectations*

Jonas D. M. Fisher[†] Leonardo Melosi[‡] Sebastian Rast[§]

March 2025

Abstract

Professional forecasters' long-run inflation expectations overreact to news and exhibit persistent, predictable biases in forecast errors. A model incorporating overconfidence in private information and a persistent expectations bias—which generates persistent forecast errors across most forecasters—accounts for these two features of the data, offering a valuable tool for studying long-run inflation expectations. Our analysis highlights substantial, time-varying heterogeneity in forecasters' responses to public information, with sensitivity declining across all forecasters when monetary policy is constrained by the effective lower bound. The model provides a framework to evaluate whether policymakers' communicated inflation paths are consistent with anchored long-run expectations.

Keywords: *Panel survey data, long-run inflation expectations, rationality, expectation bias, overconfidence, overreaction, central bank communications, anchoring.*

JEL classification: *E31, D83, E52, E37*

*Emails: jonas.fisher@chi.frb.org, leonardo.melosi@eui.eu, and s.g.rast@dnb.nl. We thank Fabio Canova, Todd Clark, Yuriy Gorodnichenko, Emi Nakamura, Barbara Rossi, Tom Stark, Allan Timmermann, Mark Watson, and Mirko Wiederholt for their helpful comments. We also thank participants at the 2022 NBER Summer Institute, NBER Conference on Inflation Expectations, the Federal Reserve Bank of San Francisco's 2023 Macroeconomics and Monetary Policy Conference, the Federal Reserve Bank of St.Louis' Information and Expectations in Macroeconomics Workshop, the 2024 SED Annual Meeting, the University of Mannheim-EABCN conference on New Challenges in Monetary Economics 2024, the 2022 IAAE Annual Conference, the EEA-ESEM congress 2022, the Midwest Macro Meetings, the 2022 Dolomiti Macro Meetings, the 49th OeNB Economic Conference, the 2024 Annual Meeting of the ECSB Research Cluster on Monetary Economics, the 35th SUERF Colloquium, and at seminars held at the European University Institute, Rice University, LMU of Munich, the Bank of Italy, the Bundesbank, the Danmarks Nationalbank, De Nederlandsche Bank, and the Norges Bank for their helpful comments and questions. Any views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System, De Nederlandsche Bank or the European System of Central Banks.

[†]FRB Chicago.

[‡]University of Warwick, European University Institute, De Nederlandsche Bank, and CEPR.

[§]De Nederlandsche Bank.

1 Introduction

Understanding how long-run inflation expectations are formed is essential for modern monetary policy theory and practice. In this paper, we develop a flexible framework for analyzing how U.S. professional forecasters form expectations about long-run inflation. A model featuring overconfidence in private information and persistent expectations bias that engenders highly persistent forecast errors explains the key time-series and cross-sectional features of the data. Our analysis offers new insights into the types of cognitive distortions that shape the formation of beliefs about long-run price dynamics. Moreover, the model allows the estimation of the heterogeneous responsiveness of long-run inflation expectations to public and private information. Finally, we show how the model can identify short-term inflation trajectories that are consistent with anchoring of long-run expectations, a central goal for monetary policy.

We identify two key challenges in explaining the combined cross-sectional and time-series dynamics of long-run inflation expectations. First, forecast errors exhibit a significant degree of serial correlation, substantially higher than that observed for short-term inflation expectations. Second, long-run expectations significantly overreact to news, in contrast to little to no overreaction observed in short-term inflation expectations. These findings highlight the unique complexities involved in modeling long-run inflation expectations, setting them apart from their short-run counterparts.

We posit that forecasters solve a dynamic signal extraction problem, taking as given that inflation comes from a time-varying-parameter, trend-cycle time-series model. We build on the standard noisy information framework, e.g., Sims (2003), Woodford (2003), Maćkowiak and Wiederholt (2009).¹ Each forecaster observes three sources of information about the trend component of inflation: the quarterly inflation rate, a *coordinating signal*, and an *idiosyncratic signal*. Including inflation as a signal enables us to estimate the sensitivity of long-run inflation expectations to changes in short-run inflation, a measure of anchoring proposed by Bernanke (2007). It also allows forecasters to estimate the parameters of the trend-cycle model in real time. The coordinating signals—which are perfectly correlated across forecasters—explain co-movement in forecasters’ expectations. We interpret these signals as representing forward-looking public information, including policymakers’ communications, widely followed media, and any events unrelated to past inflation that happen to align forecasters’ beliefs about the central bank’s commitment to price stability.² The idiosyncratic signals capture

¹Coibion and Gorodnichenko (2012) show that noisy information models match key empirical facts regarding short-run expectations of professional forecasters, consumers, firms, and central banks.

²Nimark (2014), Chahrour, Nimark and Pitschner (2021) examine how the media industry selects signals observed by economic agents and the implications of this selection for business cycles.

individual judgments or private information to help account for the heterogeneity in forecasters' expectations.

A model of long-run inflation expectations should incorporate the ability to capture nonstationary dynamics, including periods of drift. Trend-cycle models allow drift by design. Historical episodes, such as the sustained inflation surge of the 1970s, the long period of inflation below the Fed's target in the 2010s, or the resurgence of inflation following the pandemic, highlight the importance of capturing nonstationary dynamics. Additionally, the long-run inflation expectations of many forecasters exhibit little evidence of mean reversion, underscoring the need for models that can accommodate extended deviations from stability. By contrast, most models of short-term inflation expectations are built on the assumption of mean reversion so that inflation expectations are always anchored. Our trend-cycle model also has time-varying parameters to capture structural breaks that may have influenced inflation over the long sample period required to accurately estimate the trend.

We examine the model under different assumptions about the forecasters' cognitive abilities. This enables us to test for potential deviations from rationality in the way professional forecasters form their expectations. Specifically, we begin by assuming that forecasters have rational expectations and then evaluate whether this assumption is consistent with the data. Under rational expectations, forecasters are presumed to know the time-varying parameters of the trend-cycle model at each point in time, as well as the precision of each signal in conveying information about trend inflation. However, with our signal structure, they are imperfectly informed, as they cannot fully disentangle the drivers of inflation within the trend-cycle model from the signals they observe.

We estimate the model in two steps. First, we estimate the time-varying parameters of our trend-cycle model using quarterly inflation data from the first quarter of 1959 (1959Q1) to 2023Q2. In the second step, we assume the forecasters in our panel data (that begins in 1991Q3) are given the parameters of the trend-cycle model and then estimate the parameters of the signal extraction problem they face. The second-step estimation is performed using the time series of inflation, the two-sided (smoothed) estimates of the trend obtained from the first step, and the long-run expectations of the panel of forecasters. This two-step estimation approach offers a significant advantage by facilitating identification of the sensitivity of expectations to coordinating and idiosyncratic signals, which are unobservable to the econometrician.

The model with rational forecasters does not fit the data well and reveals two major sources of misspecification. First, the in-sample estimates of the i.i.d. innovations to the coordinating signals exhibit serial correlation. This deviation from the white-noise assumption underscores the model's inability to account for the persistent forecast errors characterizing the long-run inflation expectations. Second, the in-sample estimates of the innovations to the

idiosyncratic signals display excess volatility, skewness, and kurtosis, highlighting the model’s failure to explain the overreaction in long-run inflation expectations.

We then relax the assumption of rationality to investigate the behavioral features a theory of long-run inflation expectations should include to explain the data. To this end, we allow for the possibility that forecasters missperceive certain parameters of their signal extraction problem, introducing the possibility of cognitive biases. Crucially, we do not impose any specific biases *a priori*, nor do we target any particular misspecification of the rational model. Instead, the likelihood estimation of the model identifies the deviations necessary to correct the misspecification of the rational model, which is a special case of the behavioral model.

Relaxing the rationality assumptions allows the model to overcome the aforementioned misspecification, revealing two major deviations from rationality. First, in the estimated behavioral model, the coordinating signal leads to highly persistent forecast errors, a phenomenon we term *persistent expectations bias*. This cognitive bias allows the behavioral model to account for the highly persistent forecast errors characterizing the long-run inflation expectations in the data. Second, in the estimated model, forecasters overestimate the precision of their idiosyncratic signals. *Overconfidence in private information* is essential for the model to explain overreaction in long-run inflation expectations. Notably, we find that almost all professional forecasters in our sample are influenced by these two cognitive biases.

In the behavioral model, forecasters’ expectations about long-run inflation exhibit minimal sensitivity to inflation, which implies that transitory shocks to inflation, such as cyclical and i.i.d. shocks in the trend-cycle model, have negligible effects on forecasters’ long-run inflation expectations. This result suggests that temporary deviations of inflation from its target do not undermine professional forecasters’ confidence in the central bank’s ability to stabilize inflation over the long run.³

In contrast, the sensitivity of long-run inflation expectations to the coordinating signal is relatively large, on average, and varies significantly across forecasters. During periods when the federal funds rate is constrained by the effective lower bound (ELB), the median sensitivity declines dramatically. This suggests that the policy rate is a valuable communications tool for a central bank to shape long-run inflation expectations. At the ELB, the central bank loses this tool. This result is also consistent with theoretical insights from the literature on models with dispersed information, dating back to Morris and Shin (2002) and Woodford (2003), who demonstrate that public signals serve as focal points for coordinating dispersed expectations about the economy’s fundamentals.

³This property aligns closely with Bernanke (2007)’s definition of anchored inflation expectations. This concept of anchoring has been tested extensively in several empirical studies, e.g. Gürkaynak, Levin, Marder and Swanson (2007), Dräger and Lamla (2014), Corsello, Neri and Tagliabracchi (2021), Barlevy, Fisher and Tysinger (2021), Armantier, Sbordon, Topa, Van der Klaauw and Williams (2022).

Our approach leverages the entire distribution of forecasters’ long-run inflation expectations to estimate their sensitivity to public or private sources of information. By estimating heterogeneous sensitivities, we control for compositional effects when we measure average expectations. Accounting for compositional effects is particularly important in light of the critique of conditional mean forecasts highlighted by Engelberg, Manski and Williams (2010).

Finally, our model provides a framework to assess whether the inflation paths communicated by policymakers align with their objective of maintaining anchored long-run inflation expectations. To illustrate how this analysis can be carried out, we use the estimated behavioral model to predict the inflation path consistent with stable long-run inflation expectations from the fourth quarter of 2022 onward, a period when trend inflation reached one of its highest levels in the U.S. over the past 30 years. Despite this, mean long-run CPI inflation expectations remained close to target. According to our model, the inflation path consistent with stable long-run expectations closely mirrors the rapid decline in U.S. inflation observed during the first four quarters of 2023. Intriguingly, the *anchoring-compatible inflation path* also aligns closely with the inflation projections released by the Federal Reserve’s Federal Open Market Committee (FOMC) in December 2022 for the subsequent three years.

The behavioral model still rests on strong assumptions, such as forecasters agreeing on the trend-cycle model being the correct underlying model of inflation. This assumption is in practice less stringent than it may appear, since the model features innovations that can have persistent effects on the idiosyncratic signals. The estimated degree of persistence is forecaster-specific and can be large (potentially near permanent), allowing our model to capture long-lasting disagreement about the process driving inflation or about judgment calls concerning long-run inflation outcomes.⁴ We also study a version of the model in which forecasters estimate the parameters of the trend-cycle model in real time. This change to the rational model alone would not resolve its two key sources of misspecification.

While the model considered in this paper could, in principle, be estimated using data on households’ or firms’ expectations, focusing on professional forecasters is more convenient. Since professional forecasters do not influence price-setting in the economy directly, we can avoid modeling feedback from expectations to the underlying inflation dynamics. Incorporating such feedback would be complex and could obscure the primary objective of this paper: understanding how long-run inflation expectations are formed.

⁴Indeed, the latter seems to be more likely, as Stark (2013) documents that forecasters use both models and subjective considerations in reporting their projections, rather than just models. Moreover, fewer forecasters rely on models to forecast at long horizons (three or more years out) than at short horizons. In the Survey of Professional Forecasters administered by the European Central Bank, only 10 percent of forecasters rely completely on statistical models to forecast long-run inflation expectations (see de Vincent-Humphreys, Dimitrova, Falck, Henkel and Meyler (2019)).

Understanding how professional forecasters, as opposed to households or firms, form long-run inflation expectations is crucial, even though these agents are not directly involved in price setting. As attentive observers of central bank actions and communications, their expectations can serve as early indicators of inflation expectations becoming unanchored more broadly. This underscores why scholars and policymakers devote significant time and resources to studying these expectations.⁵

Previous research has shown that forecasters may strategically misreport their information or exhibit herd behavior for various reasons (Ehrbeck and Waldmann, 1996). If such strategic behaviors contribute to the coordination of inflation expectations, they would be captured by the coordinating signals in our model. However, the strategic behaviors documented in this literature primarily pertain to short-term expectations and may not necessarily apply to longer-run expectations. Additionally, as it takes 10 years for the forecast error to be realized, some of the incentives to misreport identified in these studies may not be relevant for our analysis.

2 Relation to the literature

A large and growing literature documents that professional forecasters' short-run expectations about a wide range of macroeconomic variables violate the Full Information Rational Expectations (FIRE) assumption. Notable examples include Coibion and Gorodnichenko (2015), Bordalo, Gennaioli, Ma and Shleifer (2020), Kohlhas and Walther (2021), Bianchi, Ilut and Saijo (2023), Kučinskas and Peters (2022), Rossi and Sekhposyan (2016), Afrouzi, Kwon, Landier, Ma and Thesmar (2023), and Born, Enders and Müller (2023). Farmer, Nakamura and Steinsson (2023) show that imperfect knowledge of the forecasting model, particularly its long-run properties, and reasonable initial beliefs can account for several anomalies in short-term consensus expectations about GDP growth, nominal interest rates, and deviations from the expectations hypothesis of the term structure of interest rates. However, they do not examine how professional forecasters form expectations about long-run inflation or estimate their model using cross-sectional information. Angeletos, Huo and Sastry (2021) show that a framework where agents misperceive the precision or persistence of certain variable movements explains why U.S. professional forecasters and households initially underreact and later overreact to business cycle shocks. While this literature has largely focused on short-run expectations, our focus is on long-run inflation expectations, which are

⁵For example, the Federal Reserve has developed the Index of Common Inflation Expectations, which captures the co-movement of various inflation expectations measures, including the professional forecasters' expectations used in this paper (Ahn and Fulton, 2021). Academic studies examining different aspects of professional forecasters' expectations are discussed below.

central to modern monetary policy theory and practice.

Bordalo, Gennaioli, Ma and Shleifer (2020) apply the methodology introduced by Coibion and Gorodnichenko (2015) to individual forecaster data from the U.S. *Survey of Professional Forecasters* (SPF) and document that overreaction is a pervasive feature of short-run macroeconomic expectations. We employ their methodology (as best we can) to show that long-run inflation expectations display a large overreaction. Our noisy information model attributes this overreaction to forecasters’ overconfidence in private information.

Broer and Kohlhas (2022) provide evidence of overreaction in SPF expectations for the GNP/GDP deflator over the short run (up to six quarters ahead) and show that this phenomenon cannot be explained by a stylized noisy information model with rational expectations. Our study is complementary to theirs. While we find that overconfidence alone is sufficient to explain overreaction in long-run inflation expectations, Broer and Kohlhas (2022) show that a broad notion of overconfidence and the presence of endogenous public signals is necessary for the noisy information model to explain short-run inflation expectations. We also show that overconfidence alone cannot successfully explain persistent forecast errors, a characteristic of long-run inflation expectations that is not shared with short-term expectations.

Similarly, Adam, Kuang and Xie (2024) use a noisy information model to show that biases consistent with overconfidence in private information are pervasive in short-term SPF forecasts. Bianchi et al. (2022) use a machine-learning algorithm to build an appropriate benchmark for quantifying biases in survey responses. They find that survey respondents typically place too much weight on the private or judgmental component of their forecasts and too little weight on objective, publicly available economic information. This finding echoes our results regarding overconfidence in private information.

An alternative explanation for overreaction, which has been extensively documented in the context of *short-term* expectations, arises from models with diagnostic expectations (Bordalo, Gennaioli, Ma and Shleifer, 2020).⁶ Our analysis is conducted within the class of models with noisy information, and evaluating whether a model with diagnostic expectations can explain the observed long-run inflation expectations lies beyond the scope of this paper.

Starting with Mankiw, Reis and Wolfers (2004), a growing body of literature has used individual forecasters’ expectations to explain how the cross-section of expectations evolves over time, across forecast horizons, and across countries (e.g. Patton and Timmermann (2010), Andrade and Bihan (2013), Andrade, Crump, Eusepi and Moench (2016), and Dovern, Fritsche and Slacalek (2012)), also finding evidence of persistent forecast disagreement among U.S. professional forecasters. More recently, Goldstein and Gorodnichenko (2022) develop

⁶Diagnostic expectations refer to a belief formation process in which individuals overemphasize recent or salient information when making forecasts or predictions.

a model where forecasters receive noisy signals about the future and estimate it using the cross-sectional distribution of the SPF short-term forecasts of CPI inflation. They find that forward-looking signals are crucial to explaining the cross-sectional dispersion in short-term forecasts. In our model, signals also contain forward-looking information, as they pertain to trend inflation, which is estimated using a two-sided procedure that conditions on past and future inflation. Consequently, in every period, forecasters in our model receive *future* information about low-frequency movements in inflation.

Giacomini, Skreta and Turen (2020) use professional forecasts on U.S. inflation in the 18 months prior to the inflation release and present evidence that forecast disagreement persists as uncertainty resolves. In their theory forecasters are rational and update their expectations using Bayes’ rule and disagreement stems from heterogeneous priors about the initial forecast, heterogeneous models to interpret public information, and heterogeneous inattention. Goldstein (2023) presents evidence of persistent forecast disagreement among U.S. professional forecasters, highlighting variation over time, across individuals, and across forecast horizons – particularly for average 10-year inflation expectations, where persistence is especially pronounced. His analysis primarily attributes the persistence to differences in the degree of inattention among rational forecasters. In contrast, we introduce behavioral biases into an imperfect information model and show that two key biases enable the model to explain the significant overreactions and the persistent, predictable forecast errors that characterize professional forecasters’ long-run inflation expectations.

Kozicki and Tinsley (2001) were among the first to estimate the term structure of professional forecasters’ inflation expectations by using reduced-form models. Aruoba (2020) estimates a model of the term structure of expectations featuring three factors (level, slope, and curvature) using survey data. Crump, Eusepi, Moench and Preston (2023) document the behavior of the term structure of expectations of GDP growth, inflation, and the policy rate using multiple surveys of professional forecasters for the U.S. and fit this rich data set with a multivariate reduced-form model. Herbst and Winkler (2021) and Ahn and Farmer (2024) study dynamic factor models of the term structure of disagreement among inflation forecasts using the SPF micro data. Neither of these two studies evaluates the rationality of professional forecasters’ beliefs about long-run inflation.

Our trend-cycle model builds on Stock and Watson (2007) and Chan, Clark and Koop (2018) and is closely related to the work of Mertens and Nason (2020), Cogley, Primiceri and Sargent (2010), and Hasenzagl, Pellegrino, Reichlin and Ricco (2020).⁷ Papers estimating trend-cycle models most similar to ours link trend inflation to survey-based expectations to explore the implications of changes in the inflation process for inflation forecasts and the

⁷Faust and Wright (2013) review the earlier literature on trend-cycle models.

anchoring of expectations. Henzel (2013), Mertens and Nason (2020), and Nason and Smith (2021) focus on average short-run inflation forecasts from the SPF, treating trend inflation as the long-run forecast of inflation, though long-run inflation expectations themselves are not explicitly included as observables in their estimations. Mertens (2016) estimates Beveridge-Nelson decompositions to obtain trend inflation estimates, incorporating information from inflation, survey forecasts, and long-term interest rates. These studies focus primarily on the time series dynamics of average inflation expectations, with limited focus on the cross-sectional dimension of the data.

Our paper is also connected to the literature on the anchoring of inflation expectations. Influential contributions in this area include Orphanides and Williams (2005), Beechey, Johannsen and Levin (2011), and Carvalho, Eusepi, Moench and Preston (2023). These studies examine signal extraction problems where agents attempt to infer the central bank’s inflation target based on past data. However, by focusing on a representative agent, they do not incorporate the cross-sectional information that is central to our analysis. While those papers define anchoring as the stability of average long-horizon inflation forecasts around the central bank’s inflation target, other definitions of anchoring have also been proposed in the literature.

