Long-Run Inflation Expectations*

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

August 2026

Abstract

Professional forecasters' long-run inflation expectations overreact to news and exhibit persistent, predictable forecast errors. We develop an imperfect-information model in which forecasters infer long-run inflation from observed inflation and forward-looking public and private information. Using individual forecasts from the U.S. Survey of Professional Forecasters, we identify two departures from rationality required by the data: overconfidence in private information, which explains overreaction, and persistent expectations bias, which accounts for the persistent, predictable forecast errors. We also uncover 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 real-time framework for assessing whether policymakers' communicated inflation paths are consistent with long-run expectations remaining anchored at the central bank's target.

Keywords: *Panel survey data, 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, 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, University of Ottawa, Richmond Fed, 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.

‡European University Institute and CEPR.

§De Nederlandsche Bank.

1 Introduction

Long-run expectations are central to modern macroeconomics and finance. Equity prices depend on beliefs about the long-run growth of future dividends. Long-dated government bond yields are shaped by beliefs about long-run real interest rates and inflation. Central banks use long-run inflation expectations as a key diagnostic of whether their inflation target remains credible. A common way to measure these expectations is through surveys, in particular panels of professional forecasters such as the *U.S. Survey of Professional Forecasters* (SPF), Blue Chip, and Consensus Economics.

We develop an imperfect information model to analyze the individual behavior of long-run expectations in such panel data settings. In the model, long-run expectations are beliefs about a persistent, slow-moving trend that individual forecasters must infer from limited information and may process differently and imperfectly. We take the model to the long-run inflation expectations of individual forecasters in the SPF. Two salient features characterize the data: forecast errors are highly persistent, often lingering for years without being corrected, and forecasters display a widespread tendency to overreact to news.

The model can explain these features with two departures from rationality: expectations being persistently biased — a phenomenon we refer to as *persistent expectations bias* — which accounts for the serial correlation in forecast errors and *overconfidence in private information*, which explains the overreaction to news.[1] 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 (ELB) on nominal interest rates. The model lends itself to real-time, out-of-sample predictions of the path of inflation that is consistent with the anchoring of long-run inflation expectations at the central bank's target. We present an application of the model to the recent post-pandemic bout of inflation in the US.

Forecasters solve a dynamic signal-extraction problem in which inflation follows a trend-cycle model. Each forecaster seeks to learn about the trend component, the model's counterpart to long-run inflation, which we estimate from inflation data. We therefore evaluate forecast errors against estimated trend inflation rather than realized future inflation. This helps overcome a key empirical obstacle in studying the rationality of long-run expectations. Because these expectations are available over a relatively short sample to begin with, evaluating them against future inflation requires discarding a substantial fraction of the already limited time series. This is particularly costly for inflation expectations, as it excludes the

[1] Under certain conditions, there is a direct mapping between overconfidence and diagnostic expectations—a widely used framework to explain overreaction in the context of short-run expectation formation. We discuss the formal link between the two in Appendix H and characterize the conditions under which they coincide.

post-pandemic episode, when inflation departed sharply from its behavior over the preceding decades. The difficulty of overcoming this limitation has likely contributed to the scant evidence on the rationality of long-run expectations.

Forecasters observe three sources of information about the trend: the quarterly inflation rate, a *coordinating signal*, and an *idiosyncratic signal*. Treating inflation as a signal allows us to estimate the sensitivity of long-term inflation expectations to short-term inflation movements, a measure of anchoring proposed by Bernanke (2007). The coordinating signals represent a generalization of a public signal by allowing for heterogeneity in how forecasters respond to public news. Since these signals are correlated across forecasters, the model can explain co-movement in forecasters' expectations. We interpret these signals as representing forward-looking public information, such as policymakers' communications or widely followed media, that can influence forecasters' beliefs about the central bank's commitment to price stability. The idiosyncratic signals capture individual judgments or private information, helping account for heterogeneity in forecasters' expectations.

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.

We begin by estimating the trend-cycle model using quarterly inflation data from the first quarter of 1959 (1959Q1) to 2023Q2. We then estimate the signal extraction parameters using our panel of forecasters (which starts in 1991Q3), conditioning on the parameters of the trend-cycle model. This second step uses the inflation time series, its estimated trend and cyclical components from the first step, and the panel structure of SPF long-run inflation expectations.[2]

This estimation design also gives the signals their *forward-looking* interpretation. The latent trend is estimated using the entire inflation series, including observations that become available after expectations are formed, whereas forecasters operate in real time using current inflation and noisy coordinating and idiosyncratic signals about that trend. The latter can

[2]As a robustness, we jointly estimated the parameters of the trend-cycle model and the signal extraction problem. Our main conclusions are robust. A description of this alternative estimation approach is in Appendix G.

therefore contain information about low-frequency inflation movements not yet reflected in realized inflation. This provides a tractable way to incorporate *forward-looking information* into noisy-information models used to evaluate the rationality of expectations.

The model with rational forecasters is clearly rejected by the data and reveals two major types of misspecification. First, the in-sample estimates of the innovations to the coordinating signals' auto-regressive noise turn out to be serially correlated, which is a violation of the model's assumptions. This serial correlation 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' noise 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. These misperceptions are analyzed and evaluated within a specification of distorted beliefs that remains agnostic about their microfoundation, encompassing several mechanisms proposed in the literature. 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 assumption allows the model to overcome the aforementioned misspecification, revealing two major deviations from rationality. First, in the estimated behavioral model, forecasters underestimate the persistence of noise in the coordinating signal. This cognitive bias implies a persistent expectations bias, which in turn accounts for the highly persistent forecast errors observed in long-run inflation expectations. 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.

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.[3] 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 types of misspecification.

While not targeted in estimation, the behavioral model closely matches the coefficients from the regression introduced by Coibion and Gorodnichenko (2015) and adapted by Bordalo et al. (2020) to individual forecasters. This offers an external validation of the model's behavioral features in a setting with minimal parametric restrictions. Moreover, this regression analysis confirms that persistent forecast errors and overreaction to news are features of long-run inflation expectations. Our model accounts for these features via the 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.[4] This confidence in the Fed is strong not only on average but also in the SPF cross-section and appears to strengthen after the 1990s, consistent with Hazell, Herreño, Nakamura and Steinsson (2022), who estimate the New Keynesian Phillips curve by exploiting regional (state-level) variation in inflation and unemployment.

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 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.

We use the behavioral model to derive the inflation path that ensures average long-run expectations are at the Fed's target in the post-pandemic period. We call this path the

[3] Stark (2013) documents that forecasters use different types of forecasting models and mostly combine both models and subjective considerations in reporting their projections, rather than just models. The important role of judgment has also been documented for the European Central Bank's Survey of Professional Forecasters where 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)).

[4] This property aligns closely with Bernanke (2007)'s definition of anchored inflation expectations and has been tested extensively, e.g. Gürkaynak, Levin, Marder and Swanson (2007), Dräger and Lamla (2014), Corsello, Neri and Tagliabracci (2021), Barlevy, Fisher and Tysinger (2021), Armantier, Sbordone, Topa, Van der Klaauw and Williams (2022).

anchoring-compatible inflation path. This application illustrates how the model can inform central bank decisions regarding the pace of disinflation and how quickly projected inflation must fall to prevent expectations from becoming deanchored during a high-inflation episode and its aftermath. Because trend inflation is nonstationary, the model can predict persistent deviations of long-run expectations from target when inflation is announced to remain above target for too long.

Our model shows that, as of late 2022, average long-run inflation expectations remained anchored at the target despite elevated trend inflation, due to a persistent downward bias in professional forecasts. This bias, possibly reflecting the Fed's strong reputation as an inflation fighter, eliminates the need to undershoot the inflation target to keep expectations at the Fed's target. Notably, both the inflation projections of the participants at the December 2022 Federal Open Market Committee (FOMC) meeting and the realized inflation path closely align with the model's anchoring-compatible trajectory.

Crucially, however, this outcome is conditional on the prevailing economic conditions. If inflation rises again in the coming quarters and forecasters begin to internalize this renewed increase in trend inflation, the model dynamically adjusts, calling for a faster and deeper disinflation path to preserve anchoring. In that case, tighter monetary policy — even at the cost of temporarily undershooting inflation — may become necessary to reassert credibility and prevent expectations from drifting away from the Fed's target.

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). Bordalo, Gennaioli, Ma and Shleifer (2020) document that overreaction is a pervasive feature in short-run expectations and can be explained by a model with diagnostic expectations. Kohlhas and Walther (2021) show that a model with rational but inattentive agents can also generate overreaction. 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 central to modern monetary policy theory and practice.[5]

Bordalo, Gennaioli, Ma and Shleifer (2020) extend the methodology introduced by Coibion and Gorodnichenko (2015) to individual forecaster data from the 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. We 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.

Farmer, Nakamura and Steinsson (2023) show that rational Bayesian learning about hard-to-learn long-run dynamics of the economy can explain anomalies in short-rate and GDP-growth forecasts, as well as deviations from the expectations hypothesis of the term structure. Like us, they place uncertainty about the long run at the center of the forecasting problem, but our objective is different. We develop a framework for evaluating whether individual forecasters use private and public forward-looking information rationally when forming long-run inflation expectations. Our analysis also differs in its data and model. They focus on shorter-horizon consensus forecasts, whereas we use a panel of individual long-run inflation forecasts, whose cross-sectional dimension is essential for uncovering overconfidence in private information. In their model, forecasts are formed through learning about unknown parameters and latent components from the history of the series being forecast, whereas our forecasters infer the latent trend from current inflation and forward-looking public and private signals.[6]

[5] An exception is Bordalo et al. (2024), who document overreaction in analysts' long-term profit expectations using diagnostic expectations.

[6] The conclusions are complementary: rational learning accounts for the anomalies in the consensus forecasts they study, whereas our panel evidence requires both overconfidence in private information and a persistent downward expectations bias, which can be interpreted as reflecting the anchoring of long-run

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 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 but the central focus is on understanding long-run forecasts of inflation.

Giacomini, Skreta and Turen (2020) use professional forecasts of 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) introduces a novel individual-level measure of inattention based on the persistence of forecasters' deviations from the consensus and documents particularly persistent disagreement in SPF long-run inflation forecasts. By basing his analysis on forecast disagreement, he avoids the practical hurdles involved in measuring forecast errors for long-run expectations. We instead estimate trend inflation using a state-of-the-art time-varying-parameter trend-cycle model and use the resulting forecast errors to evaluate departures from rationality.[7]

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 panel data. Neither of these two studies evaluates the rationality of professional forecasters' beliefs about long-run inflation.

expectations induced by central-bank credibility.