Kurmar, Afrouzi, Coibion and Gorodnichenko (2015) offer a taxonomy of definitions of expectations anchoring and evaluate them for New Zealand using survey data. By showing how to estimate the dynamics of expectations’ sensitivity to coordinating signals, we offer a new way to assess how anchored expectations are in a panel of professional forecasters. A novel feature of our analysis is the presence of nonstationary dynamics – specifically, trend inflation that can become ingrained in inflation expectations. This allows us to assess the risk of de-anchoring in real time and identify appropriate communications to mitigate it.

Reis (2022) relates inflation anchoring to changes in the cross-sectional variance and skewness of survey measures of inflation expectations. Grishchenko, Mouabbi and Renne (2019) use a trend-cycle model with time-varying volatility to relate anchoring to the probability of future inflation being in a certain range of the inflation target as measured by survey expectations. Binder, Janson and Verbrugge (2023) develop a measure to assess the degree of anchoring of expectations in a panel of professional forecasters. Their measure takes into account the coordination of expectations at the individual level. Consistent with our estimated sensitivity of expectations to public information at and away from the ELB, they provide evidence that the increase in the federal funds rate that ended the first ELB period enhanced the anchoring of long-run inflation expectations in the SPF.

3 A model of long-run inflation expectations

The model consists of a panel of forecasters aiming to predict long-run inflation while understanding that inflation is generated by a time-varying-parameter, trend-cycle time-series model. Forecasters observe inflation (the *inflation signal*) and receive two additional signals about trend inflation—a *coordinating signal* and an *idiosyncratic signal*—which they use to form their long-run expectations. Forecasters update their expectations about trend inflation by minimizing the variance of their estimates of the underlying state variables that govern the dynamics of inflation. Since our model is linear and its shocks are normally distributed, it is optimal for forecasters to use the Kalman filter to update their expectations.

We consider two alternative assumptions about the cognitive abilities of forecasters. In one case, we assume that forecasters understand all the parameters of the model, which we refer to as the *rational model*. In the other case, we assume that forecasters may misperceive key parameters of the signals, allowing for cognitive biases to arise. We refer to this case as the *behavioral model*. We will explore more types of deviations from rationality in subsequent analysis (Section 6).

3.1 The inflation generating process

Inflation, π_t , is generated by the following time-varying-parameter trend-cycle model:

$$\pi_t = \bar{\pi}_t + \psi_t + \sigma_\omega \omega_t; \quad (1)$$

$$\bar{\pi}_t = \bar{\pi}_{t-1} + \sigma_{\lambda,t} \lambda_t; \quad (2)$$

$$\psi_t = \phi_t \psi_{t-1} + \sigma_{\eta,t} \eta_t. \quad (3)$$

Inflation is decomposed into a trend component ($\bar{\pi}_t$), a cyclical component (ψ_t), and an i.i.d. component (ω_t). Trend inflation reflects the long-run drivers of inflation. The cyclical component captures persistent variations of inflation around its long-term trend, for example, due to Phillips curve dynamics. The i.i.d. component captures high-frequency variation in inflation that does not have persistent effects, for example, due to food and energy prices. The random variables ω_t , λ_t , and η_t are i.i.d. standard normal, and σ_ω , $\sigma_{\lambda,t}$, and $\sigma_{\eta,t}$ are strictly positive. The variances of the innovations to the trend and cycle components follow log random walk processes:

$$\ln(\sigma_{\lambda,t}^2) = \ln(\sigma_{\lambda,t-1}^2) + \gamma_\lambda \varepsilon_{\lambda,t}, \quad (4)$$

$$\ln(\sigma_{\eta,t}^2) = \ln(\sigma_{\eta,t-1}^2) + \gamma_\eta \varepsilon_{\eta,t}; \quad (5)$$

where $\varepsilon_{\eta,t}$ and $\varepsilon_{\lambda,t}$ are i.i.d. standard normal and γ_λ and γ_η are both strictly positive. The cyclical component's auto-regressive parameter is also stochastic and is modeled similarly:

$$\phi_t = \phi_{t-1} + \gamma_\phi \varepsilon_{\phi,t}, \quad (6)$$

where the $\varepsilon_{\phi,t}$ are drawn from truncated standard normal distributions with thresholds $-\phi_{t-1}/\gamma_\phi$ and $(1 - \phi_{t-1})/\gamma_\phi$ so that ϕ_t is stationary.

Our trend-cycle model is univariate. Stock and Watson (2016) show that univariate estimates of the trend in core inflation are nearly as accurate as multivariate estimates.⁸ Employing a multivariate trend-cycle model would significantly increase the already substantial computational burden of solving the forecasters' signal extraction problem. In Appendix C, Table 3, we provide evidence on the forecast performance of the above model and in particular, highlight the importance of allowing for time-varying parameters and persistence in the cyclical component.

3.2 Forecasters' information

Forecasters base their long-run inflation forecasts on limited information about the three drivers of inflation ($\bar{\pi}_t$, ψ_t , and ω_t). In both the rational and behavioral models, they are assumed to know the trend-cycle model described by equations (1)–(6) and the values of its parameters in every period. We will relax this assumption later in the context of the rational model. Additionally, forecasters observe inflation, a *coordinating signal*, and an *idiosyncratic signal*, which shape their beliefs about long-run inflation. Inflation is a public signal observed by all the forecasters, implying that today's long-run inflation expectations are conditional on current and past inflation. We interpret the coordinating signal as representing forward-looking public information, potentially due to policymakers' communications, widely followed media, and any events unrelated to past inflation that happen to align forecasters' beliefs about the central bank's commitment to price stability.⁹ In contrast, we interpret the idiosyncratic signals as forward-looking private information held by forecasters. Examples include forecasters' personal judgment about the central bank's credibility, their interpretation of factors driving inflation, and their individual experiences with past inflation episodes.¹⁰

⁸Crump et al. (2023) report that a multivariate trend-cycle model with time-varying parameters offers modest improvements over a univariate model for forecasting inflation at long horizons.

⁹Nimark (2014), Chahrour, Nimark and Pitschner (2021) examine how the media industry selects signals observed by economic agents and the implications of this selection for business cycles.

¹⁰Rich and Tracy (2017) find that forecasters tend to revise their forecasts toward the median forecast of the previous period in the European Central Bank's Survey of Professional Forecasts. Fuhrer (2018) finds similar evidence for short-run macroeconomic expectations of professional forecasters, firms, and households, both in the United States and in the euro area. While this is true for short-run expectations, it's not clear

The coordinating signals are perfectly correlated across forecasters and play the critical role of coordinating their expectations about trend inflation. The coordinating signal received by forecaster i in period t , $\tilde{s}_t(i)$, reads:

$$\tilde{s}_t(i) = \bar{\pi}_t + \alpha(i)v_{c,t}; \quad (7)$$

$$v_{c,t} = \rho_c v_{c,t-1} + \sigma_{c,t} \nu_{c,t}, \quad (8)$$

where $\alpha(i)$ and $\sigma_{c,t}$ are strictly positive and $v_{c,t}$ represents the realized noise in the coordinating signal with volatility of coordinating innovations $\sigma_{c,t}$. All forecasters' coordinating signals are affected by the same realizations of noise. The noise process is autoregressive with standard normal innovations, $\nu_{c,t}$. The variance of the innovations is time-varying and follows a random walk in logs: $\ln \sigma_{c,t}^2 = \ln \sigma_{c,t-1}^2 + \gamma_c \epsilon_{c,t}$, with standard normal innovations and with γ_c strictly positive.¹¹

The volatility of coordinating innovations, $\sigma_{c,t}$, is inversely related to the average precision of the coordinating signals across forecasters. In particular, a larger value of $\sigma_{c,t}$ makes the coordinating signal less precise. All else being equal, lower precision implies lower sensitivity of all forecasters' expectations to the coordinating signal and can, therefore, be a measure of forecasters' *average* sensitivity to the coordinating signals.¹²

The parameter $\alpha(i)$ captures heterogeneity in forecasters' sensitivity to coordinating signals. Forecasters with larger $\alpha(i)$ receive a relatively less precise coordinating signal, making their expectations less responsive to them. In the special case where $\alpha(i)$ is identical across all forecasters, the coordinating signals, $\tilde{s}_t(i)$, effectively become public signals. Furthermore, since $\alpha(i)$ is strictly positive, coordinating signals are perfectly correlated across forecasters.

The idiosyncratic signals represent private information, including a forecaster's subjective judgment about trend inflation. Formally, the idiosyncratic signal received by forecaster i in period t , $s_t(i)$, reads:

$$s_t(i) = \bar{\pi}_t + v_t(i); \quad (9)$$

$$v_t(i) = \rho(i)v_{t-1}(i) + \sigma_\nu(i)\nu_t(i), \quad (10)$$

whether the same should hold for our setting of long-run expectations. If this were the case, our model would explain these revisions using the coordinating and idiosyncratic signals, depending on whether these revisions contribute to making expectations more coordinated or more dispersed.

¹¹The stochastic process governing the volatility $\sigma_{c,t}$ is irrelevant to the solution of the forecasters' signal extraction problem, as they understand it will always be known at the time they form their expectations. However, this process is relevant for the estimation of the model parameters, as it influences the model's likelihood function.

¹²By sensitivity of expectations, we mean the absolute value of the response of expectations to a unit change in the signal.

where $v_t(i)$ denotes the realized noise in the idiosyncratic signal, which follows a forecaster-specific auto-regressive process with persistence $\rho(i)$ and standard normal innovations that are assumed to be i.i.d. and orthogonal across forecasters. The volatility of idiosyncratic innovations is forecaster-specific and is denoted by $\sigma_\nu(i)$.

Idiosyncratic noise has two key features. The volatility of idiosyncratic innovations, $\sigma_\nu(i)$, is inversely related to the accuracy of the private information held by each forecaster. A relatively high volatility implies that a forecaster receives relatively imprecise private information regarding trend inflation. All else being equal, a larger volatility lowers the sensitivity of a forecaster’s expectations to idiosyncratic noise, $v_t(i)$. The second key feature is the serial correlation of idiosyncratic noise, $v_t(i)$, allowing the model to account for persistent effects of individual judgment on forecasts, including, but not limited to, enduring disagreement over the inflation generating process among professional forecasters.

Since agents observe inflation and know the trend-cycle model’s parameters, they can determine the real-time (one-sided) estimate of trend inflation. As we will discuss, we, as econometricians, estimate trend inflation using the trend-cycle model with a two-sided (smoothed) approach, that is, by conditioning the estimation of trend inflation on the full sample of inflation observations. This feature of our empirical analysis has two important implications. First, forecasters in the model focus mostly on learning the forward-looking component of estimated trend inflation, as they already know past realizations of inflation and, therefore, the backward-looking drivers of trend inflation that we estimate in the data. Second, the coordinating and idiosyncratic signals jointly capture forward-looking information about trend inflation that is not yet reflected in the most recent inflation readings, which are already conveyed by the inflation signal.

3.3 Forecasters’ cognitive abilities

We consider two alternative assumptions regarding forecasters’ cognitive ability. In one version of the model—the *rational model*—we assume that agents are rational, meaning that they are capable of correctly assessing the true parameter values of the signal processes: the persistence of the coordinating and idiosyncratic noise (ρ_c and $\rho(i)$), the relative volatility of the coordinating signal ($\alpha(i)$), and the volatility of the idiosyncratic innovations ($\sigma_\nu(i)$).

We also consider a version of the model—the *behavioral model*—in which forecasters are less capable of understanding the key parameters of their signal extraction problem. Specifically, forecasters may misperceive the parameters of the signal processes. We denote the misperceived parameters with a star superscript. A forecaster may misperceive the persistence ($\rho_c^*(i) \neq \rho_c$) and the relative volatility ($\alpha^*(i) \neq \alpha(i)$) of the coordinating signals,

as well as the persistence ($\rho^*(i) \neq \rho(i)$) and the innovation volatility ($\sigma_\nu^*(i) \neq \sigma_\nu(i)$) of their idiosyncratic signals.¹³ Importantly, the degree of misperception can vary across forecasters and parameters. In both versions of our model, agents are assumed to know the parameters of the trend-cycle model. Later in the paper, we relax this assumption to show that this deviation from rationality does not resolve the two critical sources of misspecification in the rational model that we find.

4 Estimation

The rational and behavioral models are estimated using inflation data and our panel data on long-run inflation expectations, using a two-step estimation strategy. The rational model is nested within the behavioral model, representing a special case where $\rho_c^*(i) = \rho_c$, $\alpha^*(i) = \alpha(i)$, $\rho^*(i) = \rho(i)$, and $\sigma_\nu^*(i) = \sigma_\nu(i)$. For the behavioral model, we do not impose specific cognitive biases. Instead, the estimation procedure identifies the biases required to improve the model’s fit relative to the rational model.

4.1 The two-step estimation strategy

In the first step, we estimate the time-varying parameters of the trend-cycle model (σ_ω , γ_η , γ_λ , and γ_ϕ), summarized by equations (1)–(6), using our inflation data. This is achieved through Bayesian state-space methods. Additionally, we employ Markov chain Monte Carlo techniques to obtain smoothed estimates of the trend, cycle, and i.i.d. components of inflation.

In the second step, we estimate the parameters of the rational and behavioral models described in section 3. This estimation uses our panel of long-run inflation expectations, inflation, and the smoothed (two-sided) estimates of the trend and cyclical component of inflation obtained in the first step. In both models, forecasters are assumed to use parameter values for the trend-cycle model equal to their posterior means estimated in the first step, which are therefore not re-estimated in the second step.

More specifically, in the second step we estimate a state-space model that combines equations (1) to (6) with N sets of equations corresponding to the updating of the N forecasters’ expectations about the state via the Kalman filter. These equations are the solution to the signal extraction problems solved by the forecasters.¹⁴ As the trend-cycle model implies that the expected value of inflation at long horizons equals the trend, $\bar{\pi}_t$, we equate

¹³Since the average precision of the coordinating signals—as captured by $1/\sigma_{c,t}$ —is perfectly observed in this specification of the behavioral model, misperception regarding the parameter $\alpha(i)$ effectively captures misperception of the precision of the coordinating signal.

¹⁴Appendix A.1 describes the state-space model we use for the panel estimation in detail.

forecasters’ long-run inflation expectations at a point in time with their contemporaneous estimate of $\bar{\pi}_t$. The second-step estimation yields estimates of the signal process (ρ_c , $\sigma_{c,t}$, $\alpha(i)$, $\rho(i)$, and $\sigma_\nu(i)$, $i = 1, 2, \dots, N$). When we estimate the behavioral model, the second-step estimation also returns the parameters capturing the perceived relative volatility of the coordinating signals ($\alpha^*(i)$, $i = 1, 2, \dots, N$), the perceived volatility of idiosyncratic innovations ($\sigma_\nu^*(i)$, $i = 1, 2, \dots, N$), and the perceived noise persistences ($\rho_c^*(i)$ and $\rho^*(i)$, $i = 1, 2, \dots, N$).

In the second step, we, as econometricians, cannot observe the realization of the coordinating and idiosyncratic signals. However, jointly observing the panel of individual expectations and the dynamics of trend inflation estimated in the first step allows for the identification of the precision of the coordinating and idiosyncratic signals or their perceived precision in the case of the behavioral model. To illustrate, if the coordinating and idiosyncratic noise are, or are perceived to be, highly volatile, the two non-inflation signals convey little information. This occurs when the models can explain the panel of long-run inflation expectations almost entirely with the behavior of inflation, that is, the inflation signal.

The rationale for the two-step approach The trend and cyclical components of inflation are estimated in the first step by conditioning on inflation observed over the entire sample period (1959Q1–2023Q2). The long sample is a critical feature, as we treat the trend estimated in the first step as the “true” trend forecasters track when solving their signal extraction problem. While obviously no one knows the true trend inflation with absolute certainty, we consider the first step to deliver the best available estimate of trend inflation as it comes from a state-of-the-art trend-cycle model conditioning on a very long sample period. In Section 6, we use alternative approaches to estimate the “true” trend in order to assess the robustness of our findings.

We want to evaluate the relative empirical performance of the rational model and the behavioral model to gain insights into potential deviations from rationality that may influence professional forecasters’ long-run inflation expectations. In comparing these two models, the three latent components of inflation and the parameters of the trend-cycle model are left unchanged. This approach ensures that the comparison is based solely on the empirical performance of the two models in fitting the data. Such a comparison would not be feasible if the parameters of the trend-cycle model and the signal-extraction problem were estimated jointly.

A further advantage of our two-step estimation strategy is that trend inflation is estimated solely based on the actual dynamics of inflation. This approach ensures that the estimated trend inflation is not influenced by the need to explain the joint dynamics of forecasters’

inflation expectations. We consider this a desirable feature of our estimation strategy for at least two reasons. First, it eliminates the possibility of a circular argument. The estimated trend component, $\bar{\pi}_t$, represents the variable forecasters aim to track when forming their expectations in our model. Therefore, it is conceptually desirable to avoid contaminating the estimation of this trend component with information derived from the panel of observed expectations of professional forecasters. Second, the two-step estimation strategy is econometrically advantageous, as it facilitates the identification of the sensitivity of expectations to the coordinating and idiosyncratic signals, which, as noted earlier, cannot be directly observed by the econometrician.

4.2 Data

Our panel dataset is expectations of long-run headline CPI from the SPF, which spans the sample 1991Q3–2023Q2. We use these data because they are the longest available time series of professional forecasters’ long-run expectations. They also have the benefit of never being revised. Quarterly headline CPI inflation is from the U.S. Bureau of Labor Statistics and spans the sample 1959Q1–2023Q2.

Since we equate forecasters’ long-run inflation expectations in a quarter with their estimate of the contemporaneous trend component, ideally we would use data on long-run expectations that excludes near term forecasts that are mostly driven by the cyclical and high-frequency components of inflation. The measure that most closely matches our ideal is the 5-Year, 5-Year forward inflation expectations, which is generally not affected by cyclical and high-frequency shocks to inflation. However, these data only become available in 2011 in the SPF.¹⁵ To maximize the length of our sample, we splice these data to forecasts of CPI inflation over the next ten years that go back to 1991Q3. When shocks to inflation are small and transitory, as they were from 2000 to before the pandemic, ten-year average expectations are a good approximation to our ideal measure. We use the 5-Year, 5-Year forward inflation expectations in the later part of the sample that includes the 2021–2022 surge of inflation.¹⁶

To have a sufficient number of observations to measure the variance and serial correlation of the idiosyncratic noise, we consider only those forecasters who submitted at least 32 forecasts in the available sample period. This leaves us with an unbalanced panel of 51

¹⁵In 2011, the survey began including a check to verify that forecasters correctly understand their 5-year, 5-year forward inflation forecast, which is a derived variable. The *Blue Chip Economic Forecast* has a longer semiannual time series for six to ten years average inflation expectations than the SPF. However, individual long-run forecasts are not reported.

¹⁶The forecasters in the SPF do not observe inflation in the quarter they are surveyed because they submit their forecasts in the second month of each quarter. We address this by lagging SPF forecasts by one quarter when we estimate the model. For example, we measure long-run expectations in 2018Q2 using forecasts from the survey that was conducted in February 2018.

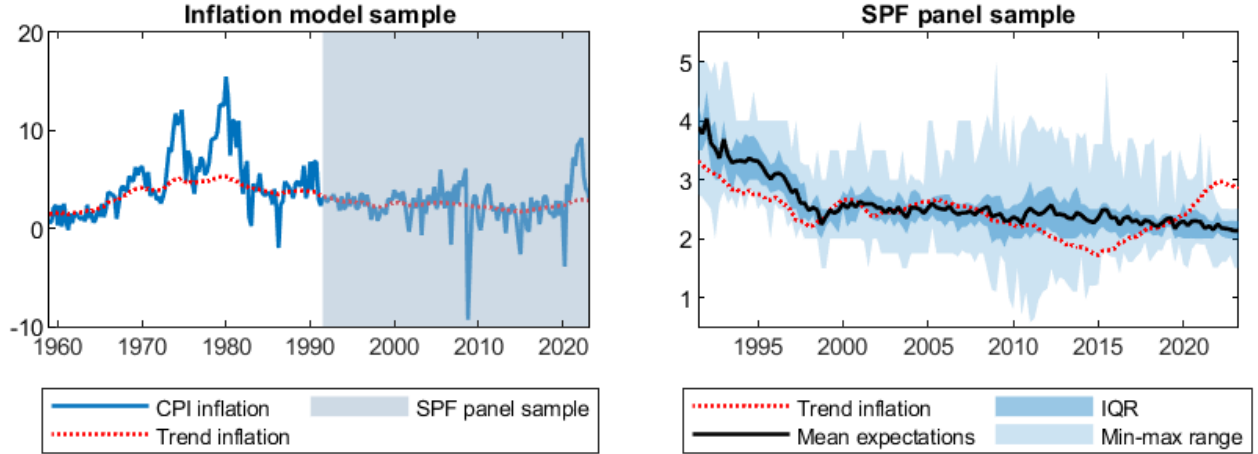


Figure 1: **Inflation, trend inflation, and long-run inflation expectations.** *Left panel:* U.S. quarterly headline CPI inflation rate (blue solid line) and its trend component (red dotted line) estimated using the trend-cycle model conditional on the entire data set (1959Q1-2023Q2). The shaded gray area indicates the sample period for the second-step panel estimation (1991Q3-2023Q2). *The right panel:* Trend component of U.S. quarterly headline CPI inflation rate estimated using the trend-cycle model (red dotted line) and the mean (black solid line), the interquartile range (the dark blue bands), and the min-max range (the light blue bands) of the distribution of SPF long-run CPI inflation expectations, constructed as explained in Section 4.2.

forecasters. Note that in some cases, there are gaps in the time series of forecasts for individual forecasters.¹⁷ In Appendix B, we show that average and median long-run expectations in our sample of forecasters correspond closely to their values in the complete SPF sample.

The estimated trend-cycle model yields the inflation trend that is data to the second step of the estimation of the rational and behavioral models.¹⁸ The left chart of Figure 1 shows the time series of CPI inflation and our smoothed estimate of the inflation trend over the full estimation sample. The shaded area shows the sample period we use to estimate the panel model. The right chart shows the mean, interquartile range, and the min-max range of the long-run inflation expectations from the SPF panel, along with the inflation trend in red. The shaded area in this chart illustrates the data the rational and behavioral models are asked to explain in the second step of the estimation. Average expectations are not used in the estimation, but we study them in the estimated model. In the early part of the sample, average expectations lag behind the decline in the trend. Between 1999 and 2008, the two series are roughly aligned and stable at around 2.5 percent. This alignment is particularly striking, given that the estimation of the trend is *not* conditioned on SPF expectations. Although the models used by professional forecasters to predict long-run

¹⁷The Philadelphia Fed must decide whether a forecaster ID should follow a forecaster when they change employer. Information on the Philadelphia Fed's website indicates that such decisions are based on judgments as to whether the forecasts represent the firm's or the individual's beliefs. See <http://www.phil.frb.org/econ/spf/Caveat.pdf>.

¹⁸The estimated parameters of the trend-cycle model are shown in Appendix C.

inflation are unknown, this finding is reassuring as it suggests that our estimated trend-cycle model aligns closely with the beliefs of an average forecaster during a period of relative macroeconomic stability. Around the beginning of the Great Recession, the trend begins a long downward drift to its nadir of 1.7 percent in 2015, and ends the sample near 3 percent. Our model will have to explain why average inflation expectations evolve so slowly in the face of these pronounced swings in trend inflation over the past 15 years.

Notice that the distribution of expectations expands quite considerably during the first period the federal funds rate was at its ELB, 2008Q4–2015Q4.¹⁹ More generally, the large and time-varying cross-sectional dispersion in the individual SPF long-run inflation expectations (the blue ranges in the right chart of Figure 1) may also represent a challenge for the rational and behavioral models.

5 Estimated models and deviations from rationality

We show the prior moments of the models’ parameters in Appendix D. The priors are fairly uninformative, reflecting our belief that the SPF data should be primarily driving the estimation. The estimation drives the auto-correlation parameter of the noise in the coordinating signal, ρ_c , to a value close to one in both models. In the estimated rational model, this process is essentially a random walk, whereas in the behavioral model this persistence is somewhat lower (0.96).

In Figure 2, we show the distributions of the posterior modes of the forecaster-specific parameters that govern the signal structure of the model. The left column shows the distribution of estimated parameters across forecasters for the rational model. The right column reports the distributions for the *true* (and unknown to the forecasters) value of the same parameters in the behavioral model. The distributions are broadly similar across the two models for the relative volatility of the coordinating signals and the persistence of the idiosyncratic noise. In the behavioral case, there are some large outliers compared to the rational model for the volatility of the idiosyncratic innovations.

These charts illustrate the significant heterogeneity of the estimated parameters of the signal structure. For the rational model, the wide distributions suggest that the degree of sensitivity of forecasters’ expectations to the signals is quite heterogeneous. To draw conclusions on the sensitivity of expectations in the behavioral model, we need first to examine

¹⁹The shrinking range during the second ELB period and toward the end of the sample period, particularly the min-max range, is an artifact of nearing the sample’s end. Specifically, as the sample period concludes, the number of new forecasters entering the survey and meeting the 32-observation cutoff is insufficient to offset the number of dropouts. A chart without the cutoff is provided in Appendix B. This shows no significant shrinkage in the min-max range toward the end of the sample period.

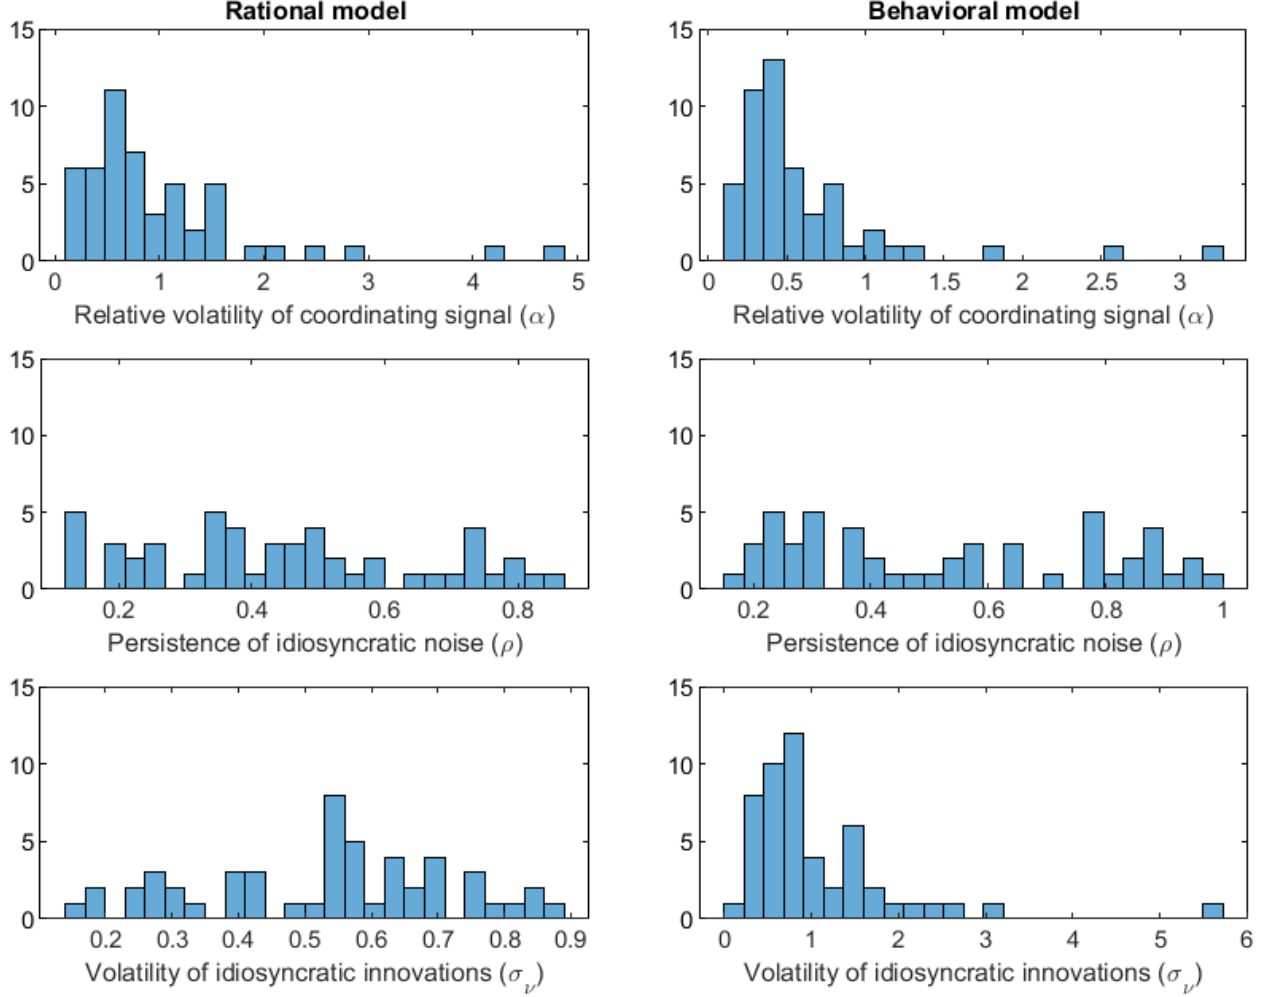


Figure 2: **Estimated model parameters.** Posterior mode for the rational (left charts) and behavioral (right charts) models' forecaster-specific parameters: the relative volatility of each forecaster's coordinating signal $\alpha(i)$ (first row), the persistence of each forecaster's idiosyncratic noise, $\rho(i)$, (second row), and the volatility of each forecaster's idiosyncratic innovations, $\sigma_\nu(i)$ (third row). Bars indicate number of forecasters.

the distributions of the estimated parameters that control the *perceived* volatility of the signals' innovations and persistence of the noise. We show the perceived parameters in Figure 3. Specifically, we present the ratio of the perceived parameters ($\rho_c^*(i)$, $\alpha^*(i)$, $\rho^*(i)$, and $\sigma_\nu^*(i)$) to the their true values (ρ_c , $\alpha(i)$, $\rho(i)$, and $\sigma_\nu(i)$). The vertical red line marks the unity line. Parameters estimated close to this line should be interpreted as reflecting small cognitive bias by the forecaster.

One of the most remarkable findings from Figure 3 concerns the perceived persistence of coordinating noise: All but one forecaster underestimates ρ_c . This systematic deviation from rationality leads to a persistent expectations bias that results in highly persistent forecast errors across almost all forecasters. This is illustrated in the left chart of Figure 4, which

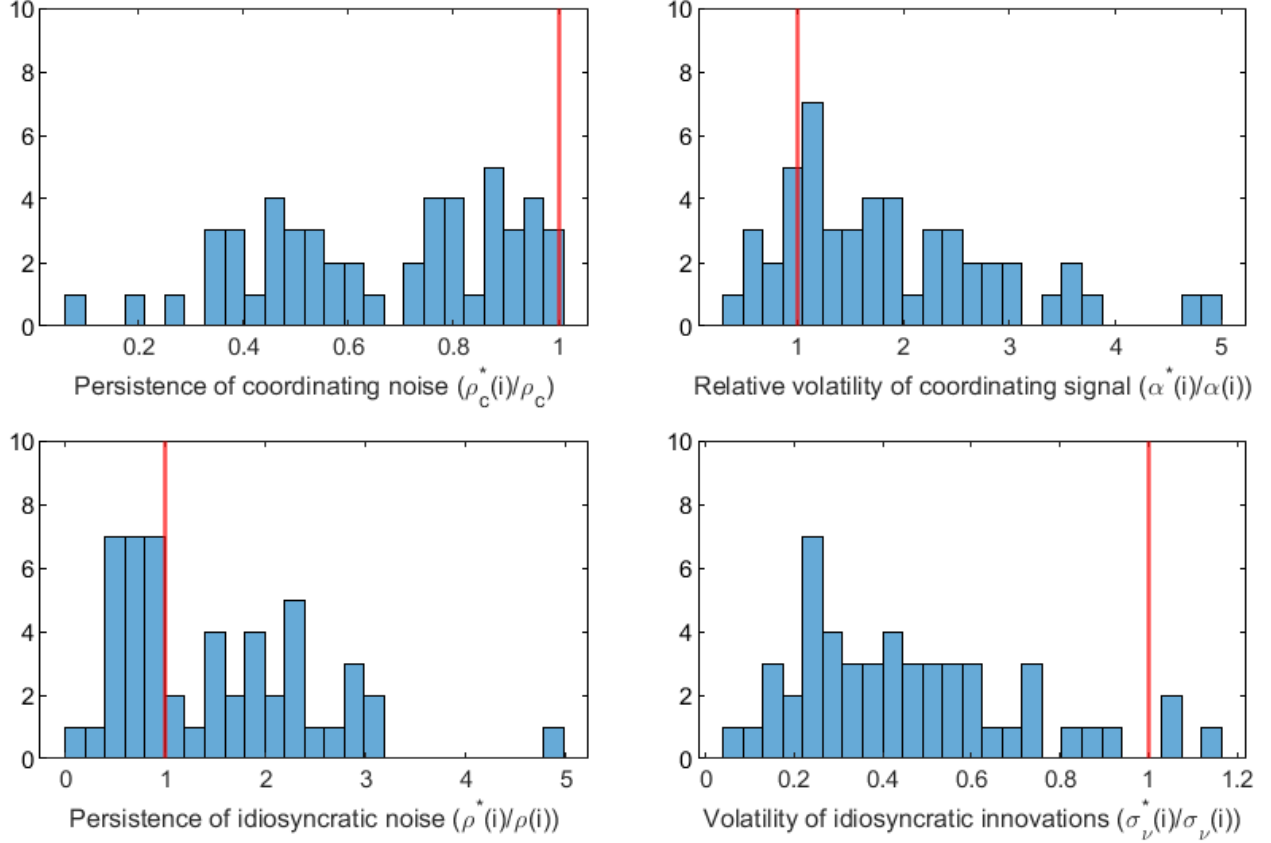


Figure 3: **Deviations from rationality.** Ratio of the posterior mode of the behavioral parameters ($\rho_c^*(i)$, $\alpha^*(i)$, $\rho^*(i)$, $\sigma_\nu^*(i)$) to the posterior mode of the true value of the corresponding structural parameters (ρ_c , $\alpha(i)$, $\rho(i)$, and $\sigma_\nu(i)$). The red vertical line denotes the line of rationality where the ratio is one and the perceived parameters coincide with their true value.

shows the response of long-run inflation expectations to a one-standard-deviation innovation to the coordinating noise, $\nu_{c,t}$. The substantial persistence in forecast error in the behavioral model is evident from the difference between the green line (the median response of long-run expectations) and the red dashed line, which is the response of trend inflation that overlaps the zero line because the trend is not affected by the innovations.²⁰

To understand why forecast errors are so persistent in the behavioral model, recall that in this model forecasters mistakenly believe that the noise can influence the coordinating signal only temporarily ($\rho_c^*(i) < \rho_c$ for all but one forecaster). However, the noise is actually very persistent ($\rho_c = 0.96$). Over time, forecasters begin to infer that only a change in trend inflation can account for the persistent variation in the coordinating signal they observe. Consequently, an innovation to the coordinating noise, $\nu_{c,t}$, leads to a long-lasting adjustment

²⁰In both the left and right charts of Figure 4, the impulse responses are computed for every period in the sample and then averaged across periods to account for the time-varying parameters of the trend-cycle model and the volatility of the innovations in the coordinating signals, $\sigma_{c,t}$.

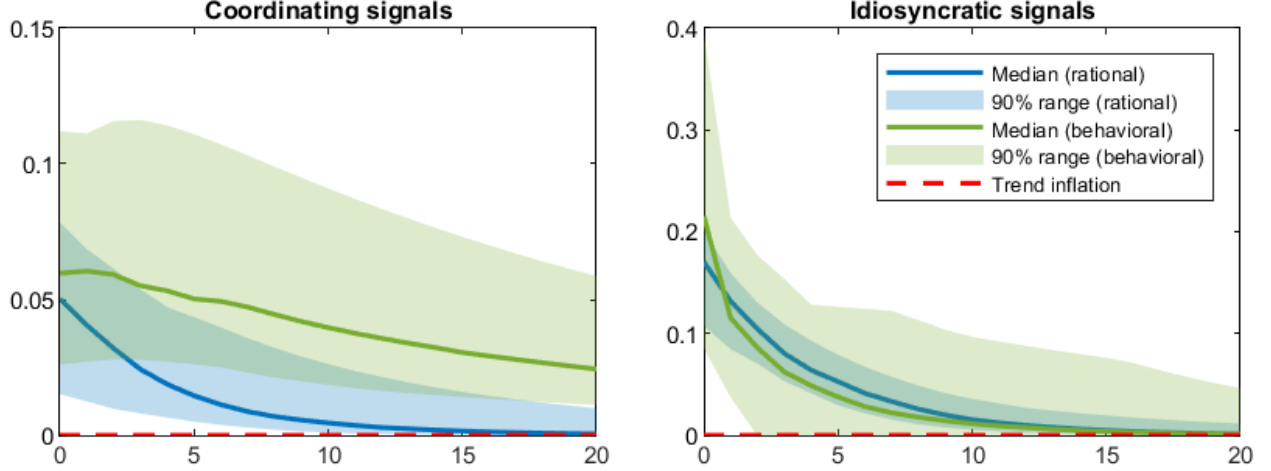


Figure 4: **Propagation of noise innovations to expectations.** Impulse response functions of long-run inflation expectations of every forecasters to a one-standard-deviation noise innovation to the coordinating signal (left chart) and to every idiosyncratic signal (right chart). Responses in the rational model are denoted in blue and those in the behavioral model are denoted in green. The solid lines denote the median response across forecasters and the shaded areas denote the 90 percent range of responses across forecasts. The red dashed line shows the true response of trend inflation. The impulse response functions in both graphs are computed in every period of the sample (1991Q3-2023Q2) and then averaged across sample periods.

in forecasters' expectations.

The behavioral model is clearly capable of capturing autonomous, long-lasting deviations of long-run inflation expectations from trend inflation. This feature enables the model to explain the persistent gaps between inflation expectations and trend inflation for most forecasters over large parts of the sample, as shown in the right chart of Figure 1. In contrast, the rational model does not have a mechanism allowing it to account for the persistent gaps in the data due to its inability to generate a persistent expectations bias.

In the left chart of Figure 4, the responses of individual expectations to innovations to the coordinating noise, $\nu_{c,t}$, are more dispersed and exhibit greater persistence in the behavioral model compared to the rational model. This is indicated by the wider 90 percent range of responses for the behavioral model (green shaded area) compared to the rational model (blue shaded area). This feature helps the behavioral model account for the large heterogeneity in the SPF data.

Another remarkable finding is shown in the lower right chart of Figure 3. In the behavioral model, almost every forecaster mistakenly believes that the volatility of their idiosyncratic innovation is lower than it actually is; that is $\sigma_\nu^*(i) < \sigma_\nu(i)$ for almost every forecaster i . This misperception implies the forecasters believe their idiosyncratic signal is more precise than it actually is. In other words, forecasters exhibit overconfidence in private information.

The implications of this are shown in the right chart of Figure 4, which shows forecasters' impulse responses to a one-standard-deviation idiosyncratic innovation, $\nu_t(i)$. In this chart,

the green shaded area is much wider than the blue one at all horizons. Overconfidence leads expectations to be more sensitive to idiosyncratic signals compared to the rational model. As a result, smaller in-sample realizations of innovations to the idiosyncratic noise are needed to explain the large and time-varying heterogeneity of forecasts. As we will show below, this feature allows the behavioral model to overcome a key source of misspecification that plagues the rational model. The wider green shaded areas in the two charts of Figure 4 highlight another significant property of the behavioral model: coordinating and idiosyncratic innovations cause some forecasters to disagree with their peers for extended periods.