[7]Section 7 shows that our conclusions are unchanged when forecast errors are constructed using subsequently realized inflation, at the cost of shortening the sample by ten years. This suggests the estimated trend is a reliable proxy for computing forecast errors for long-run inflation expectations.

A related literature has estimated trend-cycle models similar to ours and studied the link between trend inflation and survey-based inflation 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. Chan et al. (2018) document that taking into account a link between trend inflation and average inflation expectations improves the precision of trend estimates and show that this link varies over time. 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 mostly 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.

Kumar, Afrouzi, Coibion and Gorodnichenko (2015) offer a taxonomy of definitions of expectations anchoring and evaluate them for New Zealand using firm 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.

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).

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. In the model, the trend component is the best forecast of long-run inflation for forecasters. However, the trend component is unobservable and forecasters rely on observing inflation (the *inflation signal*) and 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 estimates of trend inflation using Bayes' rule. 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 a time-varying-parameter trend-cycle model building on Stock and Watson (2007) but allowing for an autoregressive cyclical component similar to Cogley et al. (2010), Chan et al. (2018) and Mertens and Nason (2020):

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

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

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

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 cyclical components follow log random walk processes:

$$\ln(\sigma^2_{\lambda,t}) = \ln(\sigma^2_{\lambda,t-1}) + \gamma_\lambda \varepsilon_{\lambda,t}; \tag{4}$$

$$\ln(\sigma^2_{\eta,t}) = \ln(\sigma^2_{\eta,t-1}) + \gamma_\eta \varepsilon_{\eta,t}, \tag{5}$$

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:

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

where we follow Chan et al. (2018) and assume 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 based on sectoral inflation data.[8] Employing a multivariate trend-cycle model would significantly increase the already substantial computational burden of solving the forecasters' signal extraction problem. Our trend-cycle model also abstracts from modeling an explicit link between expectations and trend inflation. Chan et al. (2018) introduce a model that links trend inflation and average long-run inflation expectations. Like us, they find persistent gaps between trend and average expectations.[9] In Appendix C, Table 2, we provide evidence on the better forecast performance of our trend-cycle model relative to a range of alternative models and highlight the importance of allowing for time-varying parameters and persistence in the cyclical component.

[8] Crump 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.

[9] Chan et al. (2018) also show that such a model improves the precision of trend estimates. However, the forecasting performance of the estimated trend at long horizons is not necessarily better than the one from a model without long-run expectations and depends on the sample period and type of inflation measure used.

3.2 Forecasters' information

Forecasters base their long-run inflation forecasts on imperfect information about the three drivers of inflation ($\bar{\pi}_t$, ψ_t, and ω_t). We model imperfect information following the tradition of noisy information models, e.g., Sims (2003), Woodford (2003), Maćkowiak and Wiederholt (2009). In both the rational and behavioral models, forecasters 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.[10] 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.[11]

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

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

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

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: $\ln \sigma^2_{c,t} = \ln \sigma^2_{c,t-1} + \gamma_c \epsilon_{c,t}$, with standard normal innovations and with γ_c strictly

[10] Nimark (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.

[11] Rich 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 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.

positive.[12]

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.[13]

The parameter $\alpha(i)$ captures heterogeneity in the realizations of the coordinating signals across forecasters, as well as differences in their sensitivity to that signal. Forecasters with larger $\alpha(i)$ receive a relatively less precise coordinating signal, making their expectations less responsive to them. Since $\alpha(i)$ is strictly positive, the noise process pushes all coordinating signals across forecasters in the same direction. As the dispersion of $\alpha(i)$ across forecasters increases, the correlation between forecasters' coordinating signals decreases. In the special case where $\alpha(i)$ is identical across all forecasters, the coordinating signals, $\tilde{s}_t(i)$, are perfectly correlated across forecasters and effectively collapse into a public signal.

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:

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

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

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

[12] The 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.

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

persistent effects of individual judgment on forecasts, including, but not limited to, enduring disagreement over the inflation generating process among professional forecasters.[14]

Since agents observe inflation and know the trend-cycle model's parameters, they can determine the Kalman filtered estimate of trend inflation, that is, the estimate based only on data available up to the current time period. As we will discuss, we, as econometricians, estimate trend inflation using the trend-cycle model with a Kalman smoother, 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.

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 do not necessarily apply to longer-run expectations. As it takes many 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.

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,

[14] Andre et al. (2022) provide direct evidence that even among experts, predictions about the effects of macroeconomic shocks on inflation are widely dispersed, reflecting heterogeneous subjective views about the relevant economic mechanisms.

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.[15] Importantly, the degree of misperception can vary across forecasters and parameters.

The misperceived parameters make the behavioral model a flexible framework for distorted beliefs, nesting several microfoundations proposed in the literature, including overconfidence in private information and diagnostic expectations (Appendix H). Our purpose is therefore not to take a stand on which microfoundation is correct, but to let the likelihood identify which distortions of the signal-extraction problem the data require.

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 types 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 A 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 and by employing Markov chain Monte Carlo techniques to obtain smoothed estimates of the trend, cyclical, and i.i.d. components of inflation.

In the second step, we use Bayesian likelihood estimation to obtain posterior estimates of 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 estimates of the trend and cyclical components 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

[15] Since the average precision of the coordinating signals, as captured by $1/\sigma_{c,t}$, is perfectly observed, misperception regarding the parameter $\alpha(i)$ effectively captures misperception of the precision of the coordinating signal.

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' beliefs about the state via the Kalman filter. These equations are the solution to the signal extraction problems solved by the forecasters.[16] As the trend-cycle model implies that the expected value of inflation at long horizons equals the trend, $\bar{\pi}_t$, we equate 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 $\sigma_\nu(i)$, $i = 1, 2, \ldots, N$). When we estimate the behavioral model, the second-step estimation also returns the parameters capturing the perceived relative volatility of the coordinating signals ($\alpha^*(i)$, $i = 1, 2, \ldots, N$), the perceived volatility of idiosyncratic innovations ($\sigma^*_\nu(i)$, $i = 1, 2, \ldots, N$), and the perceived noise persistences ($\rho^*_c(i)$ and $\rho^*(i)$, $i = 1, 2, \ldots, N$).

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 choice of a two-step estimation procedure is motivated by both conceptual and econometric considerations.

Since trend inflation is a latent variable that forecasters seek to infer, it would be ideal to estimate the model conditional on the true value of trend inflation. As this value is unobservable, we instead rely on the most accurate proxy available: the Kalman-smoother estimate obtained from a state-of-the-art trend-cycle model with time-varying parameters. This estimate leverages nearly seventy years of quarterly inflation data, drawing on the full sample to improve precision. The use of time-varying parameters allows the model to account for potential structural breaks over the long estimation period. The two-step procedure is designed to first recover this smoothed trend and then condition the second-step estimation of the model parameters on the resulting trend estimate. Instead, a single-step estimation would restrict our sample to the much shorter SPF sample and would not allow us to use

[16] Appendix A.1 describes the state-space model we use for the panel estimation in detail and Appendix A.2 describes the initial conditions we use for the estimation.

the Kalman smoother that incorporates forward-looking information to get the best trend estimate.

Importantly, the use of the smoothed trend gives the signals a genuinely forward-looking role. Because the trend is inferred using subsequent inflation realizations, contemporaneous signals can be informative about low-frequency movements that have not yet appeared in current inflation. This provides a tractable way to embed forward-looking information in noisy-information tests of expectation rationality.

If we were to jointly estimate all model parameters, then observed long-run inflation expectations would influence the estimated trend inflation alongside short-run inflation realizations. While this may appear appealing in principle, such an approach risks distorting the estimated trend inflation as the model adjusts it to compensate for potential misspecification. The two-step procedure resolves this concern by holding the trend fixed across both the rational and behavioral models. This ensures a fair model comparison and allows for a clearer analysis of how cognitive biases shape expectations.

This last issue is particularly pronounced in the rational expectations model, where joint estimation proves challenging: the estimated trend inflation tends to shift considerably in an attempt to reconcile the model with the data, frequently resulting in implausible estimates of trend inflation and the model parameters. By contrast, when applied to the better-fitting behavioral model, joint estimation of the parameters confirms the evidence of overconfidence in private information and persistent expectations bias (see a description of this alternative joint estimation approach in Appendix G).

Does the two-step procedure overstate the bias? The deviations from rationality we document do not hinge on this design: overreaction and the negative forecast-error bias are essentially unchanged when forecast errors are computed from realized future inflation, dispensing with the estimated trend altogether (Section 7), and the rational model remains misspecified along the same two dimensions when forecasters estimate the trend-cycle parameters in real time (Appendix F).

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 expectation, which refers to the average expected inflation between years 5 and 10 from today and is generally not affected by cyclical and high-frequency shocks to inflation. However, these data only become available in 2011 in the SPF.[17] 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.

The CPI data is typically released 10-15 days after the end of the reference month and the panelists in the SPF submit their projections in the second month of the quarter. This means that forecasters in the SPF do not (fully) observe inflation in the quarter they are surveyed. We address this by lagging SPF forecasts by one quarter when we estimate the model. For example, we measure long-run expectations in 2017Q4 using forecasts from the survey that was conducted in February 2018.

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 forecasters. Note that in some cases, there are gaps in the time series of forecasts for individual forecasters.[18] In Appendix B, we report some statistics on the subset of forecasters used in our analysis versus the full set of forecasters in the SPF survey over time. On average our analysis includes 72% of all forecasters. We also demonstrate that average and median long-run expectations in our sample of forecasters correspond closely to their values in the complete SPF sample.[19]

[17]In 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.

[18]The 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.

[19]There is a trade-off between including as many forecasters as possible — to better identify parameters over the time-series dimension — and ensuring a sufficient number of observations per forecaster to accurately estimate forecaster-specific parameters. In our view, a threshold of 32 quarters strikes a good balance in this regard. However, our results are robust to using a lower threshold of at least 16 forecasts instead of 32. This lower threshold increases the number of forecasters from 51 to 78. These results are available upon request.

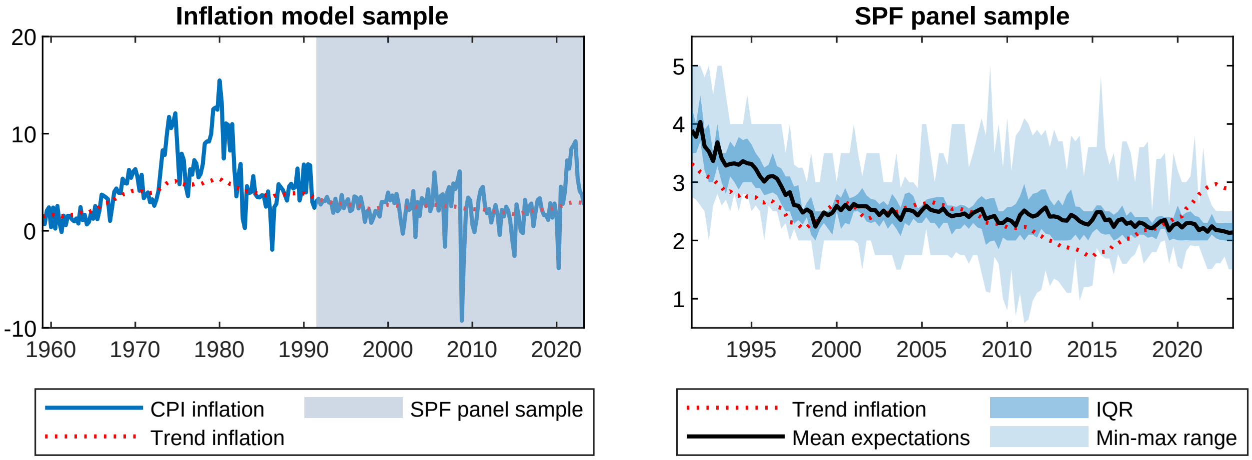

Figure 1: **Inflation, trend inflation, and long-run inflation expectations.** *Left panel*: U.S. quarterly headline CPI inflation rate (blue solid line) and its trend component (red dotted line) estimated using the trend-cycle model conditional on the entire data set (1959Q1-2023Q2). The shaded gray 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.

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.[20] The left panel 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 panel 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 panel illustrates the data the rational and behavioral models are asked to explain in the second step of the estimation. Average expectations are only implicitly used in the estimation. 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 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 considerably during the first period

[20] The estimated parameters of the trend-cycle model are shown in Appendix C.

the federal funds rate was at its ELB, 2008Q4–2015Q4.[21] More generally, the large and time-varying cross-sectional dispersion in the individual SPF long-run inflation expectations (the blue ranges in the right panel of Figure 1) may also represent a challenge for the rational and behavioral models. While the interquartile range appears visually stable over the sample, its composition is not.

While it is feasible to include short-term inflation expectations in the estimation, we have decided not to do so for a number of reasons. First, a 2009 SPF special survey on forecasting methods reveals that forecasters change their forecasting approach depending on the forecast horizon, effectively treating short-run and long-run forecasting as distinct tasks (Stark, 2013).[22] Second, jointly modeling forecasts for these two activities is not straightforward. Forecast errors of long-run inflation expectations tend to be highly persistent and expectations significantly overreact to news, whereas short-run inflation expectations, as we will show in Section 7, do not exhibit the same biases.[23] Therefore, incorporating both short- and long-run expectations into an imperfect information model presents a significant challenge. At a minimum, it would require a substantial expansion of the information or signal structure. In light of these considerations, we have chosen not to include short-run SPF expectations in the estimation.

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

[21]The 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 figure 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.

[22]de Vincent-Humphreys et al. (2019) provide similar evidence for the European Central Bank's Survey of Professional Forecasters. In particular, they document that forecasters use different types of models for short- and medium-term forecasts than for longer-term forecasts, and that the relative weight given to models versus judgment varies with the forecast horizon.

[23]Halperin and Mazlish (2025) also highlight the important role of the forecast horizon for forecast biases. Using consensus forecasts across 89 countries they document that long-run expectations overreact while short-run expectations underreact.

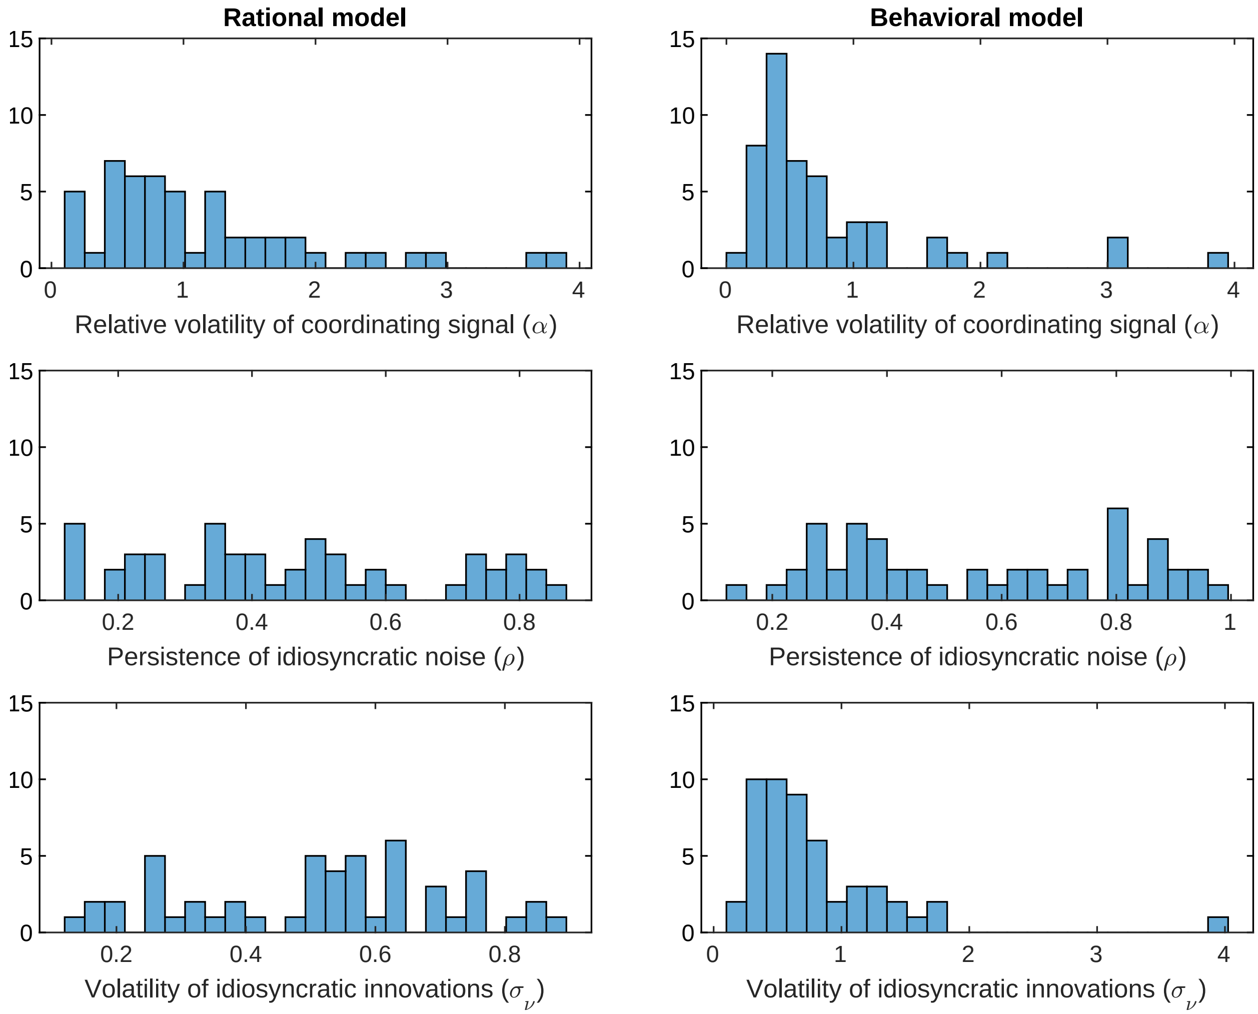

Figure 2: **Estimated model parameters.** Posterior mode for the rational (left panels) and behavioral (right panels) models' forecaster-specific parameters: the relative volatility of each forecaster's coordinating signal $\alpha(i)$ (first row), the persistence of each forecaster's idiosyncratic noise,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

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. The estimates of α allow us to assess how similar the coordinating signals are across forecasters. Given our estimates, the average correlation of coordinating signals across forecasters is above 90 percent making them close to a public signal. 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

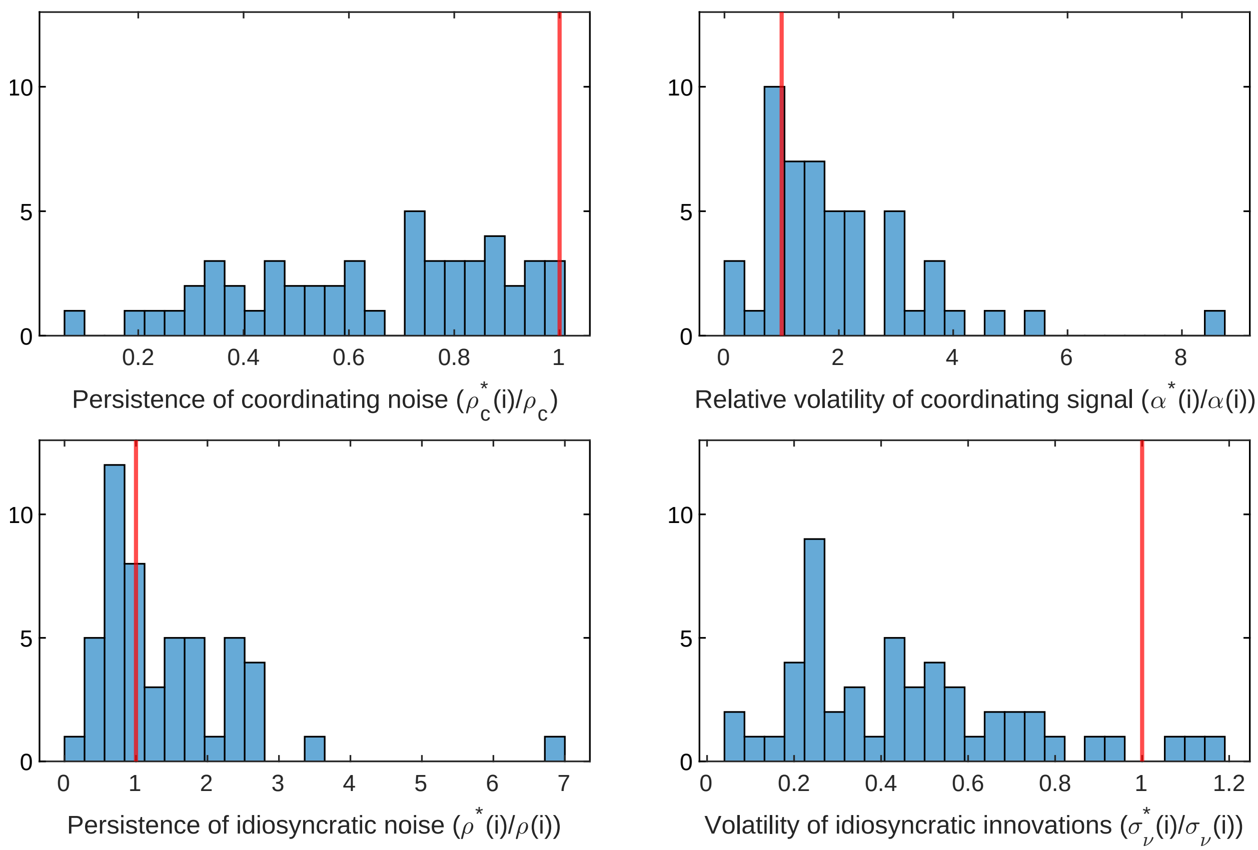

Figure 3: **Deviations from rationality.** Ratio of the posterior mode of the behavioral parameters ($\rho_c^*(i)$, $\alpha^*(i)$, $\rho^*(i)$, $\sigma_\nu^*(i)$) to the posterior mode of the true value of the corresponding structural parameters (ρ_c, $\alpha(i)$, $\rho(i)$, and $\sigma_\nu(i)$). The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