Finally, the two off-diagonal charts in Figure 3 show that, in the behavioral model, most forecasters tend to overestimate the volatility of the coordinating signals, as evidenced by $\alpha^*(i) > \alpha(i)$ for the majority of forecasters, so they believe the coordinating signal is less precise than it actually is. In contrast, the share of forecasters who overestimate the persistence of the idiosyncratic noise ($\rho^*(i) > \rho(i)$) is more evenly balanced.

6 Empirical evaluation of the models

To evaluate the empirical performance of the rational and behavioral models, we examine the in-sample properties of the estimated coordinating innovations ($\nu_{c,t}$) and idiosyncratic innovations ($\nu_t(i)$). The upper charts of Figure 5 present the autocorrelation functions of the estimated innovations to the coordinating noise in the rational model (left) and the behavioral model (right). In the rational model, the assumed i.i.d. innovations to the coordinating noise exhibit significant in-sample serial correlation. This arises because the rational model requires persistent innovations to the coordinating signals to align with persistent deviations of expectations from trend inflation shown in the right chart of Figure 1. As shown in the left chart of Figure 4, the impulse response to changes in the coordinating signal in the rational model is insufficiently persistent to account for this phenomenon.²¹ Consequently, the estimated rational model can only explain the persistent gap between average inflation expectations and trend inflation by violating the white-noise assumption for innovations to the coordinating signals. This is a clear indication that the rational model is misspecified.

The upper right chart of Figure 5 shows the autocorrelation function of the coordinating innovations in the behavioral model. These innovations are much closer to white noise because, in this model, forecasters are influenced by a persistent expectations bias when forming their long-run inflation expectations. This bias enables the coordinating signals to generate highly persistent forecast errors, as illustrated by the green line in the left chart of Figure 4. Consequently, the behavioral model does not need to violate the white noise

²¹The impulse responses to all four shocks of the models are shown in Appendix E.

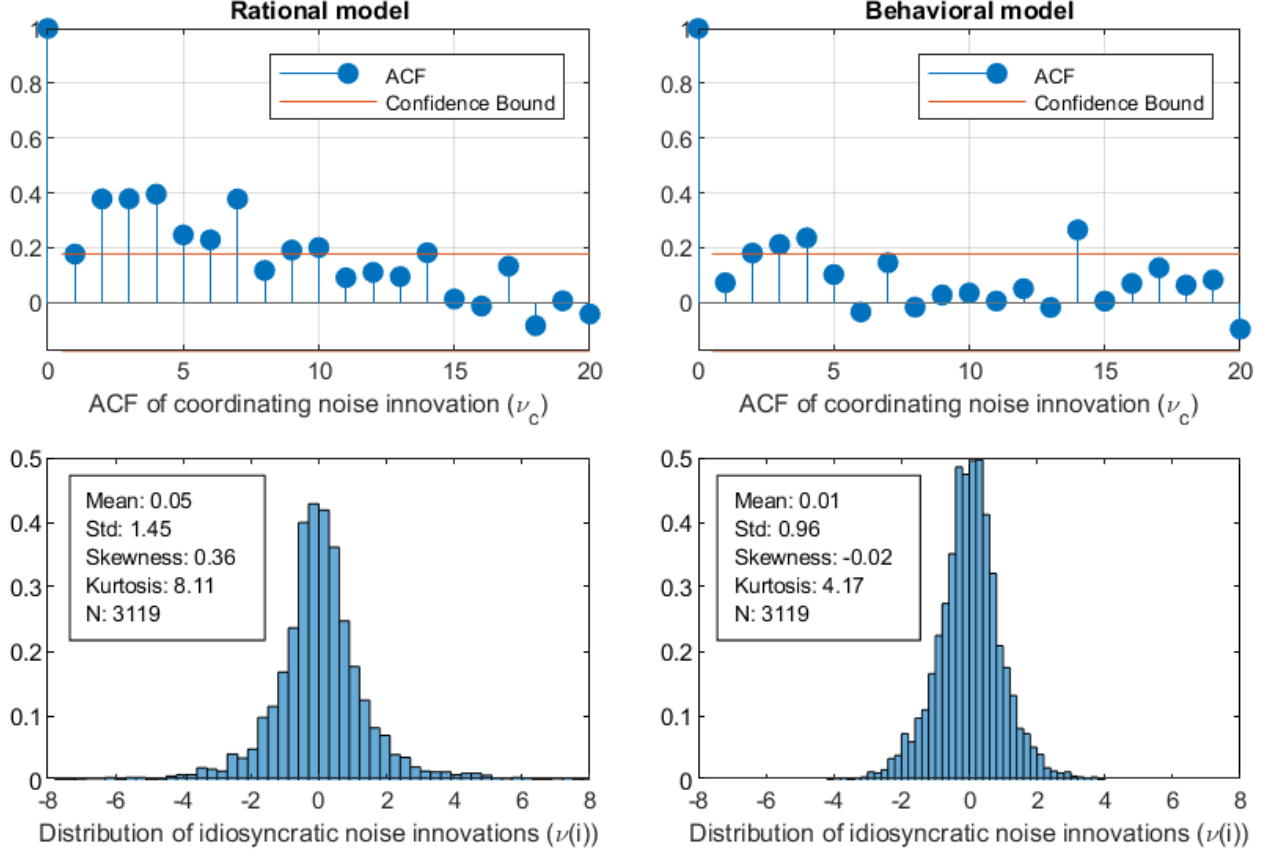


Figure 5: **Assessing misspecification in the two models.** The upper charts show the autocorrelogram of the estimated in-sample noise innovations to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The lower charts show the distribution of the estimated in-sample noise innovations to idiosyncratic signals over time and across forecasters. The left panels refer to the rational model. The right panels refer to the behavioral model.

assumption for the in-sample realizations of the coordinating innovations to account for the persistent gap between average long-run inflation expectations and the trend component of inflation.

The idiosyncratic innovations, $\nu_t(i)$, are drawn from a standard normal distribution in both models. The lower charts of Figure 5 display the distribution of the estimated in-sample realizations of the idiosyncratic innovations for all forecasters. In the rational model (left chart), the estimated idiosyncratic innovations exhibit excessive volatility, skewness, and kurtosis. In contrast, the idiosyncratic innovations estimated in the behavioral model (right chart) align much more closely with a standard normal distribution. The mean is near zero, the standard deviation is close to unity and there is virtually zero skewness. Moreover, the behavioral assumptions reduce the kurtosis of the distribution from 8 to 4, demonstrating a much closer fit to the assumed Gaussian properties.

The excess volatility, skewness, and kurtosis of the in-sample realizations of the id-

idiosyncratic innovations in the rational model highlight its struggle to account for the large cross-sectional dispersion observed in the SPF. This failure stems from the rigid link imposed by the rationality assumption between the standard deviation of the idiosyncratic innovations and the sensitivity of expectations to these innovations. On the one hand, fitting the large heterogeneity in the SPF requires a high standard deviation of idiosyncratic innovations. On the other hand, rational forecasters' expectations become less sensitive to idiosyncratic signals as the standard deviation of these innovations increases. However, reduced sensitivity to these signals causes expectations to become less dispersed, creating a tension. This tension underpins the misspecification of the rational model.

The behavioral model performs better by relaxing the tension between the volatility of the idiosyncratic innovations and the sensitivity of forecasters' expectations to idiosyncratic signals. Due to overconfidence ($\sigma_\nu^*(i) < \sigma_\nu(i)$ for almost all forecasters i), individual expectations in the behavioral model are generally more responsive to idiosyncratic shocks than in the rational model, as reflected in the size of the shaded area in green compared to the blue in the right chart of Figure 4. Consequently, the behavioral model requires smaller in-sample realizations of the idiosyncratic innovations to account for the large heterogeneity observed in the SPF.

Robustness So far, we have assumed that forecasters have perfect knowledge of the parameters of the trend-cycle model. While this assumption is consistent with the assumption of rationality, it is undoubtedly a strong and unrealistic one. In practice, forecasters must simultaneously forecast trend inflation and estimate the parameters of their inflation model in real time. Consequently, the only parameters available to forecasters are those estimated using data up to the point of their forecast. We find that this alternative model, where forecasters solve their signal extraction problem each period based on real-time parameter estimates of the trend-cycle model, suffers from the same type of misspecification that undermines the rational model.²²

We also explore another deviation from rationality: Forecasters using an incorrect model to predict trend inflation. Specifically, we assume that the “true” trend inflation is calculated using a centered 5-year moving average. However, forecasters mistakenly believe that the “true” model is the trend-cycle model with time-varying parameters, as specified in equations (1) to (6). When we estimate this behavioral model, we do not see substantial improvement in addressing the misspecification that affects the rational model. Similar conclusions are reached when the “true” trend of inflation is estimated using a 10-year centered moving average.

²²Detailed results are presented in Figure 14.

The likely reason for this failure is that moving averages are generally more responsive to cyclical fluctuations in inflation than the trend estimated using a model with time-varying parameters and heteroskedasticity. As a result, the gap between a moving average of inflation and the SPF inflation expectations becomes even larger than when trend inflation is estimated using our trend-cycle model. This, in turn, exacerbates the in-sample serial correlation of the innovations to the coordinating signals, making the issue even more pronounced. Detailed results can be found in Appendix E.

7 Forecasters' overreaction and persistent forecast errors

In this section, we delve deeper into the reasons behind the failure of the rational expectations model and further explore how the two cognitive biases significantly enhance the model's fit. Our analysis is based on examining the properties of forecast errors using the regression framework introduced by Coibion and Gorodnichenko (2015) and applied to individual forecaster data by Bordalo et al. (2020). Originally designed to test the validity of the FIRE assumption, we adapt this framework to serve as a new benchmark for evaluating the ability of the rational and behavioral models to explain key properties of forecast errors in the SPF data. This framework also allows us to compare how these properties differ between short-run and long-run inflation expectations, thereby motivating the development of a model specifically tailored to long-run expectations.

We estimate the following pooled regression:

$$\bar{\pi}_t - E_t^i \bar{\pi}_t = \beta_0^p + \beta_1^p (E_t^i \bar{\pi}_t - E_{t-1}^i \bar{\pi}_t) + \varepsilon_t^i, \quad (11)$$

where $E_t^i \bar{\pi}_t$ denotes the beliefs of forecaster i about trend inflation, which in our estimation is equated with long-run inflation expectations. The left-hand side represents the observed forecast error of a professional forecaster i in assessing trend inflation, $\bar{\pi}_t$, in period t , while the right-hand side captures the forecaster's revision between periods $t - 1$ and t . A parameter β_1^p significantly less than zero indicates *overreaction*, suggesting that the forecaster's revision of expectations in response to new signals exceeds the actual changes in trend inflation, resulting in a negative forecast error. The analysis covers our baseline sample period, spanning 1991Q3 to 2023Q2.

Overreaction in the data Panel I of Table 1 presents the results for both long-run and short-run inflation expectations. Forecast errors in predicting long-run inflation are computed using two approaches: first, as in equation (11), based on trend inflation used to estimate the

rational and behavioral models (first row of Table 1); and second, using the actual realizations of inflation (second row).²³ The latter approach requires shortening the sample period, as computing forecast errors for 5-Year, 5-Year forward inflation expectations requires observing inflation over the subsequent 10 years. Despite this limitation, the results in the table reveal that both approaches yield remarkably similar estimates of overreaction in long-run inflation expectations. This consistency mitigates any concerns regarding the use of a trend-cycle model to define professional forecasters' forecast errors in our analysis.

In the third row of the table, short-term expectations are defined as in Bordalo et al. (2020).²⁴ Notably, overreaction is statistically significant only for long-run inflation expectations.²⁵ The contrasting properties of short-term and long-run inflation expectations once again underscore the need for a model designed specifically to analyze long-run inflation expectations.

Forecasters' overreaction and overconfidence. Table 1 also includes estimates of overreaction based on data simulated using the estimated rational and behavioral models. We simulate each model for 128 periods with 51 forecasters (matching the number of observations in the sample used to estimate the models) and repeat the process 1,000 times. To ensure consistency with the data used in our model estimation, which exhibits a persistent negative forecast error, we retain only those simulations that produce a negative average forecast error. After applying this selection criterion, approximately 500 simulations remain for each case considered.

In Panel II of the table, we show the parameters in equation (11), estimated using simulated data from the rational model. When the model is simulated by drawing innovations from the true Gaussian distribution, the rational model does not reach the magnitude of overreaction observed in the data (fourth row of the table). Overreaction is a common rejection of the FIRE hypothesis, indicating that a model with agents who are jointly rational

²³This second type of forecast error is calculated consistently with how long-run inflation expectations are computed in every period of the sample. For instance, it involves averaging realized inflation rates over the subsequent 40 quarters when 10-year average inflation expectations are used as a proxy for long-run inflation expectations. This approach ensures consistency in the definition and computation of forecast errors throughout the sample period.

²⁴Specifically, forecast errors and revisions are defined with a horizon of three quarters, as follows:

$$\pi_{t+3} - E_t^i \pi_{t+3} = \beta_0^P + \beta_1^P (E_t^i \pi_{t+3} - E_{t-1}^i \pi_{t+3}) + \varepsilon_t^i,$$

where π_{t+3} denotes the annual growth rate of CPI from quarter $t - 1$ to quarter $t + 3$, with t being the survey quarter.

²⁵While Bordalo et al. (2020) find strong evidence of overreaction in SPF short-term expectations for most macroeconomic variables, they report insignificant results for CPI inflation and the GDP deflator. Our estimation, based on a different sample period, confirms their findings.

Table 1: **Regressions of forecast errors on forecast revisions**

	β_0^p	SE	β_1^p	SE	R^2
<i>I. Data</i>					
FE based on estimated trend	-0.14	0.05	-0.48	0.03	0.07
FE based on realized future inflation	-0.31	0.09	-0.48	0.06	0.03
SPF short-run inflation expectations	0.00	0.00	0.11	0.18	0.00
<i>II. Models</i>					
Rational model	-0.02	0.01	-0.39	0.03	0.07
Behavioral model	-0.12	0.03	-0.47	0.02	0.11
Behavioral model - Only overconfidence	-0.04	0.02	-0.47	0.02	0.08

Notes: Forecast Errors Analysis. Estimation results of the regression in equation (11). In Panel I, we use inflation expectations from the SPF and in Panel II we use the model-simulated beliefs about trend inflation which according to our model are equal to long-run inflation expectations. In the first row, forecast errors are computed using the trend estimated from the trend-cycle model, while the second row uses realized future inflation, with the sample restricted to end in 2013Q2. Specifically, instead of the trend, $\bar{\pi}_t$, we define the long-run inflation outcome as 5 years, 5 years forward, $\frac{1}{20} \sum_{h=21}^{40} \pi_{t+h}$, and ten year average inflation when our preferred measure is not available. In the third row, we estimate the regression specified in Bordalo et al. (2020), which is a variant of equation (11). In Panel II, we conduct 1,000 simulations for each model type. The table reports the average estimates across all simulations that yield an average negative forecast error, consistent with the estimated β_0^p in Panel I. The fourth row presents results from simulations of the rational model using innovations drawn from the true shock distribution. The fifth row reports results from simulations of the behavioral model with innovations drawn from the true shock distribution. In the sixth row, we simulate an estimated behavioral model in which the only potential deviations from rationality is overconfidence in private information—i.e. $\sigma_\nu^*(i)$ to be different from $\sigma_\nu(i)$. For the remaining parameters we do not allow any deviation from rationality. Standard errors are clustered by both forecaster and time, and the reported R^2 values correspond to the adjusted R^2 . The reported standard errors and R^2 values correspond to the average estimates across all simulations.

and fully informed cannot explain overreaction. However, our rational model with noisy information goes two thirds of the way toward its estimated value.

The behavioral model successfully accounts for the magnitude of overreaction, yielding an estimated $\beta_1^p = -0.47$, when innovations are simulated from the true distribution. This case is shown in the penultimate row of the table. Overconfidence plays a pivotal role in explaining overreaction in the SPF. Had we estimated a behavioral model with overconfidence as the only cognitive bias, the model would still have successfully replicated the overreaction observed in the data (last row). This finding demonstrates the importance of this cognitive bias in explaining overreaction in long-run inflation expectations.

As shown in the right chart of Figure 4, a noise innovation to the idiosyncratic signal leads to a negative forecast error, i.e. $\partial E_t^i \bar{\pi}_t / \partial \nu_t(i) > \partial \bar{\pi}_t / \partial \nu_t(i) = 0$. Consequently, the rational model can generate overreaction. Nevertheless, overconfidence helps the behavioral model to magnify these effects. This property is evident from the upper bound of individual expectations' responses to noise innovations in their idiosyncratic signals, as shown in the right

panel of Figure 4. At the time the noise innovation occurs, this upper bound is significantly higher than that observed in the case of the rational model.

Taking stock, overreaction emerges as a significant feature of the SPF long-run inflation expectations, contributing to the pronounced heterogeneity observed therein, a characteristic that the rational model fails to explain. A model equipped with a mechanism to generate overreaction, such as the behavioral model with overconfidence, does not need to rely on out-sized in-sample realizations of noise innovations to idiosyncratic signals to account for the large heterogeneity in the data driven by expectations' overreaction. Indeed, the distribution of these in-sample realizations, illustrated in Figure 5, is remarkably close to a standard normal distribution.

Negative forecast error bias and persistent expectation bias. The existence of a statistically significant negative persistent bias in the SPF forecast errors is supported by the estimation results displayed in the first row of Table 1. This bias is captured by the intercept parameter β_0^p , where a negative (positive) value indicates an average negative (positive) bias within our sample. This negative bias is also clearly visible in Figure 1, where the mean of SPF expectations (the black solid line) lies above the estimated trend inflation (the red dotted line) for most periods.

The baseline behavioral model, incorporating all cognitive biases, successfully accounts for this persistent negative bias in forecast errors for long-run inflation. As shown in the penultimate row of the table, the estimated β_0^p is statistically significant and closely replicates the value estimated in the SPF data. While it perfectly accounts for the overreaction observed in the data, the behavioral model with only overconfidence cannot generate persistently biased forecast errors similar to those found in the SPF panel data – see the last row of Table 1.²⁶ The rational model, the fourth row of the table, also fails to account for the negative forecast error bias observed in the data.

Clearly, both the behavioral model with only overconfidence and the rational model lack a mechanism to generate highly persistent forecast errors. In contrast, this mechanism is present in the estimated behavioral model that incorporates all cognitive biases, particularly the persistence expectations bias. This bias gives rise to the highly persistent forecast errors illustrated by the green line in the left panel of Figure 4.

The negative average bias becomes even more pronounced when forecast errors are calculated using realized future inflation rates, as highlighted in the second row of Table 1. This larger bias arises from the necessity of excluding the most recent sample observations,

²⁶Since we only consider simulations from models that imply negative average forecast errors, this parameter is, by construction, estimated to be negative when model simulations are used.