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 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 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 panel of Figure 4, which shows the response of long-run inflation expectations to a one-standard-deviation innovation

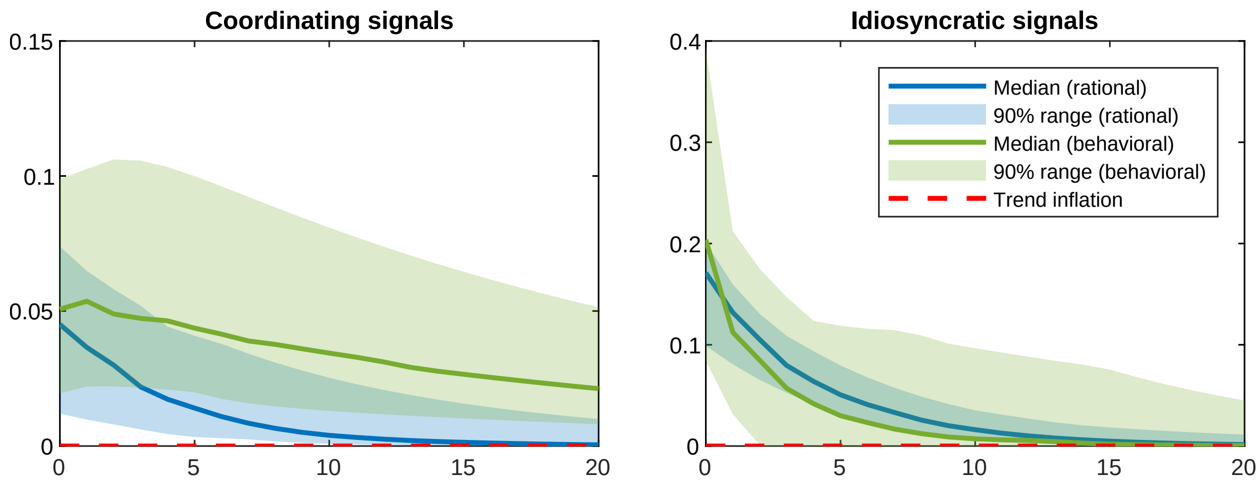

Figure 4: **Propagation of noise innovations to expectations.** Impulse response functions of long-run inflation expectations of every forecaster to a one-standard-deviation noise innovation to the coordinating signal (left panel) and to every idiosyncratic signal (right panel). Responses in the rational 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.

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.[24]

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 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 panel 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 as indicated by the

[24]In both the left and right panels 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, $\sigma_{c,t}$.

blue line in the left panel of Figure 4.

In the left panel 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 panel 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 panel of Figure 4, which shows forecasters' impulse responses to a one-standard-deviation idiosyncratic innovation, $\nu_t(i)$. In this panel, 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 misspecification that plagues the rational model. The wider green shaded areas in the two panels 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 panels in Figure 3 show that, in the behavioral model, most forecasters tend to overestimate the relative 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 *filtered* coordinating innovations ($\nu_{c,t}$) and idiosyncratic innovations ($\nu_t(i)$). The upper panels 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

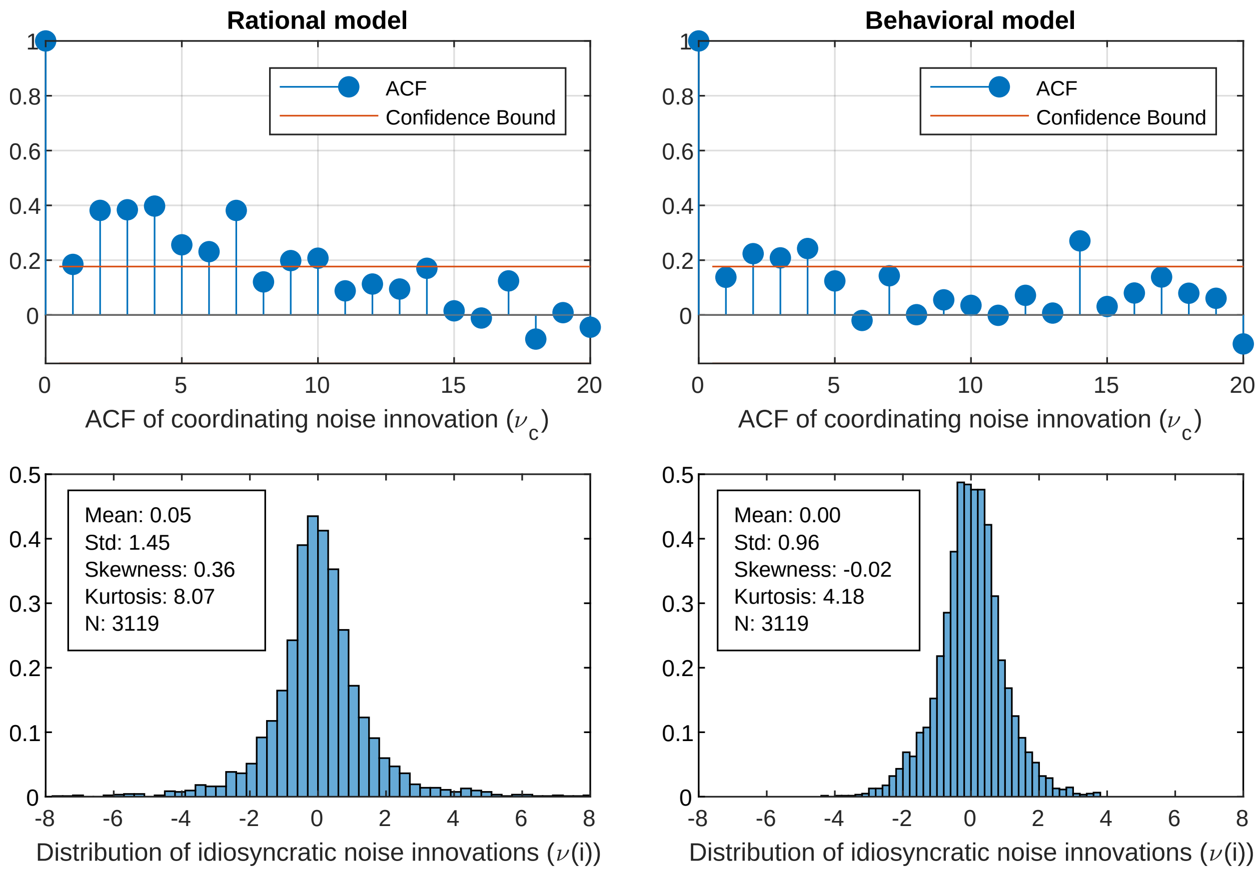

Figure 5: **Assessing misspecification in the two models.** The upper panels show the autocorrelogram of the estimated in-sample noise innovations to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The lower panels show the distribution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

noise exhibit significant in-sample serial correlation.[25] 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 panel of Figure 1. As shown in the left panel of Figure 4, the impulse response to changes in the coordinating signal in the rational model is insufficiently persistent to account for this phenomenon.[26] 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 panel of Figure 5 shows the autocorrelation function of the coordinating innovations in the behavioral model. These innovations are much closer to white noise

[25]Note that we focus on the filtered estimates since the smoothed estimates may add serial correlation that is unrelated to model misspecification.

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

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 panel of Figure 4. Consequently, the behavioral model does not need to violate the white noise 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 panels of Figure 5 display the distribution of the estimated in-sample realizations of the idiosyncratic innovations for all forecasters. In the rational model (left panel), the estimated idiosyncratic innovations exhibit excessive volatility, skewness, and kurtosis. In contrast, the idiosyncratic innovations estimated in the behavioral model (right panel) 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 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 tight link 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 panel 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. The detailed results for this real-time estimation — both in a time-varying parameter and a constant parameter setting — are presented in Appendix F.

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.

Finally, we assess the implications of alternative specifications of the trend-cycle model for the model misspecification that undermines the rational model. First, estimating the trend-cycle model with core CPI inflation instead of headline CPI inflation leads to the same type of misspecification. Second, taking our baseline model but assuming constant parameters and starting the estimation at the beginning of the Great Moderation in 1984 also leads to similar conclusions. Detailed results for all of these robustness checks can be found in Appendix F.

7 Forecasters' overreaction and persistent forecast errors

In this section, we show that the behavioral model matches untargeted properties of forecast errors using the regression framework introduced by Coibion and Gorodnichenko (2015) and applied to individual forecasters by Bordalo et al. (2020). Originally designed to test the FIRE hypothesis, we repurpose this framework as a benchmark to assess how well the rational and behavioral models capture key features of forecast errors in the SPF data. The regression results confirm two central empirical patterns in the formation of long-run inflation expectations: persistent expectations bias and overreaction to news. Here, these patterns emerge within an alternative less parametrically restricted framework, reinforcing their robustness and providing an independent validation of our findings. The framework also highlights systematic differences between short- and long-run expectations, motivating the development of models specifically tailored to the long run.

We estimate the following pooled regression:

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

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 in the SPF. 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).[27] 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 shows that our two-step design does not mechanically generate or amplify the departures from rationality we document: the realized-inflation forecast errors make no use of the estimated trend, yet they deliver essentially the same overreaction and a negative forecast-error bias that is, if anything, larger, despite losing the final ten years of the sample.