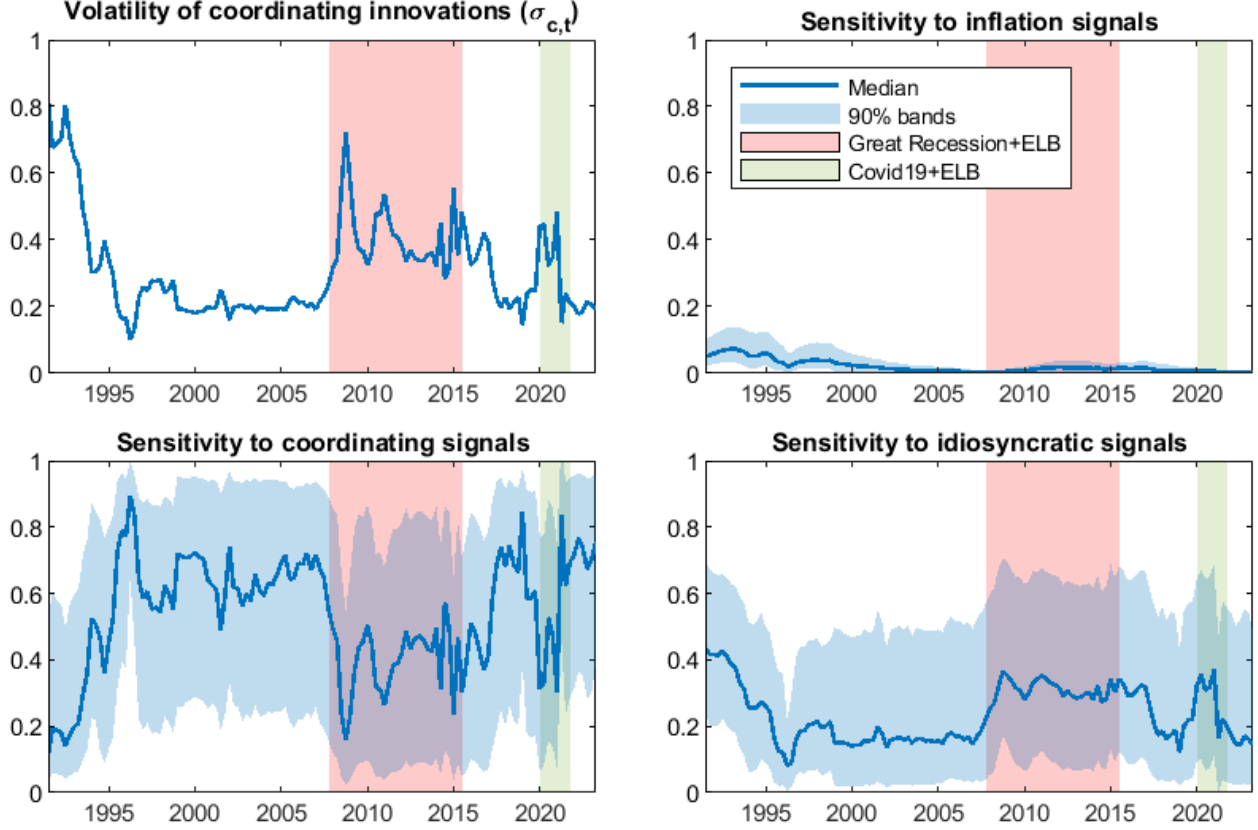


Figure 6: **Time-varying volatility of coordinating innovations and sensitivities of expectations to signals.** The upper left chart shows the posterior mode of the time-varying volatility of coordinating innovations ($\sigma_{c,t}$). The upper right chart shows the median and 90 percent range of the distribution of sensitivities of individual expectations to inflation. The lower left chart shows the median and 90 percent range of the distribution of sensitivities of individual expectations to the coordinating signals. The lower right chart shows the median and 90 percent range of the distribution of sensitivities of individual expectations to the idiosyncratic signals. Sensitivity is defined as the absolute value of the initial response of forecaster i 's expectations to a unitary innovation to the signal. Technically, the sensitivity is derived from estimated Kalman gain (based on the perceived parameters) which represents how much forecasters' expectations about long-run inflation respond to new information as captured by the three signals.

during which inflation surged sharply, leading to persistent positive forecast errors.²⁷ Using this model-free definition of forecast errors would further increase the negative bias, thereby exacerbating the struggles the rational model faces in accounting for the SPF panel data.

8 The sensitivity of expectations to private and public information

Our estimated behavioral model of long-run inflation expectations (henceforth, the model) allows us to measure the sensitivity of long-run inflation expectations to inflation, the

²⁷Recall that the need to end our sample period in 2013 is the main reason why we define forecast errors for long-run inflation expectations using the trend-cycle model introduced in Section 3.

coordinating signal, and the idiosyncratic signals. Sensitivity is defined as the absolute value of the initial response of forecaster i 's expectations to a unitary innovation to the signal. Technically, the sensitivity is equal to the estimated Kalman gains from the signal extraction problem solved by each forecaster (see Appendix A.1 for the derivation of the Kalman gains). These Kalman gains capture how forecasters update their beliefs about long-run inflation in response to new information as captured by the three signals. They are a function of the estimated parameters from both the trend-cycle model and the signal extraction problem and hence vary across forecasters and over time. Importantly, since the behavioral model allows for misperception of the signal processes' parameters, the estimated Kalman gains are based on the perceived parameters. Therefore, these Kalman gains are *subjective* and not necessarily equal to the optimal Kalman gains as in the case of the rational framework.

The upper-right chart of Figure 6 illustrates the sensitivity of forecasters' long-run inflation expectations to the inflation signal. Throughout the sample period, sensitivity to the idiosyncratic signal remains remarkably low, suggesting that long-run expectations were largely insulated from cyclical variations in inflation. This characteristic suggests that inflation expectations, particularly after 2005, were largely anchored in the sense of Bernanke (2007).²⁸ Moreover, we observe minimal variation in sensitivity across forecasters, as indicated by the narrow blue shaded area.

The lower-left chart of Figure 6 displays the sensitivity of long-run inflation expectations to coordinating signals. This sensitivity is time-varying, with its value in each period depending on the estimated values of all parameters in our model, including the time-varying parameters of the trend-cycle model. A comparison of the two charts on the left of Figure 6 reveals that this sensitivity is heavily influenced by the estimated evolution of the volatility of coordinating innovations, $\sigma_{c,t}$, which is inversely related to the average precision of coordinating signals.

The sensitivity of SPF long-run inflation expectations to coordinating signals is much larger and exhibits much greater variation across forecasters than the sensitivity to the inflation signal. The sensitivity to coordinating signals is also larger than for the inflation signal for most of the sample. Examining the dynamics of the median sensitivity (depicted by the solid blue line), we find that in the early sample period, expectations became increasingly responsive to the coordinating signal. From the late 1990s through to shortly before the onset of the Great Recession (indicated by the red shaded area), the median sensitivity stabilizes around 0.65. This value implies that a 25 basis point increase in the coordinating signal shifts expectations by approximately 16 basis points on average.

²⁸This finding aligns with empirical studies on the sensitivity of long-run inflation expectations to short-term expectations, such as Barlevy, Fisher and Tysinger (2021), Carvalho, Eusepi, Moench and Preston (2023) and references therein. Carvalho et al. (2023), Gáti (2023) propose structural models where long-run expectations exhibit endogenous sensitivity to short-term inflation surprises.

The model accounts for forecaster heterogeneity in part by allowing sensitivities to coordinating signals to vary across individuals ($\alpha(i)$), resulting in a wide range of sensitivities (represented by the blue shaded area in the lower-left chart of Figure 6). The pre-ELB period is characterized by a symmetric distribution of sensitivities. However, the onset of the first ELB episode introduces persistent changes to this distribution. Specifically, we observe notable and abrupt fluctuations in the skewness of the distribution.

At the onset of the Great Recession and the start of the first ELB period (the red shaded area), the median sensitivity falls sharply and becomes much more volatile. A similar decline occurs also in the second and most recent ELB episode (the green shaded area). When the Federal Reserve lifts off the federal funds rate to end both ELB periods, we observe an increase in the median sensitivity to coordinating signals.

The subdued sensitivity of expectations to the coordinating signals during the ELB episodes suggests a crucial role for monetary policy in coordinating public beliefs and that the interest rate tool is an important component of that. This role has been formalized by the literature on dispersed information, dating back to Morris and Shin (2002) and Woodford (2003), who showed that public signals serve as focal points for coordinating dispersed expectations about the economy’s fundamentals.²⁹

9 The anchoring-compatible inflation path

Now we show how our model can be used to evaluate whether the path of inflation communicated by policymakers is consistent with their objective of keeping long-run inflation expectations anchored. Our analysis begins in the fourth quarter of 2022 when CPI inflation significantly exceeded SPF long-run CPI inflation expectations. As illustrated in the right chart of Figure 1, despite a substantial rise in trend inflation during 2021-2022, average 5-Year, 5-Year forward CPI inflation expectations remained remarkably stable.

We use our model to address the following question: what path of inflation over the next three years would ensure that average long-run inflation expectations remain anchored, that is unchanged, at the level observed at the end of 2022? We then compare this model-derived path of inflation to the trajectory projected by the FOMC in December 2022. This exercise allows us to assess whether the inflation path communicated by policymakers was aligned with the goal of anchoring long-run expectations.

To compute this *anchoring-compatible* inflation path, we start by guessing a trajectory for trend inflation. Given this guessed path and the assumption that average long-run inflation

²⁹Other contributions to this literature include Mankiw and Reis (2002), Sims (2003), Reis (2006), Angeletos and Pavan (2007), Nimark (2008), Maćkowiak and Wiederholt (2009), Lorenzoni (2009, 2010), Angeletos and Lian (2018). Related empirical works include Melosi (2014), Falck, Hoffmann and Hürtgen (2021).

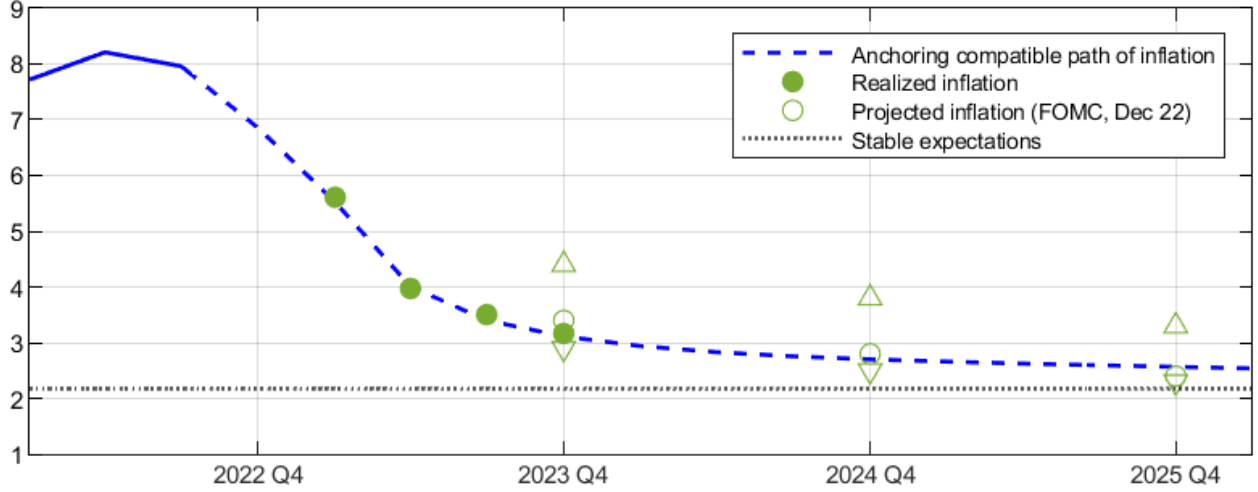


Figure 7: **Anchoring-compatible inflation path.** The solid blue line denotes actual year-on-year headline CPI inflation rate. The dashed blue line denotes the model-predicted path of year-on-year inflation rate compatible with the 5-Year, 5-Year SPF CPI inflation expectations remaining stable at their observed 2022Q4 level over the next 5 years. The green filled circles denote actual year-on-year CPI inflation rate in the first three quarters of 2023. The median Summary of Economics Projections for the fourth-quarter over fourth-quarter PCE inflation rates in 2023-2025 are denoted by the empty green circles. For comparability we map the SEP from PCE units into CPI units assuming a bias of 30 basis points. The empty green upward (downward) arrows show the upper (lower) range of the SEP projections.

expectations remain stable, the model determines the corresponding path of inflation. Using this inflation path, we then compute the smoothed, implied trend inflation to validate or refine the initial guess. The resulting anchoring-compatible inflation path is the solution to this fixed-point problem for trend inflation. All parameters are set to their estimated value in 2022Q3. In sum, finding the anchoring-compatible path essentially involves forecasting inflation, conditional on long-run expectations remaining fixed at their 2022Q4 level over the next five years and on trend inflation being consistent with the resulting path.³⁰

Although inflation expectations are not particularly sensitive to the inflation signal, they respond noticeably to the coordinating and idiosyncratic signals, which are influenced by trend inflation. The elevated trend inflation estimated at the end of 2022 exerts upward pressure on long-run inflation expectations by keeping these signals elevated. The anchoring-compatible inflation path must counteract this upward pressure, which, if unaddressed, could lead to elevated trend inflation becoming ingrained in inflation expectations, causing their de-anchoring. Communicating a shallower path of inflation would reduce trend inflation, thereby mitigating the upward pressure on inflation expectations. These dynamics underscore

³⁰Since the cross-section of expectations over the next five years is unobservable, we set the realizations of the idiosyncratic innovations to zero but assume forecasters continue to interpret idiosyncratic signals as noisy information. We also assume that inflation and trend inflation converge to the average long-run inflation expectations after five years and remain there indefinitely.

the intricate interplay between the inflation path, trend inflation, and long-run expectations in our model.

In Figure 7, the blue dashed line represents the inflation path consistent with anchored average long-run CPI inflation expectations from 2022Q4 through 2027Q4. This path features a relatively rapid decline in 2023, followed by a slower convergence to the level at which average long-run inflation expectations are assumed to remain anchored. This level, indicated by the horizontal black dotted line, corresponds to the median SPF long-run inflation expectations in 2022Q4, the starting point of this exercise.

Now we compare the anchoring-compatible path with the median and the upper-lower range of the Federal Reserve’s Summary of Economic Projections (SEP) released following the meeting of the FOMC in December 2022, which are denoted by the green dots and arrows in Figure 7.³¹ The inflation path projected by the FOMC in its SEP aligns closely with the path generated by the model. This alignment suggests that the model interprets the inflation path communicated by FOMC members as broadly consistent with maintaining anchored average long-run inflation expectations, despite the elevated inflation and trend inflation observed at the end of 2022.

In Figure 7, we also compare the anchoring-compatible path of inflation with the actual year-on-year CPI headline inflation observed in the four quarters of 2023 (solid green circles). These observations are remarkably close to the anchoring-compatible path of inflation generated by our model. Given that average SPF inflation expectations have remained very close to their 2022Q4 level throughout this period (as shown in the figure in Appendix F), we regard this alignment as an indirect empirical validation of our model’s accuracy and relevance.

10 Concluding remarks

We develop and estimate a new framework for understanding the formation of long-run inflation expectations, leveraging panel data from the U.S. Survey of Professional Forecasters. The findings underscore the role of cognitive distortions, in particular overconfidence in private information and a persistent expectations bias, in shaping forecasters’ beliefs about long-run price dynamics. By capturing key features of the data, the model reveals substantial, time-varying heterogeneity in forecasters’ responsiveness to public information, with sensitivity declining for all forecasters when monetary policy is constrained by the ELB.

Moreover, the model provides a practical framework for assessing whether the inflation

³¹The SEP is available at <https://www.federalreserve.gov/monetarypolicy/fomc.htm>. Note that the SEP reports only the Q4/Q4 headline PCE inflation rate. For comparability, we convert the SEP projections into CPI units by assuming a gap of 30 basis points.

paths communicated by policymakers align with their objective of maintaining well anchored long-run inflation expectations. These insights are helpful for understanding the formation of long-run inflation expectations and for informing effective monetary policy strategies.

While this paper primarily focuses on the cognitive biases driving the formation of forecasters’ beliefs about long-run inflation, an important avenue for future research is identifying the events and episodes that drive time-series and cross-sectional fluctuations in forecasters’ sensitivity to public information. Exploring these dynamics could deepen our understanding of how expectations evolve under varying economic conditions and policy regimes.

In future work, this framework could also be extended to other (potentially) nonstationary variables, such as expectations regarding long-run real GDP growth, productivity growth or the long-run level of interest rates, which are also asked by the SPF, albeit at a lower frequency. Investigating how public and private information shape the distribution of expectations—and their resulting macroeconomic effects—would add to the growing literature on how changes in the distribution of expectations across forecasters impact the economy, e.g. (Ascari et al., 2024).

References

- Adam, Klaus, Pei Kuang, and Shihan Xie (2024) “Overconfidence in Private Information Explains Biases in Professional Forecasts,” mimeo.
- Afrouzi, Hassan, Spencer Y Kwon, Augustin Landier, Yueran Ma, and David Thesmar (2023) “Overreaction in Expectations: Evidence and Theory,” *The Quarterly Journal of Economics*, 138 (3), 1713–1764.
- Ahn, Hie Joo and Leland E. Farmer (2024) “Disagreement About the Term Structure of Inflation Expectations,” Finance and Economics Discussion Series 2024-084, Board of Governors of the Federal Reserve System (U.S.).
- Ahn, Hie Joo and Chad Fulton (2021) “Inflation expectations and inflation forecasting,” FEDS Notes. Washington: Board of Governors of the Federal Reserve Systems, March 05, 2021.
- Andrade, Philippe and Hervé Le Bihan (2013) “Inattentive professional forecasters,” *Journal of Monetary Economics*, 60 (8), 967–982.
- Andrade, Philippe, Richard K. Crump, Stefano Eusepi, and Emanuel Moench (2016) “Fundamental disagreement,” *Journal of Monetary Economics*, 83, 106–128.
- Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry (2021) “Imperfect Macroeco-

- omic Expectations: Evidence and Theory,” *NBER Macroeconomics Annual*, 35, 1–86.
- Angeletos, George-Marios and Chen Lian (2018) “Forward Guidance without Common Knowledge,” *American Economic Review*, 108 (9), 2477–2512.
- Angeletos, George-Marios and Alessandro Pavan (2007) “Efficient Use of Information and Social Value of Information,” *Econometrica*, 75 (4), 1103–1142.
- Armantier, Olivier, Argia M Sbordone, Giorgio Topa, Wilbert Van der Klaauw, and John C Williams (2022) “A New Approach to Assess Inflation Expectations Anchoring Using Strategic Surveys,” Staff Reports 1007, Federal Reserve Bank of New York.
- Aruoba, S. Borağan (2020) “Term Structures of Inflation Expectations and Real Interest Rates,” *Journal of Business & Economic Statistics*, 38 (3), 542–553.
- Ascari, Guido, Paolo Bonomolo, and Alessandro Celani (2024) “The Macroeconomic Effects of Inflation Expectations: The Distribution Matters,” CEPR Discussion Papers 18937, C.E.P.R. Discussion Papers.
- Barlevy, Gadi, Jonas Fisher, and May Tysinger (2021) “Are long run expectations well anchored?,” *Federal Reserve Bank of Chicago Fed Letter* (458), 1059–1078.
- Beechey, Meredith J., Benjamin K. Johannsen, and Andrew T. Levin (2011) “Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area Than in the United States?” *American Economic Journal: Macroeconomics*, 3 (2), 104–29.
- Bernanke, Ben S. (2007) “Inflation expectations and inflation forecasting,” Speech given at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Available at: <https://www.federalreserve.gov/newsevents/speech/bernanke20070710a.htm>.
- Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo (2023) “Diagnostic Business Cycles,” *Review of Economic Studies*, Forthcoming.
- Bianchi, Francesco, Sydney C. Ludvigson, and Sai Ma (2022) “Belief Distortions and Macroeconomic Fluctuations,” *American Economic Review*, 112 (7), 2269–2315.
- Binder, Carola, Wesley Janson, and Randal Verbrugge (2023) “Out of Bounds: Do SPF Respondents Have Anchored Inflation Expectations?” *Journal of Money, Credit and Banking*, 55 (2-3), 559–576.
- Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer (2020) “Overreaction in Macroeconomic Expectations,” *American Economic Review*, 110 (9), 2748–82.
- Born, Benjamin, Zeno Enders, and Gernot Müller (2023) “On FIRE, news, and expectations,”

CEPR Discussion Papers 18259, C.E.P.R. Discussion Papers.

- Broer, Tobias and Alexandre N. Kohlhas (2022) “Forecaster (Mis-)Behavior,” *The Review of Economics and Statistics*, 1–45.
- Carvalho, Carlos, Stefano Eusepi, Emanuel Moench, and Bruce Preston (2023) “Anchored Inflation Expectations,” *American Economic Journal: Macroeconomics*, 15 (1), 1–47.
- Chahrour, Ryan, Kristoffer Nimark, and Stefan Pitschner (2021) “Sectoral Media Focus and Aggregate Fluctuations,” *American Economic Review*, 111 (12), 3872–3922.
- Chan, Joshua C.C., Todd E. Clark, and Gary Koop (2018) “A New Model of Inflation, Trend Inflation, and Long-Run Inflation Expectations,” *Journal of Money, Credit and Banking*, 50 (1), 5–53.
- Cogley, Timothy, Giorgio E. Primiceri, and Thomas J. Sargent (2010) “Inflation-Gap Persistence in the US,” *American Economic Journal: Macroeconomics*, 2 (1), 43–69.
- Coibion, Olivier and Yuriy Gorodnichenko (2012) “What Can Survey Forecasts Tell Us about Information Rigidities?” *Journal of Political Economy*, 120 (1), 116–159.
- (2015) “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts,” *American Economic Review*, 105 (8), 2644–78.
- Corsello, Francesco, Stefano Neri, and Alex Tagliabracci (2021) “Anchored or de-anchored? That is the question,” *European Journal of Political Economy*, 69, article 102031.
- Crump, Richard K., Stefano Eusepi, Emanuel Moench, and Bruce Preston (2023) “The term structure of expectations,” in Bachmann, Rüdiger, Giorgio Topa, and Wilbert van der Klaauw eds. *Handbook of Economic Expectations*, 507–540: Academic Press.
- Del Negro, Marco and Christopher Otrok (2008) “Dynamic factor models with time-varying parameters: measuring changes in international business cycles,” Staff Reports 326, Federal Reserve Bank of New York.
- Dovern, Jonas, Ulrich Fritsche, and Jiri Slacalek (2012) “Disagreement Among Forecasters in G7 Countries,” *The Review of Economics and Statistics*, 94 (4), 1081–1096.
- Dräger, Lena and Michael J. Lamla (2014) “Anchoring of Consumer’ Inflation Expectations: Evidence from Microdata,” manuscript, University of Essex.
- Ehrbeck, Tilman and Robert Waldmann (1996) “Why are Professional Forecasters Biased? Agency Versus Behavioral Explanations,” *The Quarterly Journal of Economics*, 111 (1), 21–40.

- Engelberg, Joseph, Charles F. Manski, and Jared Williams (2010) “Assessing the Temporal Variation of Macroeconomic Forecasts by a Panel of Changing Composition,” *Journal of Applied Econometrics*, 26 (7), 1059–1078.
- Falck, E., M. Hoffmann, and P. Hürtgen (2021) “Disagreement about inflation expectations and monetary policy transmission,” *Journal of Monetary Economics*, 118 (C), 15–31.
- Farmer, Leland, Emi Nakamura, and Jón Steinsson (2023) “Learning About the Long Run,” *Journal of Political Economy*, Forthcoming.
- Faust, Jon and Jonathan H. Wright (2013) “Chapter 1 - Forecasting Inflation,” in Elliott, Graham and Allan Timmermann eds. *Handbook of Economic Forecasting*, 2 of Handbook of Economic Forecasting, 2–56: Elsevier.
- Fuhrer, Jeffrey C. (2018) “Intrinsic expectations persistence: evidence from professional and household survey expectations,” Working Papers 18-9, Federal Reserve Bank of Boston.
- Gáti, Laura (2023) “Monetary policy anchored expectations—An endogenous gain learning model,” *Journal of Monetary Economics*, 140, S37–S47, Inflation: Drivers and Dynamics 2022.
- Giacomini, Raffaella, Vasiliki Skreta, and Javier Turen (2020) “Heterogeneity, Inattention, and Bayesian Updates,” *American Economic Journal: Macroeconomics*, 12 (1), 282–309.
- Goldstein, Nathan (2023) “Tracking Inattention,” *Journal of the European Economic Association*, 21 (6), 2682–2725.
- Goldstein, Nathan and Yuriy Gorodnichenko (2022) “Expectations Formation and Forward Information,” Working Paper 29711, National Bureau of Economic Research.
- Grishchenko, Olesya, Sarah Mouabbi, and Jean-Paul Renne (2019) “Measuring Inflation Anchoring and Uncertainty: A U.S. and Euro Area Comparison,” *Journal of Money, Credit and Banking*, 51 (5), 1053–1096.
- Gürkaynak, Refet, Andrew T. Levin, Andrew N. Marder, and Eric T. Swanson (2007) “Inflation targeting and the anchoring of inflation expectations in the Western Hemisphere,” *Economic Review*, 25–47.
- Hasenzagl, Thomas, Filippo Pellegrino, Lucrezia Reichlin, and Giovanni Ricco (2020) “A Model of the Fed’s View on Inflation,” *The Review of Economics and Statistics*, 1–45.
- Henzel, Steffen R. (2013) “Fitting survey expectations and uncertainty about trend inflation,” *Journal of Macroeconomics*, 35, 172–185.
- Herbst, Edward P. and Fabian Winkler (2021) “The Factor Structure of Disagreement,”

Finance and Economics Discussion Series 2021-046, Board of Governors of the Federal Reserve System (U.S.).

- Kim, Sangjoon, Neil Shephard, and Siddhartha Chib (1998) “Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models,” *The Review of Economic Studies*, 65 (3), 361–393.
- Kohlhas, Alexandre N. and Ansgar Walther (2021) “Asymmetric Attention,” *American Economic Review*, 111 (9), 2879–2925.
- Kozicki, Sharon and P.A. Tinsley (2001) “Shifting endpoints in the term structure of interest rates,” *Journal of Monetary Economics*, 47 (3), 613–652.
- Kurmar, Saten, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko (2015) “Inflation targeting does not anchor inflation expectations: Evidence from New Zealand,” *Brookings Papers on Economic Activity*, 43 (Fall), 151–208.
- Kučinskas, Simas and Florian S. Peters (2022) “Measuring Under- and Overreaction in Expectation Formation,” *The Review of Economics and Statistics*, 1–45.
- Lorenzoni, Guido (2009) “A Theory of Demand Shocks,” *The American Economic Review*, 99 (5), 2050–2084.
- (2010) “Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information,” *The Review of Economic Studies*, 77 (1), 305–338.
- Maćkowiak, Bartosz and Mirko Wiederholt (2009) “Optimal Sticky Prices under Rational Inattention,” *The American Economic Review*, 99 (3), 769–803.
- Mankiw, N. Gregory and Ricardo Reis (2002) “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve*,” *The Quarterly Journal of Economics*, 117 (4), 1295–1328.
- Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers (2004) “Disagreement about Inflation Expectations,” in Gertler, Mark and Kenneth Rogoff eds. *NBER Macroeconomics Annual 2003, Volume 18*, 209–270: The MIT Press.
- Melosi, Leonardo (2014) “Estimating Models with Dispersed Information,” *American Economic Journal: Macroeconomics*, 6 (1), 1–31.
- Mertens, Elmar (2016) “Measuring the Level and Uncertainty of Trend Inflation,” *The Review of Economics and Statistics*, 98 (5), 950–967.
- Mertens, Elmar and James M. Nason (2020) “Inflation and professional forecast dynamics: An evaluation of stickiness, persistence, and volatility,” *Quantitative Economics*, 11 (4),

1485–1520.

Morris, Stephen and Hyun Song Shin (2002) “Social Value of Public Information,” *American Economic Review*, 92 (5), 1521–1534.

Nason, James M. and Gregor W. Smith (2021) “Measuring the slowly evolving trend in US inflation with professional forecasts,” *Journal of Applied Econometrics*, 36 (1), 1–17.

Nimark, Kristoffer (2008) “Dynamic pricing and imperfect common knowledge,” *Journal of Monetary Economics*, 55 (2), 365–382.

Nimark, Kristoffer P. (2014) “Man-Bites-Dog Business Cycles,” *American Economic Review*, 104 (8), 2320–67.

Omori, Yasuhiro, Siddhartha Chib, Neil Shephard, and Jouchi Nakajima (2007) “Stochastic volatility with leverage: Fast and efficient likelihood inference,” *Journal of Econometrics*, 140 (2), 425–449.

Orphanides, Athanasios and John Williams (2005) “The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations,” *Journal of Economic Dynamics & Control*, 29, 1927–1950.

Patton, Andrew J. and Allan Timmermann (2010) “Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion,” *Journal of Monetary Economics*, 57 (7), 803–820.

Reis, Ricardo (2006) “Inattentive Producers,” *The Review of Economic Studies*, 73 (3), 793–821.

——— (2022) “Losing the inflation anchor,” *Brookings Papers on Economic Activity*, 2021, 307–379.

Rich, Robert W. and Joseph Tracy (2017) “The behavior of uncertainty and disagreement and their roles in economic prediction: a panel analysis,” Staff Reports 808, Federal Reserve Bank of New York.

Rossi, Barbara and Tatevik Sekhposyan (2016) “Forecast Rationality Tests in the Presence of Instabilities, with Applications to Federal Reserve and Survey Forecasts,” *Journal of Applied Econometrics*, 31 (3), 507–532.

Sims, Christopher A. (2003) “Implications of rational inattention,” *Journal of Monetary Economics*, 50 (3), 665–690.

Stark, Tom (2013) “SPF Panelists’ Forecasting Methods: A Note on the Aggregate Results of a November 2009 Special Survey,” Federal Reserve of Philadelphia, Real-Time Data

Research Center, mimeo.

Stock, James H. and Mark W. Watson (2007) “Why Has U.S. Inflation Become Harder to Forecast?” *Journal of Money, Credit and Banking*, 39 (s1), 3–33.

——— (2016) “Core Inflation and Trend Inflation,” *The Review of Economics and Statistics*, 98 (4), 770–784.

de Vincent-Humphreys, Rupert, Ivelina Dimitrova, Elisabeth Falck, Lukas Henkel, and Aidan Meyler (2019) “Twenty years of the ECB Survey of Professional Forecasters,” *Economic Bulletin Articles*, 1.

Woodford, Michael (2003) “Imperfect Common Knowledge and the Effects of Monetary Policy,” in Philippe Aghion, Joseph E. Stiglitz, Roman Frydman and Michael Woodford eds. *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, 25–58, Princeton, NJ: Princeton University Press.

A Detailed model description

A.1 Derivation of model equations

The environment confronted by forecaster i has a state-space representation given by

$$\xi_t(i) = \Phi_t(\mathbf{i})\xi_{t-1}(i) + \mathbf{R}_t(\mathbf{i})e_t(i) \quad (12)$$

$$y_t(i) = \mathbf{D}(\mathbf{i})\xi_t(i) + \mathbf{Q}u_t \quad (13)$$

where

$$\begin{aligned} \xi_t(i) &= [\psi_t, \bar{\pi}_t, v_{c,t}, v_t(i)]' \\ e_t(i) &= [\eta_t, \lambda_t, \nu_{c,t}, \nu_t(i)]' \\ y_t(i) &= [\pi_t, \tilde{s}_t(i), s_t(i)]' \\ u_t &= [\omega_t, \omega_{2,t}, \omega_{3,t}]'. \end{aligned}$$

Here $\Phi_t(\mathbf{i})$ is a $k \times k$ matrix which depends on ϕ_t , ρ_c and $\rho(i)$, where $k = 4$ is the number of state variables; $\mathbf{R}_t(\mathbf{i})$ is $k \times 4$ and depends on $\sigma_{\eta,t}$, $\sigma_{\lambda,t}$, $\sigma_{c,t}$, $\sigma_{\nu}(i)$; $\mathbf{D}(\mathbf{i})$ is a $3 \times k$ matrix and \mathbf{Q} is 3×3 and depends on σ_{ω} . π_t , $\tilde{s}_t(i)$ and $s_t(i)$ refer to the inflation, coordinating and idiosyncratic signals, respectively. The measurement errors $\omega_{2,t}$ and $\omega_{3,t}$ are just added for completeness but their variance is zero so they are irrelevant. The detailed definitions are as follows:

$$\Phi_t(\mathbf{i}) = \begin{bmatrix} \phi_t & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \rho_c & 0 \\ 0 & 0 & 0 & \rho(i) \end{bmatrix}, \quad \mathbf{R}_t(i) = \begin{bmatrix} \sigma_{\eta,t} & 0 & 0 & 0 \\ 0 & \sigma_{\lambda,t} & 0 & 0 \\ 0 & 0 & \sigma_{c,t} & 0 \\ 0 & 0 & 0 & \sigma_{\nu}(i) \end{bmatrix} \quad (14)$$

$$\mathbf{D}(\mathbf{i}) = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & \alpha(i) & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} \sigma_{\omega} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (15)$$

The Kalman filter recursion for forecaster i is given by:

$$\xi_{t|t-1}(i) = \Phi_t(\mathbf{i})\xi_{t-1|t-1}(i) \quad (16)$$

$$P_{t|t-1}(i) = \Phi_t(\mathbf{i})P_{t-1|t-1}(i)\Phi_t(\mathbf{i})' + \mathbf{R}_t(i)\mathbf{R}_t(i)' \quad (17)$$

$$s_{t|t-1}(i) = \mathbf{D}(\mathbf{i})\xi_{t|t-1}(i) \quad (18)$$

$$F_{t|t-1}(i) = \mathbf{D}(\mathbf{i})P_{t|t-1}(i)\mathbf{D}(\mathbf{i})' + \mathbf{Q}\mathbf{Q}' \quad (19)$$

$$\xi_{t|t}(i) = \xi_{t|t-1}(i) + \underbrace{P_{t|t-1}(i)\mathbf{D}(\mathbf{i})' [F_{t|t-1}(i)]^{-1}}_{K_t(i)} [y_t(i) - \mathbf{D}(\mathbf{i})\xi_{t|t-1}(i)] \quad (20)$$

$$P_{t|t}(i) = P_{t|t-1}(i) - P_{t|t-1}(i)\mathbf{D}(\mathbf{i})' [F_{t|t-1}(i)]^{-1} \mathbf{D}(\mathbf{i})P_{t|t-1}(i) \quad (21)$$

$K_t(i)$ is the Kalman gain of forecaster i which represents how sensitive beliefs about the different state variables are to new information captured by the three signals. In this rational

model the Kalman gain captures the optimal weight given to new information.

Re-arrange the Kalman filter recursions as follows to obtain equation:

$$\xi_{t|t}(i) = \xi_{t|t-1}(i) + K_t(i) [y_t(i) - \mathbf{D}(\mathbf{i})\xi_{t|t-1}(i)] \quad (22)$$

$$= [\mathbf{I}_4 - K_t(i) \mathbf{D}(\mathbf{i})] \Phi_t(\mathbf{i}) \xi_{t-1|t-1}(i) + K_t(i) y_t(i) \quad (23)$$

$$= [\mathbf{I}_4 - K_t(i) \mathbf{D}(\mathbf{i})] \Phi_t(\mathbf{i}) \xi_{t-1|t-1}(i) + K_t(i) [\mathbf{D}(\mathbf{i})\xi_t(i) + \mathbf{Q}u_t] \quad (24)$$

$$= [\mathbf{I}_4 - K_t(i) \mathbf{D}(\mathbf{i})] \Phi_t(\mathbf{i}) \xi_{t-1|t-1}(i) + K_t(i) [\mathbf{D}(\mathbf{i})(\Phi_t(i)\xi_{t-1}(i) + \mathbf{R}_t(\mathbf{i})e_t(i)) + \mathbf{Q}u_t] \quad (25)$$

Putting the trend-cycle model and the combined the signal extraction problems of all forecasters together gives our state space model of the econometrician.

The transition equation we use in our panel estimation reads

$$\begin{bmatrix} \xi_t \\ \vec{\xi}_{t|t} \\ \omega_t \end{bmatrix} = \tilde{\Phi}_t \begin{bmatrix} \xi_{t-1} \\ \vec{\xi}_{t-1|t-1} \\ 0 \end{bmatrix} + \tilde{\mathbf{R}}_t \begin{bmatrix} \eta_t \\ \lambda_t \\ \nu_{c,t} \\ \vec{\nu}_{v,t} \\ \omega_t \end{bmatrix} \quad (26)$$

where $\vec{\xi}_{t|t}$ and $\vec{\nu}_t$ are column vectors stacking $\xi_{t|t}(i)$ and $\nu_t(i)$ of every forecaster. Note that ξ_t contains the idiosyncratic noise processes for all forecasters, i.e.

$$\xi_t = \begin{bmatrix} \psi_t & \bar{\pi}_t & v_{c,t} & \vec{v}_t \end{bmatrix}'$$

The measurement equations for our panel estimation are

$$\begin{bmatrix} \pi_t^{cpi} \\ \psi_t^{est} \\ \bar{\pi}_t^{est} \\ \mathbb{E}_t \pi_t^{long}(1) \\ \mathbb{E}_t \pi_t^{long}(2) \\ \vdots \\ \mathbb{E}_t \pi_t^{long}(N) \end{bmatrix} = \begin{bmatrix} \mathbf{D}_{\mathbf{CPI}} & \mathbf{0}_{1 \times k} & \mathbf{0}_{1 \times k} & \dots & \mathbf{0}_{1 \times k} & \sigma_\omega \\ \mathbf{1}_1 & \mathbf{0}_{1 \times k} & \mathbf{0}_{1 \times k} & \dots & \mathbf{0}_{1 \times k} & 0 \\ \mathbf{1}_2 & \mathbf{0}_{1 \times k} & \mathbf{0}_{1 \times k} & \dots & \mathbf{0}_{1 \times k} & 0 \\ \mathbf{0}_{1 \times k} & \mathbf{1}_2 & \mathbf{0}_{1 \times k} & \dots & \mathbf{0}_{1 \times k} & 0 \\ \mathbf{0}_{1 \times k} & \mathbf{0}_{1 \times k} & \mathbf{1}_2 & \dots & \mathbf{0}_{1 \times k} & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0}_{1 \times k} & \mathbf{0}_{1 \times k} & \mathbf{0}_{1 \times k} & \dots & \mathbf{1}_2 & 0 \end{bmatrix} \begin{bmatrix} \xi_t \\ \xi_{t|t}(1) \\ \xi_{t|t}(2) \\ \vdots \\ \xi_{t|t}(N) \\ \omega_t \end{bmatrix}, \quad (27)$$

where $\mathbf{D}_{\mathbf{CPI}}$ is a zero row vector of length $N+k-1$ with elements 1 and 2 equal to 1. $\mathbf{1}_n$ denotes the $1 \times n$ row vector with elements all equal to zero except the n -th one which is equal to one. The observable variables in the vector on the left hand side of equation (27) include an empirical measure of inflation such as CPI inflation, π_t^{cpi} , our estimates of the cyclical component, ψ_t^{est} , trend inflation, $\bar{\pi}_t^{est}$, and an empirical measure of long-run inflation expectations of forecasters, $\pi_t^{long}(i)$.

Behavioral model

We allow each forecaster to have forecaster-specific views on the model parameters of the coordinating and idiosyncratic signal processes which can be different from the "true" estimates

based on the model described above. More specifically, we can denote these parameters with $*$ so that we have the following additional parameters: $\rho_c^*(i), \rho^*(i), \sigma_\nu^*(i)$ and $\alpha^*(i)$. These parameters only enter the Kalman filter recursion for each forecaster and thereby also influence the matrix $\tilde{\Phi}_t$ in the transition equation defined in equation (26). $\tilde{\mathbf{R}}_t$ and the measurement equation remains unchanged. More specifically, in equations (16)-(20) the forecaster's beliefs about the parameters are used instead of the "true" estimates so that

$$\begin{aligned}\xi_{t|t}(i) &= \xi_{t|t-1}(i) + K_t^*(i) [y_t(i) - \mathbf{D}^*(\mathbf{i})\xi_{t|t-1}(i)] \\ &= [\mathbf{I}_4 - K_t^*(i) \mathbf{D}^*(\mathbf{i})] \Phi_t^*(\mathbf{i})\xi_{t-1|t-1}(i) + K_t^*(i) [\mathbf{D}(\mathbf{i})(\Phi_t(\mathbf{i})\xi_{t-1}(i) + \mathbf{R}_t(\mathbf{i})e_t(i)) + \mathbf{Q}u_t]\end{aligned}$$

Note that the variables with $*$ do not appear in $y_t(i)$ since these are the true signals. $K_t^*(i)$ is the *subjective* Kalman gain by forecaster i . As before this represents how sensitive beliefs about the different state variables are to new information captured by the three signals. Differently from the rational model, the parameter misperceptions imply that this Kalman gain is subjective and does not necessarily capture the optimal weight given to new information.