In the third row of the table, short-term expectations are defined as in Bordalo et al.

[27] This 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.

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

	β_0^p	SE	β_1^p	SE	R^2
I. Data					
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
II. Models					
Rational model	-0.02	0.01	-0.39***	0.03	0.07
Behavioral model	-0.11***	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 * indicate statistical significance at the 1%, 5% and 10% level, respectively.

(2020).[28] Notably, overreaction is statistically significant only for long-run inflation expectations.[29] 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

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

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

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.

[29] While 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.

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.[30]

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 and fully informed cannot explain overreaction. However, our rational model with noisy information goes a long way toward its estimated value.

The behavioral model successfully accounts for the magnitude of overreaction when innovations are simulated from the true distribution, yielding an estimated $\beta_1^p = -0.47$. 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 panel 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, 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

[30] The estimated β_1^p coefficients are essentially unchanged when all 1,000 simulations are used. The average β_0^p across all 1,000 simulations is close to zero for every model, as the models' innovations are drawn from symmetric distributions; conditioning on simulations with a negative average forecast error is therefore necessary to evaluate the models against a data sample that, like ours, features a persistent negative bias.

of these in-sample realizations, illustrated in Figure 5, is remarkably close to a standard normal distribution.

Negative forecast error bias and persistent expectations bias. Consistent with Figure 1, Table 1 reveals that the SPF forecast errors exhibit a negative bias. 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.

The baseline behavioral model, incorporating all cognitive biases, can successfully account for such a 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 across all simulations 90% of the estimated β_0^p lie between -0.27 and -0.01 which includes 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.The rational model, the fourth row of the table, also fails to account for the negative forecast error bias observed in the data. Across all simulations 90% of the estimated β_0^p lie between -0.05 and -0.00.

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 persistent 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, during which inflation surged sharply, leading to persistent positive forecast errors. 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.

Diagnostic expectations and the central bank's credibility. An alternative explanation for overreaction, extensively documented in the context of short-term expectations, is offered by diagnostic expectations (Bordalo, Gennaioli and Shleifer, 2018, Bordalo, Gennaioli, Ma and Shleifer, 2020, Bordalo, Gennaioli and Shleifer, 2022). In Appendix H, we show that diagnostic expectations and overconfidence are closely related mechanisms; in multi-signal environments, the discrepancy between them primarily depends on how asymmetric overconfidence is across signals. However, like overconfidence (last row of Table 1), diagnostic expectations alone cannot account for the persistent bias in forecast errors that is central to the SPF panel of long-run inflation expectations. As we show in Appendix H.3, when noise persistence is asymmetric across signals, diagnostic expectations do not generate persistent forecast errors. This is the empirically relevant case in our setting, where the inflation signal carries transitory noise while the coordinating signal carries more persistent noise. Matching that pattern in our estimation requires an additional deviation from rationality: the persistent expectations bias.

A natural economic interpretation of this bias is that the coordinating signal partly conveys the central bank's commitment to price stability, so that forecasters' slow recognition of its persistent component is what keeps long-run expectations from tracking movements in trend inflation for long periods. Over significant parts of our sample, this has kept average long-run expectations close to the Fed's target even as observed inflation has diverged from it for a prolonged period. This represents a novel finding of our paper on long-run inflation expectations, complementing the existing literature on survey-based short-term macroeconomic expectations.

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 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 panel of Figure 6 illustrates the sensitivity of forecasters' long-run inflation expectations to the inflation signal. Throughout the sample period, sensitivity to the inflation 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 2000, were largely anchored in the sense of Bernanke (2007).[31] This finding

[31]This finding is consistent with empirical evidence for the U.S. on the sensitivity of long-run inflation

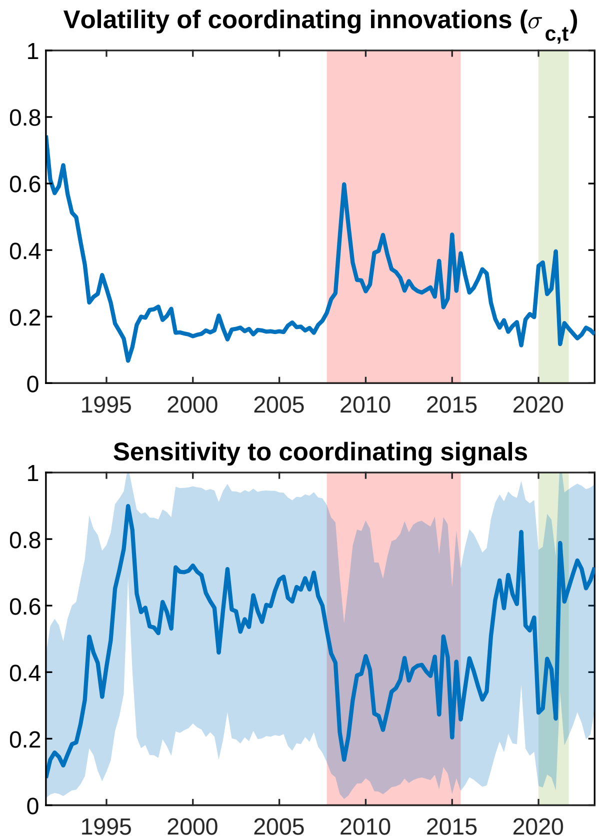

Figure 6: **Time-varying volatility of coordinating innovations and sensitivities of expectations to signals.** The upper left panel shows the posterior mode of the time-varying volatility of coordinating innovations ($\sigma_{c,t}$). The upper right panel shows the median and 90 percent range of the distribution of sensitivities of individual expectations to inflation. The lower left panel shows the median and 90 percent range of the distribution of sensitivities of individual expectations to the coordinating signals. The lower right panel 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.

aligns with Hazell, Herreño, Nakamura and Steinsson (2022), which shows that the stability of inflation since 1990 is largely driven by the anchoring of long-run expectations rather than by a flattening of the Phillips curve. Moreover, we observe minimal variation in sensitivity across forecasters, as indicated by the narrow blue shaded area.

The lower-left panel of Figure 6 displays the sensitivity of long-run inflation expectations to the coordinating signal. 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-

expectations to short-term expectations (e.g., Barlevy, Fisher and Tysinger (2021), Carvalho, Eusepi, Moench and Preston (2023) and references therein).

varying parameters of the trend-cycle model. A comparison of the two panels 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 the coordinating signal is much larger and exhibits much greater variation across forecasters than the sensitivity to the inflation signal. The sensitivity to the coordinating signal is also larger than for the idiosyncratic signal for most of the sample (lower right panel). 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 (median).

The model accounts for forecaster heterogeneity in part by allowing sensitivities to coordinating signals to vary across individuals, resulting in a wide range of sensitivities (represented by the blue shaded area in the lower-left panel 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 Fed lifts off the federal funds rate from near zero to end both ELB periods, we observe an increase in the median sensitivity to the coordinating signal.

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.[32]

[32] Other 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).

9 The anchoring-compatible inflation path

We use our model to derive the inflation path consistent with average long-run expectations equal to the central bank's target going forward, which is the notion of anchoring we use in this section. In the context of the post-pandemic period, this analysis helps answer two important questions: How quickly should the central bank disinflate? And how far must inflation fall to prevent expectations from drifting? The faster and deeper the disinflation, the greater the risk that the required tightening triggers a recession. We compute this *anchoring-compatible path* starting from 2022Q4 and then compare it to the projections of the participants at the December 2022 FOMC meeting as well as actual inflation outcomes.

To compute the anchoring-compatible path of inflation, we begin by guessing a trajectory for trend inflation over the period 2023-2027. We use the model to find the path of inflation that keeps long-run expectations at the target level over the subsequent five years, given the guessed evolution of the trend.[33] Then we assess whether the model-implied path of inflation is consistent with the guessed dynamics of trend inflation. To this end, we use the estimated trend-cycle model to obtain the trend of the model-implied path of inflation. If this trend aligns with the initial guess, the path of inflation is deemed anchoring-compatible. Otherwise, the procedure is iterated, using the newly estimated trend as the updated input. A detailed description of the algorithm is provided in Appendix I. In sum, finding the anchoring-compatible path involves forecasting inflation from 2022Q4 under the assumption that average long-run inflation expectations stay fixed at their level at that time for the next five years, and that trend inflation is consistent with the forecasted path of inflation.[34]

The level at which expectations are considered to be anchored is 2 percent in terms of PCE inflation, the FOMC's stated inflation objective. However, our analysis uses CPI inflation, and the gap between CPI and PCE inflation is time-varying. This makes it nontrivial to translate the FOMC target into CPI terms. As a working assumption, we set the long-run CPI inflation target at 2.17 percent – the level of 5-year, 5-year forward CPI inflation expectations from the SPF observed in 2022Q4. We think this is a reasonable assumption given that average 5-year, 5-year forward PCE inflation expectations in 2022Q4 are 2.0 and, hence, at the FOMC's target.

Figure 7 displays various aspects of the anchoring-compatible path and information from the FOMC's December 2022 Summary of Economic Projections (SEP).[35] We first consider the

[33]Due to the notion of anchoring used in this section, which focuses on average long-run inflation expectations, we turn off the innovations to the idiosyncratic signals. For simplicity, we also set the i.i.d. component of inflation to zero.

[34]Given our estimation sample ends in 2023Q2, this is not a fully out-of sample exercise. For that we would have to end our estimation in 2022Q3.

[35]The SEP is released quarterly by the FOMC. The SEP presents individual Committee participants'

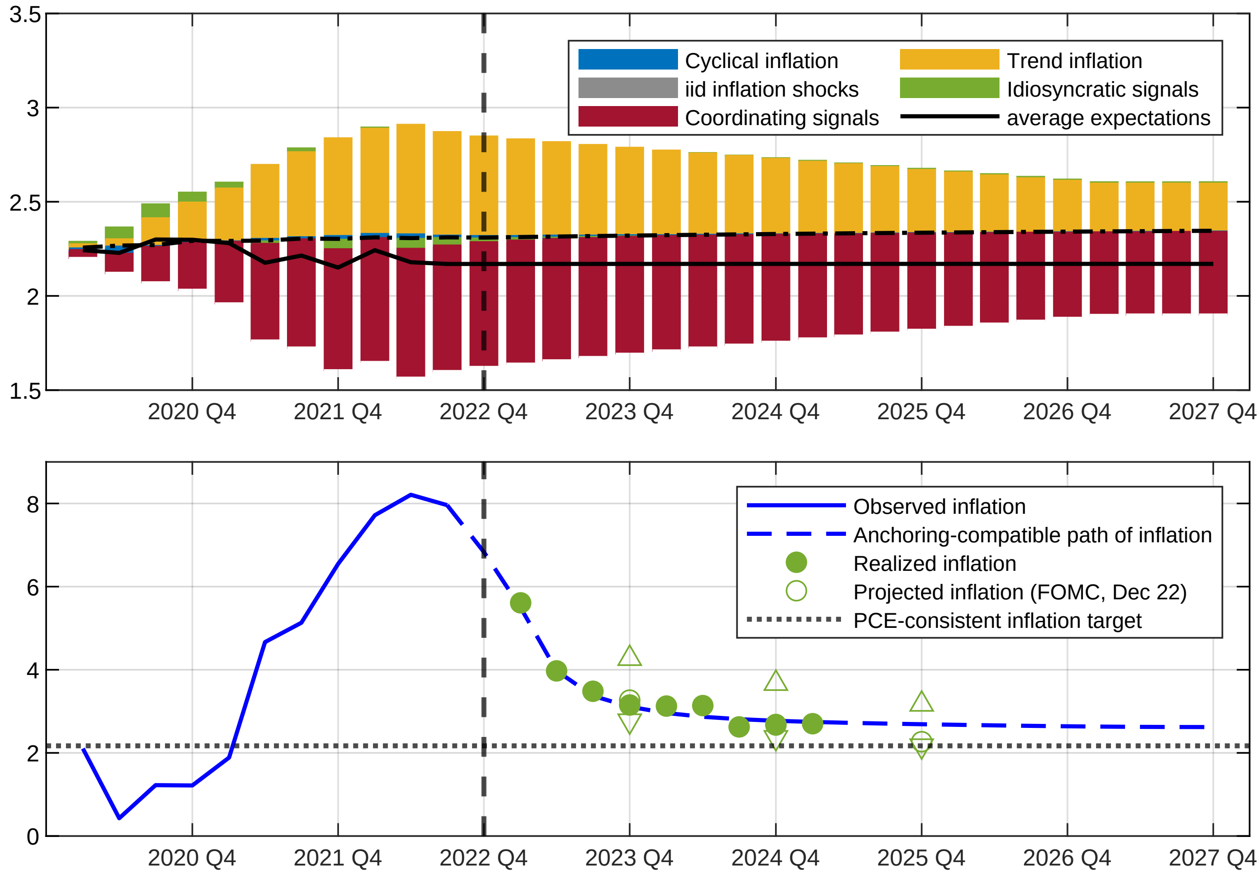

Figure 7: **Anchoring-compatible inflation path.** The upper panel shows the shock decomposition of the gap between average SPF inflation expectations (black solid line) and the model's forecast of average expectations in the absence of any shocks as of 2019Q4 (black dashed-dotted line). The bars show the contribution of each model shock through 2022Q3, when SPF expectations are observed, and from 2022Q4 onward, when average expectations are assumed to remain anchored at their level from the 2022Q4 survey (2.17%). The lower panel shows observed year-on-year CPI inflation (solid blue line) and the model-implied path consistent with 5-year, 5-year forward SPF inflation expectations remaining anchored at their end-2022 level (dashed blue line). Green filled circles denote actual CPI inflation. Open green circles indicate the SEP median projections for PCE inflation (2023-2025), converted to CPI units using a 17 basis point adjustment. Upward and downward triangles show the SEP projection range. In both panels the vertical black dashed line marks the start of the anchoring-compatible projection period.

upper panel. The black line left of the vertical line displays actual SPF average expectations and to the right the level of average expectations being targeted. The bars capture how each of the model's shocks contributed to changes in average expectations that the model could not forecast as of 2019Q4 through the end of the simulation period. The decomposition is conditioned on information available in 2019Q4 which is before the pandemic was reflected in

forecasts for key economic indicators, including GDP growth, the unemployment rate, total and core PCE inflation, and the federal funds rate, over the next three calendar years and in the longer run, based on their assessments of appropriate monetary policy. The SEP is available at https://www.federalreserve.gov/monetarypolicy/fomc.htm.

professional forecasters' expectations.

The model explains the relative insensitivity of SPF inflation expectations to the outburst of inflation over the 2021-2022 period through the persistent expectations bias.[36] The increase in trend inflation in 2021-2022 contributed to pushing average expectations upward (yellow bars). Likely reflecting the Federal Reserve's strong anti-inflation reputation, the coordinating signal played a stabilizing role (red bars) by offsetting the upward pressure from rising trend inflation on average expectations, effectively preventing their de-anchoring.

Against this backdrop, what inflation path announced in 2022Q4 would keep average long-run expectations stable over the next five years? This path is depicted by the dashed blue line in the bottom panel of Figure 7. The path features a rather steep decline in inflation during 2023, followed by a more gradual convergence toward the level at which long-run inflation expectations are assumed to remain anchored.

The elevated trend inflation estimated at the end of 2022 continues to exert upward pressure on long-run inflation expectations throughout the five-year simulation horizon and absent a countervailing factor would be projected to de-anchor expectations. Despite this positive contribution to the potential de-anchoring of expectations, the model does not recommend communicating a policy stance aimed at undershooting the inflation target. This prediction is mainly driven by the fact that the rise in trend inflation has not become embedded in forecasters' expectations due to the negative contributions of the persistent expectations bias present before 2022Q4.

Avoiding an inflation undershoot grants the central bank greater flexibility and allows for a less hawkish policy stance, thereby increasing the likelihood of a soft landing, that is, bringing inflation back to target without triggering a recession. Of course, this is not a general result. It reflects a specific condition estimated by our model at the end of 2022. Namely, a large and negative persistent expectations bias. Our model incorporates this bias when determining the inflation path consistent with anchored expectations.

However, if, going forward, SPF inflation expectations were to become more sensitive to the heightened trend inflation and began to drift upward, the scenario implied by our model could rapidly deteriorate. If the elevated trend inflation became embedded in professional forecasters' expectations, the impact of the persistent expectations bias would dwindle and, as a result, maintaining anchoring would require the central bank to communicate a faster and deeper disinflation path. In this case, tighter monetary policy aimed at engineering an inflation undershoot may be required to steer trend inflation closer to the central bank's target.

[36]See Reis (2026) for a comprehensive analysis of the role of expectations in linking shocks to inflation outcomes during 2021-24.

As of 2022Q4, there is little evidence of de-anchoring. However, this could change if inflation remains persistently above target or even surges again.[37] In that case, expectations of some forecasters might start rising toward the higher trend inflation, with the distribution of expectations becoming more dispersed and skewed to the right. The model would interpret this as a weaker role for coordinating signals and a stronger positive influence of idiosyncratic signals on average inflation expectations. To counter this shift, it would recommend communicating a faster path of disinflation, potentially involving a temporary undershooting of the target.

Is the anchoring-compatible path shown in the lower panel of Figure 7 consistent with the actual evolution of inflation through 2025Q1? During this period, the observed year-on-year headline CPI inflation (solid green dots) aligns remarkably closely with the anchoring-compatible path of inflation. Given that average SPF inflation expectations have remained near their level at the end of 2022 over this period (see Appendix Figure 16), we view this alignment as an indirect empirical out-of-sample validation of the model.

Since the anchoring-compatible path reflects the inflation trajectory the central bank would need to communicate to keep expectations stable in our model, we compare it with the median and projection range from the December 2022 SEP. The median projection and the projection range are represented by the green open dots and triangles, respectively, in the lower panel of Figure 7. The median SEP inflation projection aligns remarkably closely with the anchoring-compatible path generated by the model. This alignment suggests that the model sees the FOMC's median inflation projection in December 2022 as broadly consistent with maintaining average long-run expectations anchored at target.

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 paths communicated by policymakers align with their objective of maintaining well-anchored

[37] Gennaioli et al. (2024) show that for households, inflation surges can trigger selective recall of past high-inflation experiences, amplifying the risk of de-anchoring.

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.).

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.

Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart (2022) "Subjective Models of the Macroeconomy: Evidence From Experts and Representative Samples," *The Review of Economic Studies*, 89 (6), 2958–2991.

Angeletos, George-Marios, Zhen Huo, and Karthik A. Sastry (2021) "Imperfect Macroeconomic 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.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, Matthew OBrien, and Andrei Shleifer (2024) "Long-Term Expectations and Aggregate Fluctuations," *NBER Macroeconomics Annual*, 38, 311–347.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018) "Diagnostic Expectations and Credit Cycles," *Journal of Finance*, 73 (1), 199–227.

——— (2022) "Overreaction and Diagnostic Expectations in Macroeconomics," *Journal of Economic Perspectives*, 36 (3), 223–44.

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 (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.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam (1998) "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance*, 53 (6), 1839–1885.

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.

Fuhrer, Jeffrey C. (2018) "Intrinsic expectations persistence: evidence from professional and household survey expectations," Working Papers 18-9, Federal Reserve Bank of Boston.

Gennaioli, Nicola, Marta Leva, Raphael Schoenle, and Andrei Shleifer (2024) "How Inflation Expectations De-Anchor: The Role of Selective Memory Cues," Working Paper 32633, National Bureau of Economic Research.

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.

Halperin, Basil and J. Zachary Mazlish (2025) "Overreaction and Forecast Horizon: Longer-term Expectations Overreact More, Shorter-term Expectations Drive Fluctuations," mimeo.

Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson (2022) "The Slope of the

Phillips Curve: Evidence from U.S. States," *The Quarterly Journal of Economics*, 137 (3), 1299–1344.

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.

Kumar, 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 Eco-*

nomic 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.

——— (2026) "Why Did Inflation Rise and Fall in 2021–24? Channels and Evidence from Expectations," CFM Discussion Paper CFMDP2026-03, London School of Economics and Political Science.

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.

Scheinkman, José A. and Wei Xiong (2003) "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 111 (6), 1183–1220.

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.