A.2 Initial conditions for estimation

Initial conditions of trend-cycle model

We initialize the state equations of our trend-cycle model as described in equations (1)-(3) as follows: $\epsilon_0 = 0$, $\bar{\pi}_0 = 0$. The initial uncertainty is set to the unconditional variance for ϵ_0 and to 1e6 for $\bar{\pi}_0$. The time-varying parameters are initialized as $\ln(\sigma_{\eta,1}^2) \sim N(0, \kappa)$, $\ln(\sigma_{\lambda,1}^2) \sim N(0, \kappa)$, $\phi_1 \sim N(0, \kappa)$ where $\kappa=1e6$.

Initial conditions of forecaster model

As usual in the literature on factor models the relative scale of loadings and factors is indeterminate and requires a normalization. We set $\log(\sigma_{c,0}^2) = 0$.

We assume that for each forecaster the initial variance-covariance matrix $P_{0|0}(i)$ in equation (17) is at the "steady-state"³² level given the initial parameter values. To compute this steady-state matrix we start from the following matrix

$$P_{0|0}(i) = \begin{bmatrix} \frac{\sigma_{\eta,1991Q2}^2}{1-\phi_{1991Q2}^2} & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{1-\rho_c^2} & 0 \\ 0 & 0 & 0 & \frac{\sigma_\nu^2(i)}{1-\rho(i)^2} \end{bmatrix} \quad (28)$$

and we iterate over equations (17), (19), (21) to get the "steady-state". The values from 1991Q2 are used since our first period of the panel model is 1991Q3.

³²Intuitively, this implies that forecasters have already received some signals before and are not completely uninformed about the history of inflation. Alternatively, we could assume that forecasters have never received any signals before, including inflation, and impose a wide/uninformative initial uncertainty. This would lead to large initial spikes in the Kalman gains since forecasters react a lot to the first few signals but results for rest of the sample are little affected.

Initial conditions of panel model

The initial conditions for the first part of the transition equation defined in equation (26) are as follows:

$$\xi_0 \equiv \begin{bmatrix} \epsilon_{1991Q2} \\ \bar{\pi}_{1991Q2} \\ 0 \\ \mathbf{0}_{N \times 1} \end{bmatrix} \quad (29)$$

$$\mathbf{P}_0 \equiv \begin{bmatrix} 0 & 0 & 0 & \mathbf{0}_{1 \times N} \\ 0 & 0 & 0 & \mathbf{0}_{1 \times N} \\ 0 & 0 & \frac{1}{1-\rho_c^2} & \mathbf{0}_{1 \times N} \\ \mathbf{0}_{N \times 1} & \mathbf{0}_{N \times 1} & \mathbf{0}_{N \times 1} & \mathbf{I}_{N \times N} \frac{\sigma_v^2}{1-\rho^2} \end{bmatrix} \quad (30)$$

The zeros in the upper left 2×2 part mean that we do not want the panel estimation to change the initial estimate of trend and cycle that we got from the time series estimation.

Denoting with $\xi_{0|0}^e(i)$ the initial condition of the econometrician for the state estimate of forecaster i , we assume

$$\xi_{0|0}^e(i) \equiv \begin{bmatrix} \epsilon_{1991Q2} \\ \bar{\pi}_{1991Q2} \\ 0 \\ 0 \end{bmatrix}, \quad \forall i \quad (31)$$

ω_0 is set in line with the estimate of the i.i.d. component for 1991Q2 again assuming zero uncertainty since we do not want the panel estimation to change the initial estimates of the trend-cycle model.

The initial covariance matrix is based on deriving the variance of $\xi_{t|t}(i)$ and is given by

$$\Sigma_t(i) = \text{var}(K_t(i)y_t(i)) \quad (32)$$

$$= K_t(i)\text{var}(\mathbf{D}(\mathbf{i})\xi_t(i) + \mathbf{Q}u_t)K_t(i)' \quad (33)$$

$$= K_t(i)[\mathbf{D}(\mathbf{i})\text{var}(\xi_t(i))\mathbf{D}(\mathbf{i})' + \mathbf{Q}\mathbf{Q}']K_t(i)' \quad (34)$$

$$= K_t(i)[\mathbf{D}(\mathbf{i})P_{t|t}(i)\mathbf{D}(\mathbf{i})' + \mathbf{Q}\mathbf{Q}']K_t(i)' \quad (35)$$

We can evaluate this at time 0 using $P_{0|0}(i)$ and the corresponding $K_0(i)$ as defined above, so that

$$\mathbf{P}_{0|0}^e(i) \equiv \Sigma_0(i), \quad \forall i \quad (36)$$

The remaining elements of the initial covariance matrix for equation (26) are assumed to be zero.

B Selection of forecasters

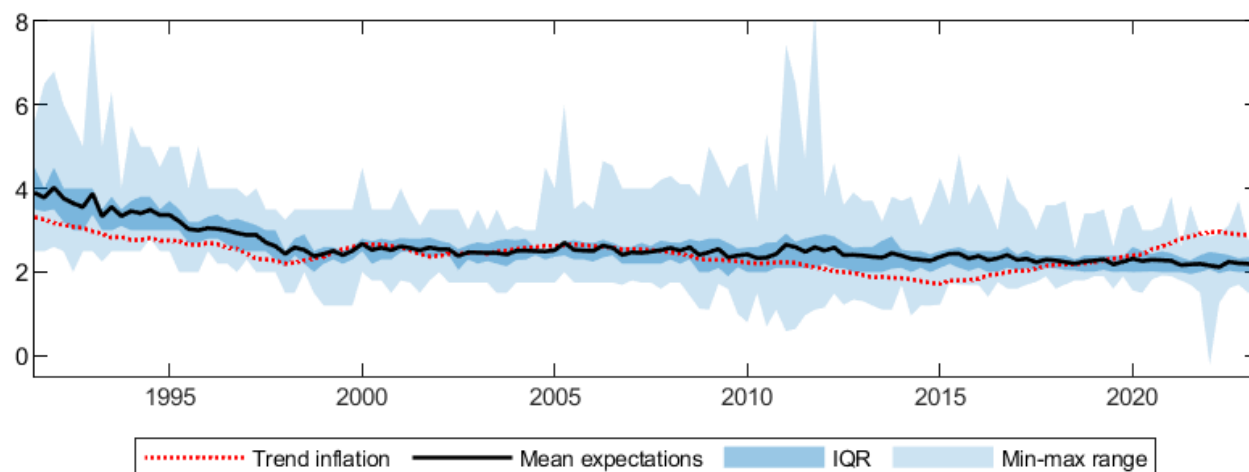


Figure 8: **Trend inflation and full sample of long-run inflation expectations.** Trend component of U.S. quarterly headline CPI inflation rate estimated using the trend-cycle model (red dotted line) and the mean (black solid line), the interquartile range (the dark blue bands), and the min-max range (the light blue bands) of the distribution of SPF long-run CPI inflation expectations including all forecasters in the survey.

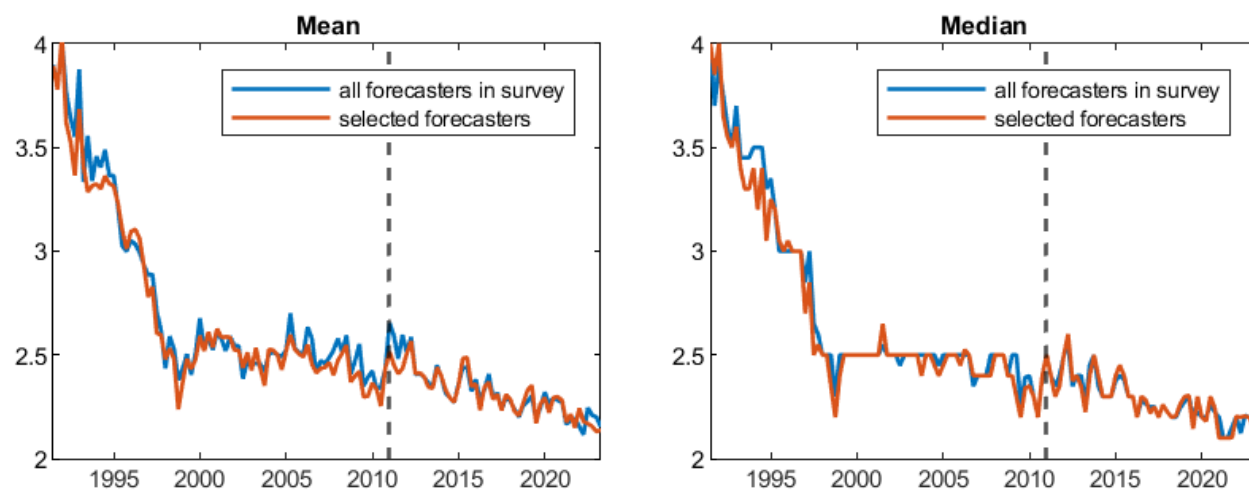


Figure 9: **Time series of aggregate inflation expectations.** The dashed vertical line indicates 2011Q1 before which we use 10-year average expectations and afterwards 5-year 5-year expectations.

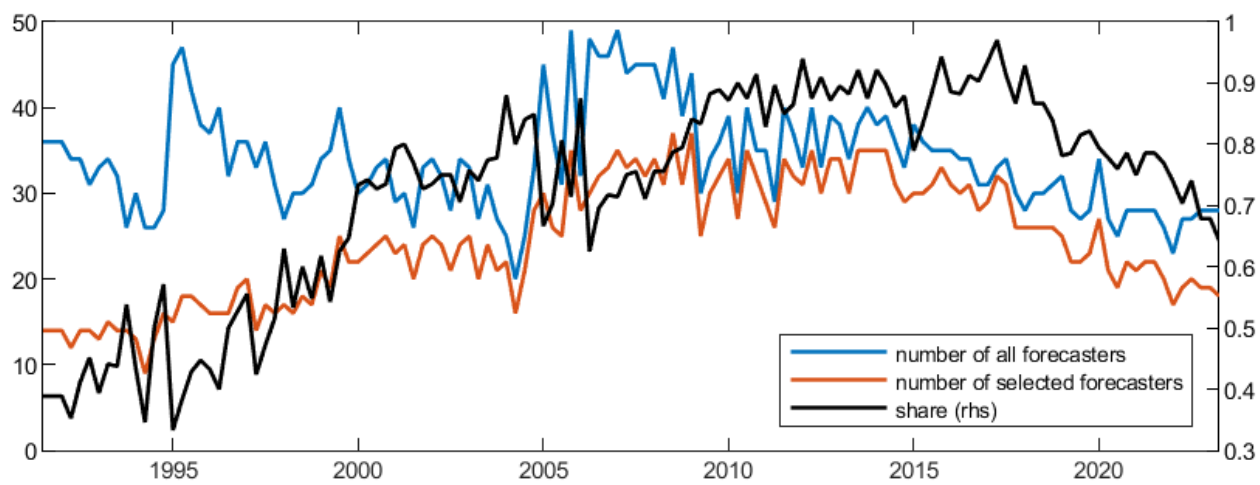


Figure 10: Number of total and selected forecasters in SPF.

C The estimated trend-cycle model

The priors and posterior mean for the time-invariant parameters of the trend-cycle model are shown in Table 2.³³ The posterior mean of the time-varying parameters are shown in Figure 11. The volatility of the trend component of inflation peaks during the Great Inflation at 44 basis points (bps) and again in the early 1980s. Trend inflation volatility declines remarkably quickly over the ensuing 20 years, leveling off around 25 bps around 2000. In the most recent period, the volatility of trend inflation has been increasing again. Nevertheless, the recent rise is not nearly as quick as the surge observe in the 1960s and early 1970s. We estimate that the volatility of trend inflation is just a touch below 30 bps at the end of our sample period (the second quarter of 2023.)

Table 2: **Prior and posterior moments for the trend-cycle model parameters**

Parameters	Prior moments					Posterior moments	
	Distr.	Shape	Scale	Mean	5%-95% Range	Mean	5%-95% Range
γ_η^2	IG	5	0.040	0.010	[0.0040,0.0200]	0.054	[0.0185,0.1174]
γ_λ^2	IG	5	0.040	0.010	[0.0040,0.0200]	0.011	[0.0045,0.0276]
γ_ϕ^2	IG	5	0.004	0.001	[0.0004,0.0020]	0.002	[0.0006,0.0032]
σ_ω^2	IG	3	0.200	0.100	[0.0320,0.2450]	0.143	[0.0382,0.3524]

Notes: IG stands for inverse gamma distribution. The shape and scale are parameters of that distribution. The posterior mean and the interquantile range are computed via *Gibbs Sampling*.

The volatility of cyclical inflation is much larger than that of the trend component. It peaks in 1980 and again during the Great Recession. Cyclical volatility of inflation fell very sharply after reaching its second peak, reaching a trough at around 1.6 percent in the 2010s. Interestingly, cyclical volatility started increasing before the onset of the pandemic. In the last quarters of the sample cyclical volatility slipped slightly below its third peak 2.5 percent reached during the pandemic crisis.

The persistence of the cyclical component starts increasing in the 1960s and peaks toward the end of the Great Inflation, topping out around 0.8. Persistence falls gradually thereafter and settles at about 0.30 in the middle of the 2000s. Since the onset of the Great Recession, cyclical persistence has been rising, following a pattern that resembles the volatility of trend inflation on the left panel. The persistence of cyclical inflation is close to 0.55 at the end of the sample period. The estimates shown in Figure 11 are very similar to what found by previous works in the literature—e.g Chan et al. (2018)—and overall seem to provide a plausible characterization of inflation over the last 60 years.

Table 3 shows the forecasting performance of the estimated trend relative to the trend estimated with alternative prior or model specifications. The evaluation is based on the root-mean squared error between the trend and long-run realized CPI inflation. First,

³³We describe how we initialize the state vector in Appendix A.2. Following Chan et al. (2018) and Stock and Watson (2007) we assume the initial value of the trend is zero and allow a very wide initial uncertainty so the trend in the second period of the sample is relatively unconstrained. The stochastic volatilities are approximated following the approach by Kim et al. (1998) and using the more accurate 10-component Gaussian mixture approximation proposed by Omori et al. (2007).

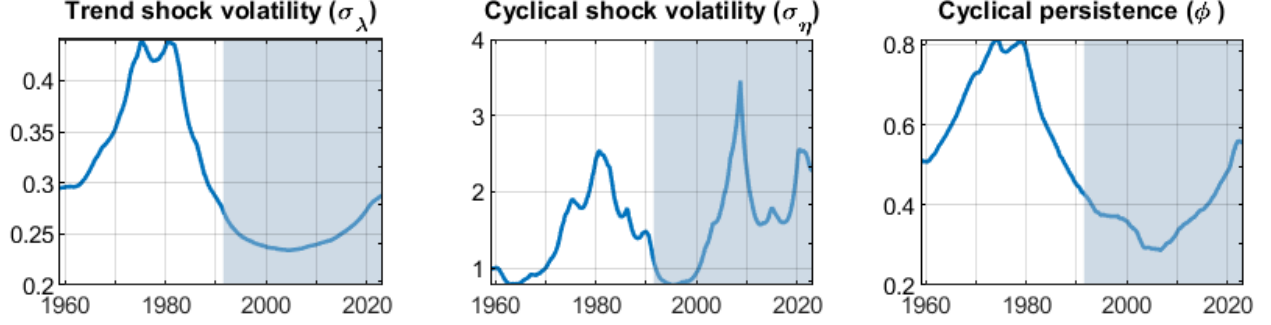


Figure 11: **Posterior mean for the time-varying parameters of the trend-cycle model.** The estimation of the three time-varying moments is conditional on quarterly U.S. headline CPI inflation rate observed over the period 1959Q1–2023Q2. Trend volatility and cyclical volatility are expressed in percentage. The shaded blue area indicates the sample period for the second-step panel estimation (1991Q3–2023Q2).

Table 3: **Forecasting performance of estimated trend**

	Full sample		SPF panel sample	
	10Y	5Y5Y	10Y	5Y5Y
<i>I. Alternative priors</i>				
Baseline	1.00	1.00	1.00	1.00
Wider prior	0.97	1.02	1.08	1.05
2x prior mean	1.00	1.00	1.01	1.00
5x prior mean	0.97	1.01	1.00	1.00
<i>II. Alternative model specifications</i>				
No noise component	0.98	1.00	1.00	1.00
Constant parameters (tight prior)	1.20	1.24	1.48	1.26
Constant parameters (wide prior)	1.11	1.19	1.30	1.17
No persistence in cyclical component	1.23	1.25	1.14	1.09

Notes: Root mean squared errors (RMSE) relative to baseline model. Root mean squared errors of different trend estimates and realized inflation. The first two columns show the results over the full sample from 1959 onwards and the last two rows for the SPF panel sample from 1991Q3 onwards. In both cases the first column shows the RMSE with respect to 10-year average CPI inflation and the second column with respect to average inflation between 6 and 10 years ahead. All RMSE are scaled relative to the baseline model so that a value lower than 1 indicates a better forecasting performance than the baseline model. Wider prior (second row) uses a shape parameter of 1.5 instead of 5 in the IG prior distribution. Rows 3 and 4 adjust the scale parameter such that the prior mean is twice or five times larger, respectively. The 5th row is based on a model without the noise component ω . The 6th and 7th row report the RMSE for our baseline model but assuming parameters are constant. The 8th row reports the RMSE for the univariate model from Stock and Watson (2016) estimated using CPI headline inflation. This model allows for stochastic volatility and outlier adjustments but the cyclical component has no persistence.

the relative differences in the root-mean squared error are very small for alternative prior specifications, highlighting the robustness of our prior choice. Second, using a model with constant parameters leads to significantly larger forecast errors underscoring the importance to allow for time-variation in parameters even if this complicates the modelling framework.

Third, modelling the cyclical component as iid - as for example done in Stock and Watson (2016) - leads to a more volatile trend estimate that removes high-frequency movements in current inflation but is less suited for forecasting inflation at longer horizons.

D Priors and posterior modes of the models' parameters

Table 4: **Prior distribution and posterior mode for the parameters**

Parameters	Prior moments					Posterior mode	
	Distr.	Par(1)	Par(2)	Mean	5%-95%	Rational	Behavioral
ρ_c	B	2.63	2.63	0.50	[0.17,0.83]	0.99	0.96
$\alpha(i)$	IG	3.00	1.00	0.50	[0.16,1.22]	Figure 2	Figure 2
$\rho(i)$	B	2.63	2.63	0.50	[0.17,0.83]	Figure 2	Figure 2
$\sigma_\nu(i)$	IG	3.00	1.00	0.50	[0.16,1.22]	Figure 2	Figure 2
$\rho_c^*(i)$	B	2.63	2.63	0.50	[0.17,0.83]	NA	Figure 3
$\alpha^*(i)$	IG	3.00	1.00	0.50	[0.16,1.22]	NA	Figure 3
$\rho^*(i)$	B	2.63	2.63	0.50	[0.17,0.83]	NA	Figure 3
$\sigma_\nu^*(i)$	IG	3.00	1.00	0.50	[0.16,1.22]	NA	Figure 3

Notes: B and IG stand for beta, and inverse gamma distributions. Par(1) and Par(2) are the shape and scale for the inverse gamma distribution, respectively, and are the shape parameters α and β for the beta distribution.

In Table 4, we summarize prior moments and the posterior mode of the models' parameters. Priors are fairly uninformative, reflecting the beliefs that the SPF data should be primarily driving the setting of the parameters values. For the time-varying volatility of coordinating innovations, $\sigma_{c,t}$, we normalize the starting value $\ln \sigma_{c,0}$ at zero. This is a standard normalization since the initial value $\sigma_{c,0}$ cannot be separately identified from the scale of the forecaster-specific $\alpha(i)$ (Del Negro and Otrok, 2008). The parameter governing the scale of innovations in the random walk process for $\ln \sigma_{c,t}^2$, which was denoted with γ_c , is set to 0.5. In Table 4, the estimation drives the auto-correlation of the coordinating noise process, ρ_c , to a value very close to one, implying that it follows a near-unit root process.

E Empirical performance of the rational model

Figure 12 shows the impulse response functions to all four shocks comparing the rational and behavioral model.