Online Appendix

Long-Run Inflation Expectations

Jonas D. M. Fisher Leonardo Melosi Sebastian Rast

A Detailed model description

A.1 Derivation of model equations

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

$$\xi_t(i) = \mathbf{\Phi_t}(\mathbf{i})\xi_{t-1}(i) + \mathbf{R_t}(\mathbf{i})e_t(i) \tag{1}$$

$$y_t\,(i) = \mathbf{D}(\mathbf{i})\xi_t(i) + \mathbf{Q}u_t \tag{2}$$

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 $\mathbf{\Phi_t}(\mathbf{i})$ 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:

$$\mathbf{\Phi_t}(\mathbf{i}) = \begin{bmatrix} \phi_t & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \rho_c & 0 \\ 0 & 0 & 0 & \rho(i) \end{bmatrix}, \ \mathbf{R}_t(i) = \begin{bmatrix} \sigma_{\eta,t} & 0 & 0 & 0 \\ 0 & \sigma_{\lambda,t} & 0 & 0 \\ 0 & 0 & \sigma_{c,t} & 0 \\ 0 & 0 & 0 & \sigma_v(i) \end{bmatrix} \tag{3}$$

$$\mathbf{D}(\mathbf{i}) = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & \alpha(i) & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}, \ \mathbf{Q} = \begin{bmatrix} \sigma_\omega & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{4}$$

The Kalman filter recursion for forecaster i is given by:

$$\xi_{t|t-1}\,(i) = \mathbf{\Phi_t}(\mathbf{i})\xi_{t-1|t-1}\,(i) \tag{5}$$

$$P_{t|t-1}\,(i) = \mathbf{\Phi_t}(\mathbf{i})P_{t-1|t-1}\,(i)\,\mathbf{\Phi_t}(\mathbf{i})' + \mathbf{R}_t(i)\mathbf{R}_t(i)' \tag{6}$$

$$s_{t|t-1}\,(i) = \mathbf{D}(\mathbf{i})\xi_{t|t-1}\,(i) \tag{7}$$

$$F_{t|t-1}\,(i) = \mathbf{D}(\mathbf{i})P_{t|t-1}\,(i)\,\mathbf{D}(\mathbf{i})' + \mathbf{Q}\mathbf{Q}' \tag{8}$$

$$\xi_{t|t}\,(i) = \xi_{t|t-1}\,(i) + \underbrace{P_{t|t-1}(i)\mathbf{D}(\mathbf{i})'\left[F_{t|t-1}\,(i)\right]^{-1}}_{K_t(i)}\left[y_t\,(i) - \mathbf{D}(\mathbf{i})\xi_{t|t-1}\,(i)\right] \tag{9}$$

$$P_{t|t}(i) = P_{t|t-1}(i) - P_{t|t-1}(i)\mathbf{D(i)}'\left[F_{t|t-1}(i)\right]^{-1}\mathbf{D(i)}P_{t|t-1}(i) \tag{10}$$

$K_t(i)$ 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:

$$\begin{aligned}
\xi_{t|t}(i) &= \xi_{t|t-1}(i) + K_t(i)\left[y_t(i) - \mathbf{D(i)}\xi_{t|t-1}(i)\right] && (11)\\
&= \left[\mathbf{I}_4 - K_t(i)\mathbf{D(i)}\right]\mathbf{\Phi_t(i)}\xi_{t-1|t-1}(i) + K_t(i)y_t(i) && (12)\\
&= \left[\mathbf{I}_4 - K_t(i)\mathbf{D(i)}\right]\mathbf{\Phi_t(i)}\xi_{t-1|t-1}(i) + K_t(i)\left[\mathbf{D(i)}\xi_t(i) + \mathbf{Q}u_t\right] && (13)\\
&= \left[\mathbf{I}_4 - K_t(i)\mathbf{D(i)}\right]\mathbf{\Phi_t(i)}\xi_{t-1|t-1}(i) + K_t(i)\left[\mathbf{D(i)}(\Phi_t(i)\xi_{t-1}(i) + \mathbf{R_t(i)}e_t(i)) + \mathbf{Q}u_t\right] && (14)
\end{aligned}$$

Putting the trend-cycle model and the combined 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

$$\begin{bmatrix} \xi_t \\ \overrightarrow{\xi}_{t|t} \\ \omega_t \end{bmatrix} = \widetilde{\mathbf{\Phi}}_t \begin{bmatrix} \xi_{t-1} \\ \overrightarrow{\xi}_{t-1|t-1} \\ 0 \end{bmatrix} + \widetilde{\mathbf{R}}_t \begin{bmatrix} \eta_t \\ \lambda_t \\ \nu_{c,t} \\ \overrightarrow{\nu_{v,t}} \\ \omega_t \end{bmatrix} \tag{15}$$

where $\overrightarrow{\xi}_{t|t}$ and $\overrightarrow{\nu_t}$ 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.

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

.

The measurement equations for our panel estimation are

$$\begin{bmatrix} \pi_t^{cpi} \\ \psi_t^{est} \\ \bar{\pi}_t^{est} \\ \mathbb{E}_t\pi_t^{long}(1) \\ \mathbb{E}_t\pi_t^{long}(2) \\ \vdots \\ \mathbb{E}_t\pi_t^{long}(N) \end{bmatrix} = \begin{bmatrix} \mathbf{D_{CPI}} & \mathbf{0}_{1\times k} & \mathbf{0}_{1\times k} & \dots & \mathbf{0}_{1\times k} & \sigma_\omega \\ \mathbf{1}_1 & \mathbf{0}_{1\times k} & \mathbf{0}_{1\times k} & \dots & \mathbf{0}_{1\times k} & 0 \\ \mathbf{1}_2 & \mathbf{0}_{1\times k} & \mathbf{0}_{1\times k} & \dots & \mathbf{0}_{1\times k} & 0 \\ \mathbf{0}_{1\times k} & \mathbf{1}_2 & \mathbf{0}_{1\times k} & \dots & \mathbf{0}_{1\times k} & 0 \\ \mathbf{0}_{1\times k} & \mathbf{0}_{1\times k} & \mathbf{1}_2 & \dots & \mathbf{0}_{1\times k} & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0}_{1\times k} & \mathbf{0}_{1\times k} & \mathbf{0}_{1\times k} & \dots & \mathbf{1}_2 & 0 \end{bmatrix} \begin{bmatrix} \xi_t \\ \xi_{t|t}(1) \\ \xi_{t|t}(2) \\ \vdots \\ \xi_{t|t}(N) \\ \omega_t \end{bmatrix}, \tag{16}$$

where $\mathbf{D_{CPI}}$ 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 (16) 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 $\widetilde{\mathbf{\Phi}}_t$ in the transition equation defined in equation (15). $\widetilde{\mathbf{R}}_t$ and the measurement equation remains unchanged. More specifically, in equations (5)-(9) the forecaster's beliefs about the parameters are used instead of the "true" estimates so that

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

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 κ=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 (6) is at the "steady-state"[38] level given the initial parameter values. To compute this steady-state matrix we start from the following matrix

$$P_{0|0}(i) = \begin{bmatrix} \frac{\sigma^2_{\eta,1991Q2}}{1-\phi^2_{1991Q2}} & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{1-\rho_c^2} & 0 \\ 0 & 0 & 0 & \frac{\sigma^2_\nu(i)}{1-\rho(i)^2} \end{bmatrix} \tag{17}$$

and we iterate over equations (6), (8), (10) to get the "steady-state". The values from 1991Q2 are used since our first period of the panel model is 1991Q3.

Initial conditions of panel model

The initial conditions for the first part of the transition equation defined in equation (15) are as follows:

$$\xi_0 \equiv \begin{bmatrix} \epsilon_{1991Q2} \\ \bar{\pi}_{1991Q2} \\ 0 \\ \mathbf{0}_{N\times 1} \end{bmatrix} \tag{18}$$

$$\mathbf{P}_0 \equiv \begin{bmatrix} 0 & 0 & 0 & \mathbf{0}_{1\times N} \\ 0 & 0 & 0 & \mathbf{0}_{1\times N} \\ 0 & 0 & \frac{1}{1-\rho_c^2} & \mathbf{0}_{1\times N} \\ \mathbf{0}_{N\times 1} & \mathbf{0}_{N\times 1} & \mathbf{0}_{N\times 1} & \mathbf{I}_{N\times N}\overrightarrow{\frac{\sigma_v^2}{1-\rho^2}} \end{bmatrix} \tag{19}$$

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 $\xi^e_{0|0}(i)$ the initial condition of the econometrician for the state estimate of forecaster i, we assume

$$\xi^e_{0|0}(i) \equiv \begin{bmatrix} \epsilon_{1991Q2} \\ \bar{\pi}_{1991Q2} \\ 0 \\ 0 \end{bmatrix}, \quad \forall i \tag{20}$$

[38]Intuitively, 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.

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

$$\begin{aligned}\Sigma_t(i) &= var(K_t(i)y_t(i)) && (21)\\ &= K_t(i)var(\mathbf{D}(\mathbf{i})\xi_t(i) + \mathbf{Q}u_t)K_t(i)' && (22)\\ &= K_t(i)[\mathbf{D}(\mathbf{i})var(\xi_t(i))\mathbf{D}(\mathbf{i})' + \mathbf{Q}\mathbf{Q}']K_t(i)' && (23)\\ &= K_t(i)[\mathbf{D}(\mathbf{i})P_{t|t}(i)\mathbf{D}(\mathbf{i})' + \mathbf{Q}\mathbf{Q}']K_t(i)' && (24)\end{aligned}$$

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

$$\mathbf{P}^e_{0|0}(i) \equiv \Sigma_0(i), \quad \forall i \tag{25}$$

The remaining elements of the initial covariance matrix for equation (15) are assumed to be zero.

B Selection of forecasters

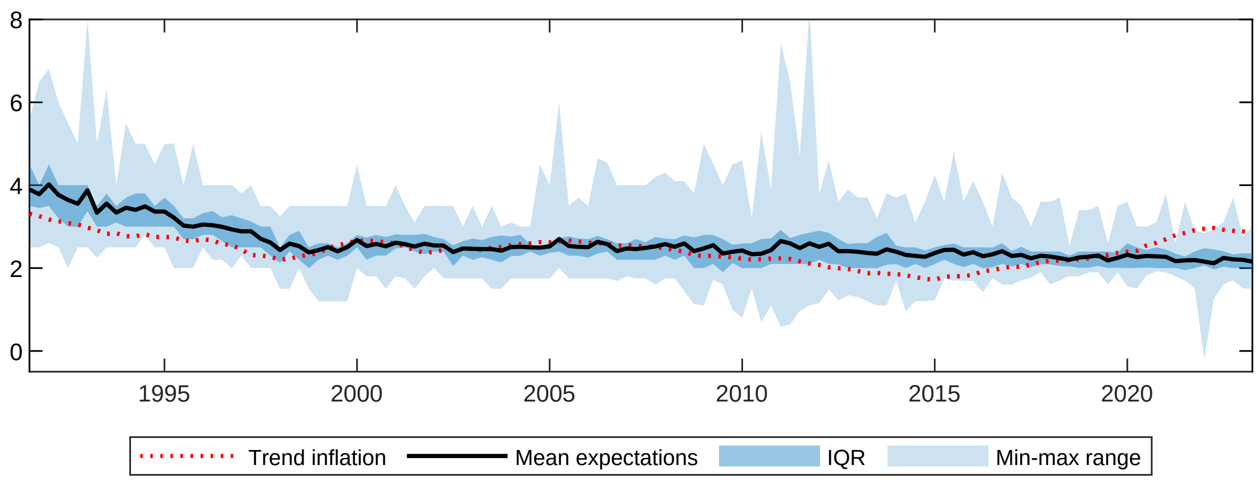

Figure 1: **Trend inflation and full sample of long-run inflation expectations.** Trend component of U.S. quarterly headline CPI inflation rate estimated using the trend-cycle model (red dotted line) and the mean (black solid line), the interquartile range (the dark blue bands), and the min-max range (the light blue bands) of the distribution of SPF long-run CPI inflation expectations including all forecasters in the survey.