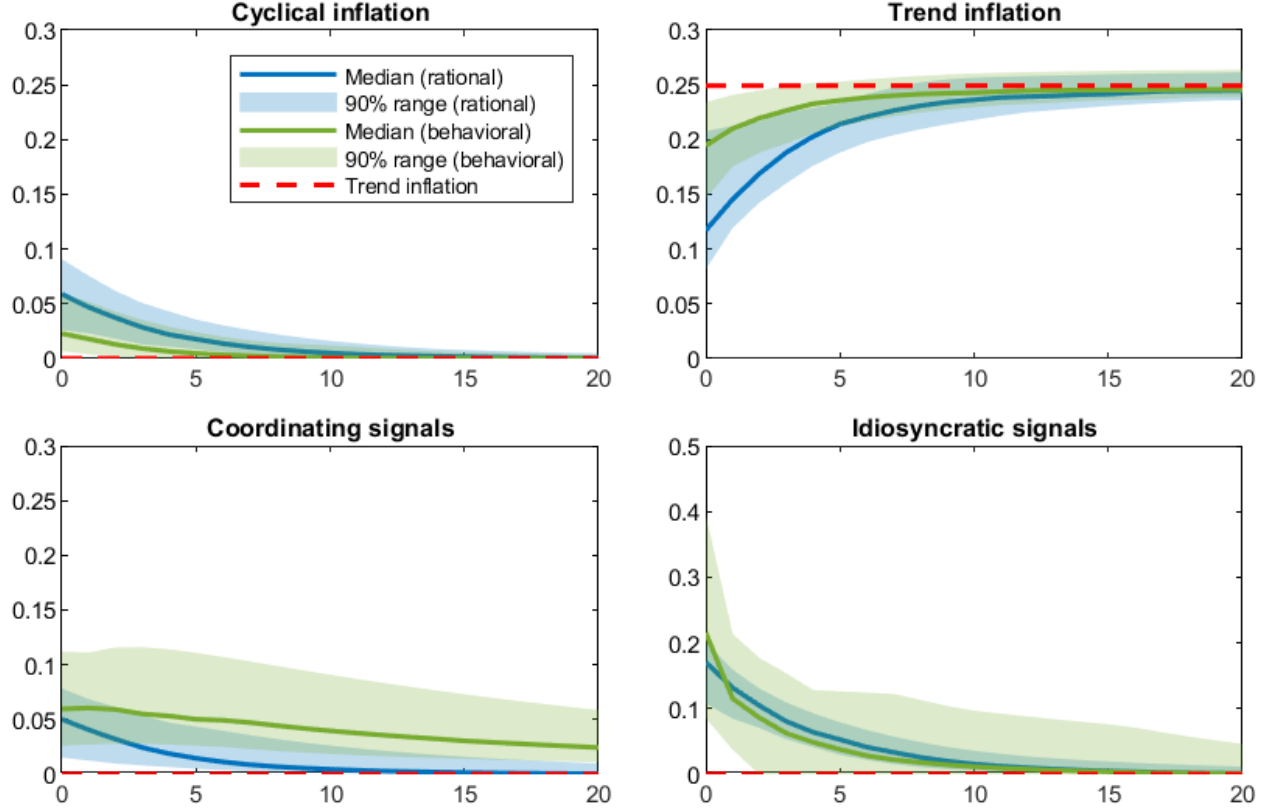


Figure 12: **Propagation of shocks to expectations.** Impulse response functions of long-run inflation expectations of every forecasters to a one-standard-deviation cyclical inflation shock (top left chart), trend inflation shock (top right chart), noise innovation to the coordinating signal (bottom left chart) and to every idiosyncratic signal (right chart). Responses in the rational model are denoted in blue and those in the behavioral model are denoted in green. The solid lines denote the median response across forecasters and the shaded areas denote the 90 percent range of responses across forecasts. The red dashed line shows the true response of trend inflation. The impulse response functions in both graphs are computed in every period of the sample (1991Q3-2023Q2) and then averaged across sample periods.

The following figures show some robustness to the misspecification of the rational model. The two dimensions of misspecification are the serial correlation of the coordinating signal innovations and the deviation from Gaussianity in the distribution of the idiosyncratic signal innovations. In our baseline model the trend comes from an estimated trend cycle model. Instead, in Figure 13 we take the same rational model but assume that the trend is simply equal to a 5-year (lhs) or 10-year (rhs) moving average of headline CPI inflation. This assumption about the trend imposes less assumptions and at the same time still captures the underlying slow-frequency movements in inflation. In the left panel of Figure 14, we estimate the trend using core CPI inflation instead of headline CPI inflation. In all three cases the trend and the parameters of the trend-cycle model are based on the full-sample information and we assume forecasters know the parameters. In the rhs panel of 14 we estimate the

trend-cycle model with headline CPI inflation in real-time so that the parameters do not only change over time but also with every vintage. The idea here is to mimic forecasters learning the parameters over time and to see whether this can fix the misspecification of the baseline rational model. In all four cases the rational model still exhibits the same two dimensions of misspecification.

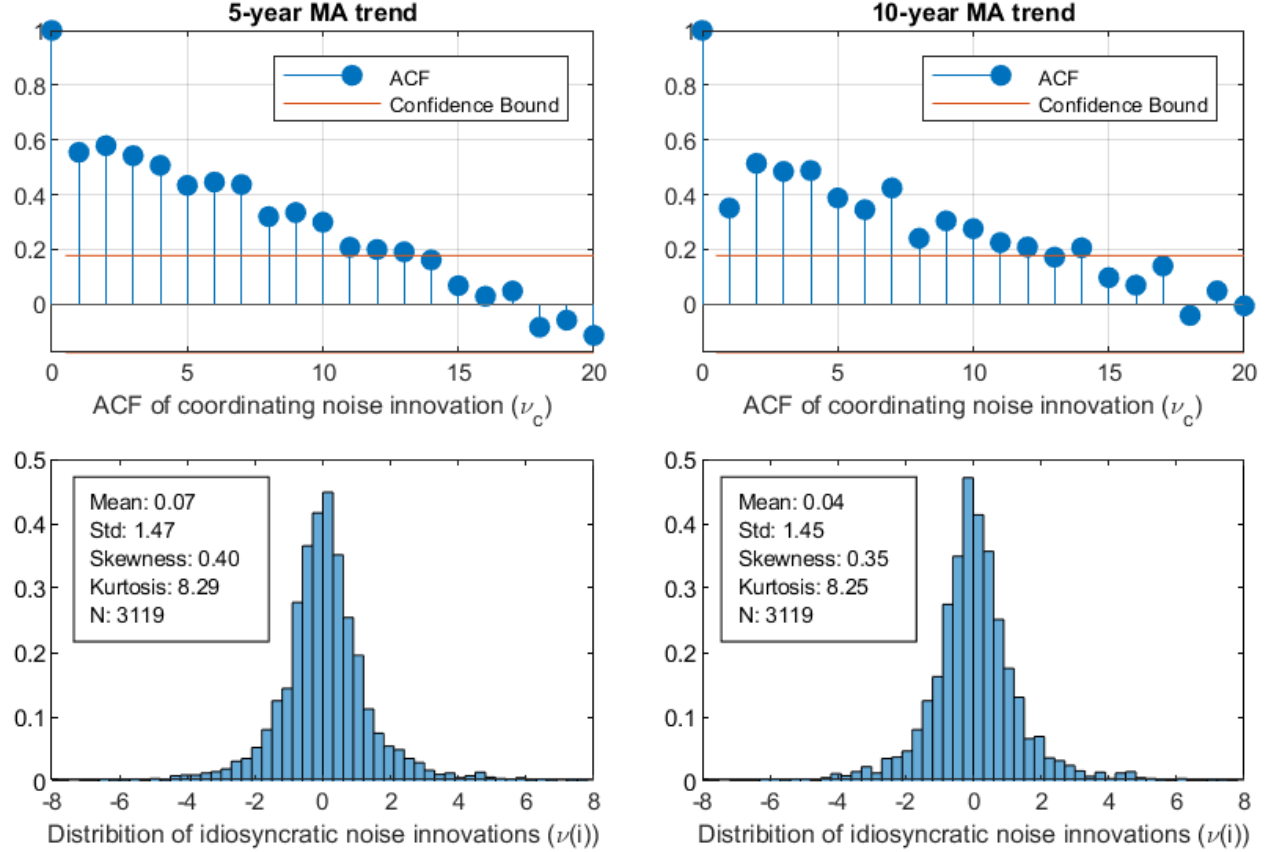


Figure 13: **Rational model: assessing the misspecification under moving average trend.** The upper charts show the autocorrelogram of the estimated in-sample innovation to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The lower charts show the distribution of all the estimated in-sample innovations to the idiosyncratic signals over time and across forecasters. The figures correspond to the rational model and the true trend being a 5-year moving average of CPI inflation (lhs) and a 10-year moving average of CPI inflation (rhs), respectively.

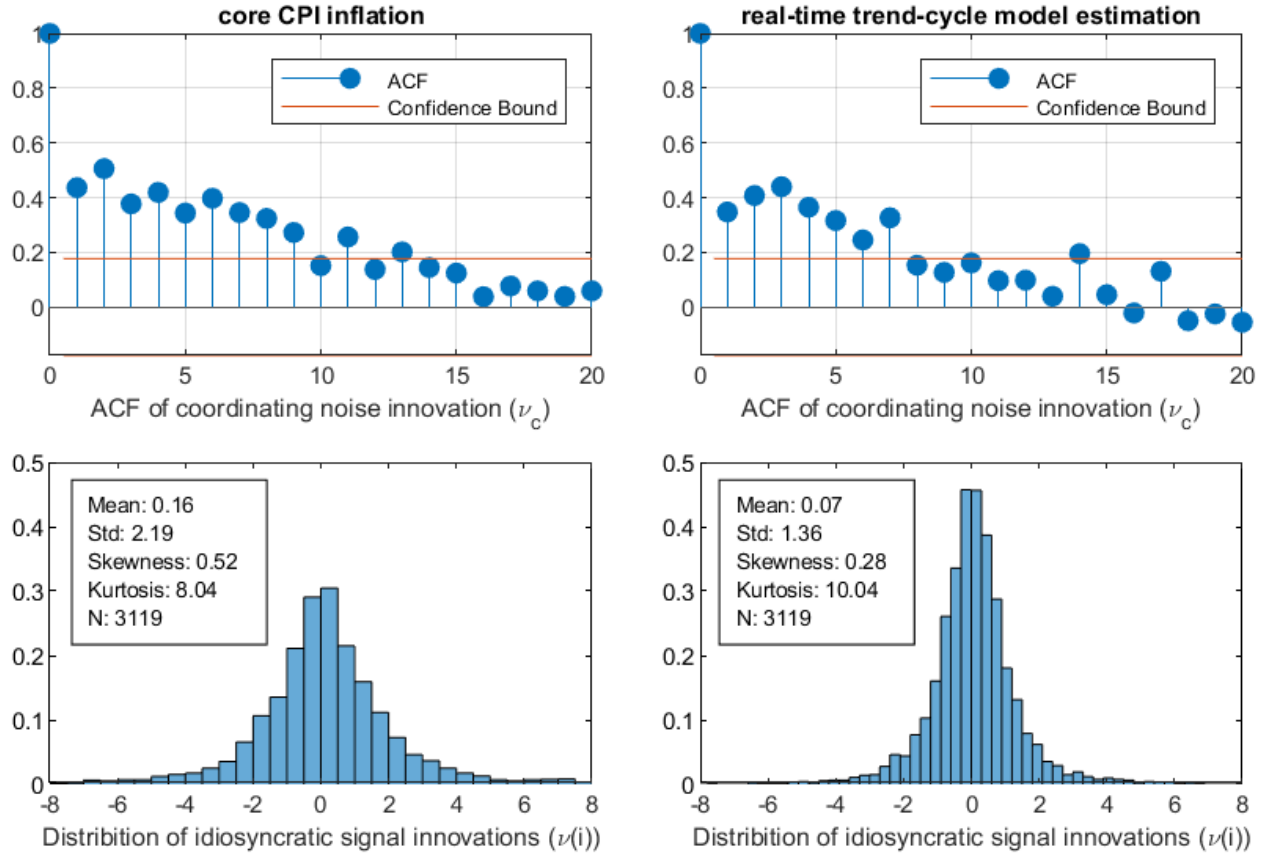


Figure 14: **Rational model: assessing the misspecification under core inflation and real-time estimation.** The upper charts show the autocorrelogram of the estimated in-sample innovation to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The lower charts show the distribution of all the estimated in-sample innovations to the idiosyncratic signals over time and across forecasters. The lhs figure corresponds to the rational model with core CPI inflation instead of headline CPI inflation. The rhs figure is based on a real time estimation of the trend-cycle model.

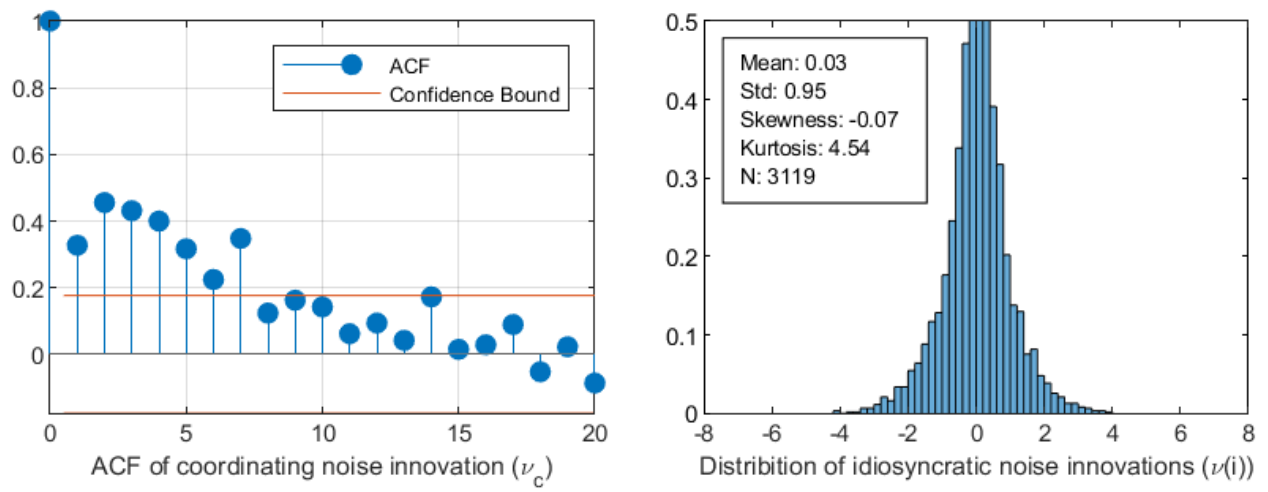


Figure 15: **Behavioral model only with overconfidence: assessing the misspecification.** The left chart shows the autocorrelogram of the estimated in-sample innovation to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The right chart shows the distribution of all the estimated in-sample innovations to the idiosyncratic signals over time and across forecasters. The model only allows for overconfidence, i.e. $\sigma_\nu^*(i) \neq \sigma_\nu(i)$ is allowed, while for all other parameters in the signal extraction problem there is no deviation from rationality.

F Anchoring-compatible inflation path derivation

In the following we describe the detailed procedure underlying the exercise in section 9. We take the perspective of a FOMC policy maker in December 2022 so that the latest available data is 2022Q3 for inflation and the SPF survey round from 2022Q4 (answered in November 2022).

- (i) **Initial guess for trend inflation:** We use the median path for PCE inflation from the December 2022 Summary of Economic Projections (SEP). For each year from 2023-2025, the forecasts refer to the year-on-year growth rate of the fourth quarter. We apply these year-on-year growth rates to the CPI index and linearly interpolate the missing quarters. Then we compute quarterly annualized growth rates and add 30 bps to be consistent with CPI inflation being on average 30 bps higher than PCE inflation. Finally, we apply a 4-quarter moving average to that projection path to avoid jumps generated by the fact that we only have one SEP projection per year. From 2026 onwards we assume PCE inflation is at 2 percent and CPI inflation at 2.3 percent accordingly. Note that is just an initial guess and in principle any inflation path can be used. Then, we use this inflation path in our trend-cycle model to estimate the inflation trend going forward from 2022Q4.
- (ii) **Estimate of inflation path:** We obtain estimates of the inflation path from 2022Q4-2025Q4³⁴ by applying the Kalman smoother to the state-space model as defined in equations (26)-(27) except that the measurement equations are modified as follows:
 - Inflation: For 2022Q4 we impose actual realized but afterwards the path of inflation is treated as missing. The assumption here is that at the time of the FOMC meeting in mid December the quarter is almost over and inflation in that quarter cannot really be influenced anymore.
 - Trend inflation: We set trend inflation equal to the initial guess from step (i)
 - Iid component of inflation: We set the i.i.d. component to zero
 - Expectations: Individual inflation expectations are not available and since we want to keep expectations stable at their current value we impose that average inflation expectations are equal to the average value from the 2022Q4 SPF round (see Figure 16 for the path of stable expectations and realized average expectations since 2022Q4). Note, this is only imposed for forecasters who have been in the sample during the two years prior to December 2022.
 - Idiosyncratic signals: We set the innovations to idiosyncratic signals to zero. The idea is to focus only on the role of coordinating sig vs inflation in this exercise.

The state-space model is initialized using the estimated states in 2022Q3 from our baseline estimation. The time-varying parameter $\sigma_{c,t}$ is fixed at the last estimate from 2022Q3 and the time-invariant parameters are kept as in our baseline estimation.

³⁴In the long run, i.e. after 2025Q4, we assume trend inflation is equal to average expectations and the cyclical+i.i.d. components are zero.

- (iii) The inflation path from the previous step is not necessarily consistent with the initial guess for trend inflation. Therefore, we apply the Kalman smoother on the trend-cycle model going forward from 2022Q4 to obtain a new trend estimate that is consistent with the inflation path from step (ii).³⁵ With this new trend estimate we restart step (ii) and solve for a "fixed point" by iterating over these steps until convergence in the inflation path.

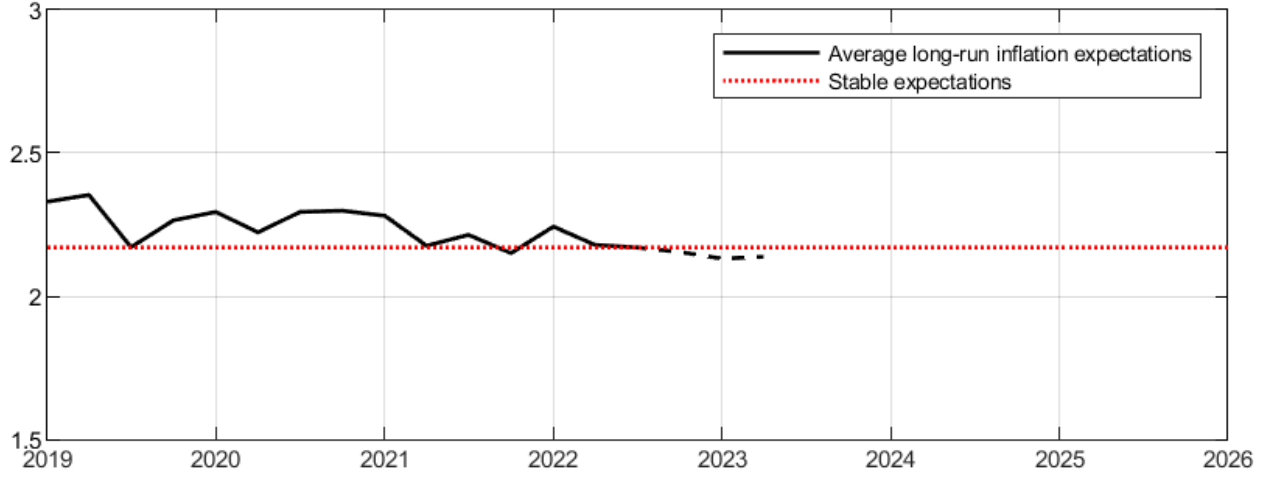


Figure 16: **Average long-run inflation expectations.** The black solid line shows average long-run inflation expectations until 2022Q4 when we start the exercise. The dashed black line shows the realizations of average SPF long-run inflation expectations observed since 2022Q4. The red dotted line corresponds to the level of stable average expectations from 2022Q4.

³⁵For simplicity we assume that the trend estimate before 2022Q4 remains unchanged and we keep all the parameters fixed at their 2022Q3 values. In line with (ii) we impose that the i.i.d. component is zero.