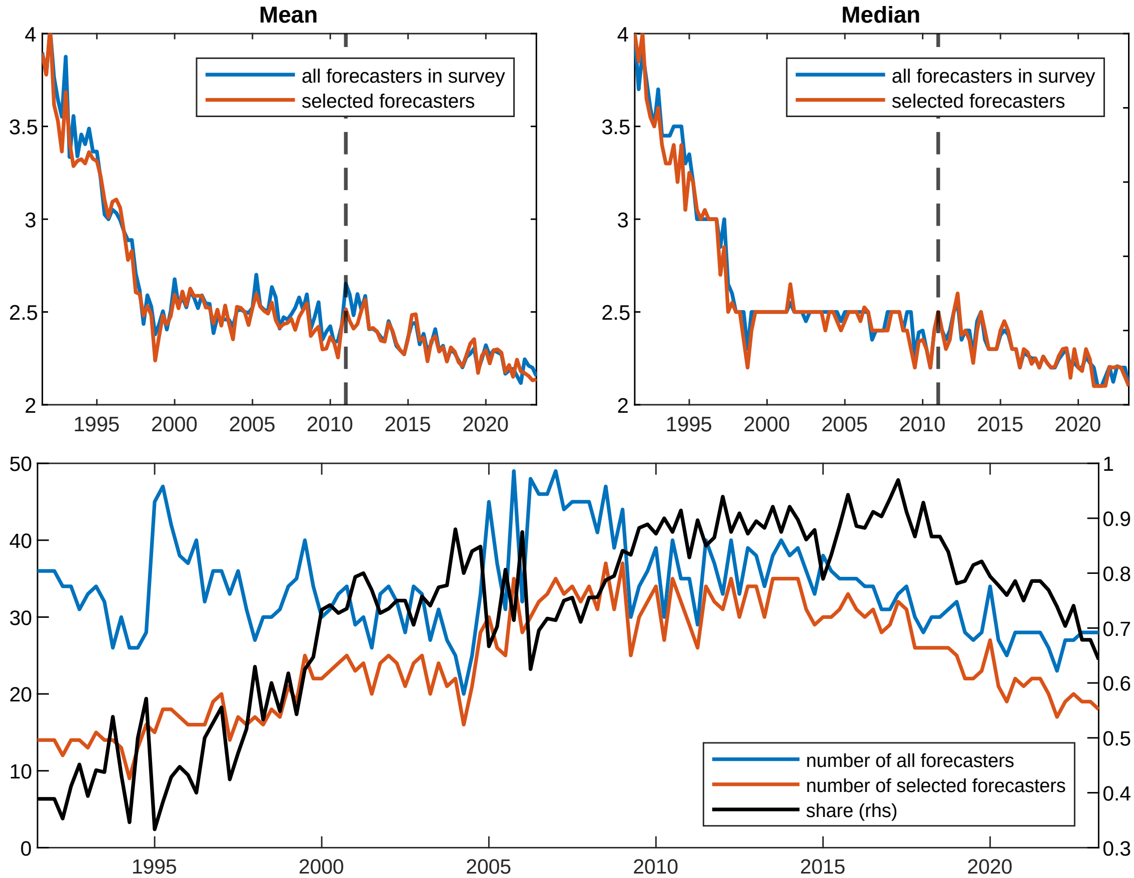

Figure 2: **Time series of aggregate inflation expectations (top) and number of total and selected forecasters in SPF (bottom).** Top panels: The dashed vertical line indicates 2011Q1 before which we use 10-year average expectations and afterwards 5-year 5-year expectations.

C The estimated trend-cycle model

The trend-cycle model is estimated using Bayesian methods via a Markov chain Monte Carlo (MCMC) algorithm. The priors and posterior mean for the time-invariant parameters of the trend-cycle model are shown in Table 1.[39] The posterior mean of the time-varying parameters are shown in Figure 3. 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 observed in the 1960s and early 1970s. We estimate that the volatility of trend inflation is just a touch below 30 bps at

[39]We 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).

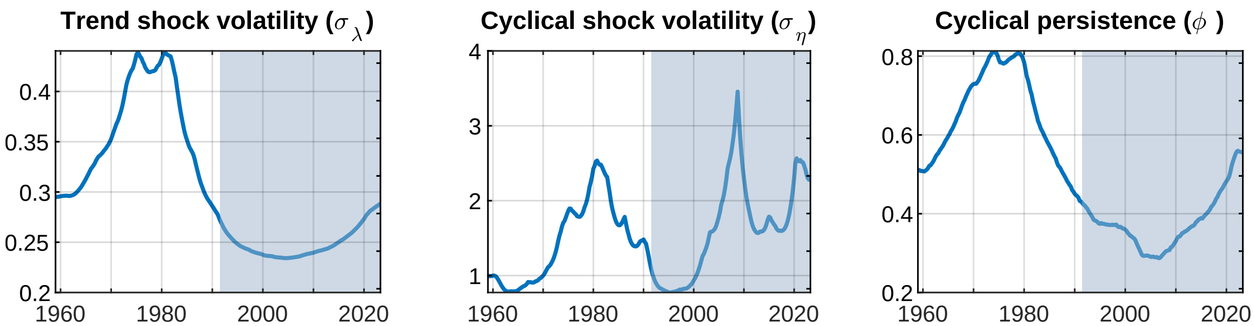

Figure 3: **Posterior mean for the time-varying parameters of the trend-cycle model.** The estimation of the three time-varying moments is conditional on quarterly U.S. headline CPI inflation rate observed over the period 1959Q1–2023Q2. Trend volatility and cyclical volatility are expressed in percentage. The shaded blue area indicates the sample period for the second-step panel estimation (1991Q3-2023Q2).

the end of our sample period (2023Q2).

Table 1: **Prior and posterior moments for the trend-cycle model parameters**

	Prior moments					Posterior moments	
Parameters	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 interquartile 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 3 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 2: **Forecasting performance of estimated trend**

	Full sample		SPF panel sample	
	10Y	5Y5Y	10Y	5Y5Y
I. Alternative priors				
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
II. Alternative model specifications				
No noise component	0.98	1.00	1.00	1.00
Constant parameters (tight prior)	1.20	1.24	1.48	1.25
Constant parameters (wide prior)	1.11	1.19	1.29	1.16
No persistence in cyclical component	1.43	1.37	2.13	1.62
5-year moving average	1.08	1.18	1.18	1.12
10-year moving average	0.93	1.08	1.11	1.10
Baseline model with core CPI	0.99	0.97	0.93	0.89
Constant parameters (from 1984)			1.04	1.05
Constant parameters (from 1984, lower σ^2_λ prior)			1.01	1.03

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 years 5 and 10 from today. This restricts the sample to end in 2013Q2. 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 9th and 10th row report the RMSE for the centered moving average of CPI inflation over a window of 5 and 10 years, respectively. The 11th row reports the RMSE for our baseline model but estimated using core CPI instead of headline CPI inflation. The last two rows report the RMSE for our baseline model with constant parameters but estimated only from 1984Q1. The 12th row assumes identical IG prior distributions for the cyclical and trend shock variance while the 13th row assumes a lower prior mean for the trend shock variance.

Table 2 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, 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 modeling framework. Third, modeling the cyclical component as i.i.d - 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. Fourth, our trend estimate performs better than a simple centered moving average which also includes forward-looking information. Fifth, using core CPI instead of headline CPI inflation leads to a slightly better forecast performance which is mainly driven by the mid 2000s. Sixth, estimating the trend with a constant parameter model only from the start of the Great Moderation in 1984 leads to a better forecasting performance than a constant parameter model which includes the Great Inflation of the 1970/80s but still performs slightly worse than our baseline model. Since the sample period for this forecast performance evaluation is restricted to end in 2013Q2, it does not use the trend estimates during the inflation surge period. During this period the trend estimates from the constant parameter models — also when estimated only from 1984 — peak at a significantly higher level relative to the estimates from our baseline model.

D Priors and posterior modes of the models' parameters

Table 3: **Prior distribution and posterior mode for the parameters**

	Prior moments					Posterior mode	
Parameters	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
$\alpha(i)$	G	2.00	1.00	2.00	[0.36,4.74]	Figure 2	Figure 2
$\rho(i)$	B	2.63	2.63	0.50	[0.17,0.83]	Figure 2	Figure 2
$\sigma_\nu(i)$	logN	0.00	1.00	1.65	[0.19,5.18]	Figure 2	Figure 2
$\rho_c^*(i)$	B	2.63	2.63	0.50	[0.17,0.83]	NA	Figure 3
$\alpha^*(i)$	G	2.00	1.00	2.00	[0.36,4.74]	NA	Figure 3
$\rho^*(i)$	B	2.63	2.63	0.50	[0.17,0.83]	NA	Figure 3
$\sigma_\nu^*(i)$	logN	0.00	1.00	1.65	[0.19,5.18]	NA	Figure 3

Notes: B, G and logN stand for beta, gamma, and log-normal distributions. Par(1) and Par(2) are the shape and scale for the gamma distribution, respectively, the shape parameters α and β for the beta distribution and the logarithm of location μ and logarithm of scale σ for the log-normal distribution.

The rational and behavioral model (second estimation step) are estimated via Bayesian likelihood estimation. In Table 3, we summarize prior moments and the posterior mode of the models' parameters. 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 $\ln \sigma^2_{c,t}$, which was denoted with γ_c, is set to 0.5.[40] In Table 3, 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 Additional estimation results

Figure 4 shows the impulse response functions to all four shocks comparing the rational and behavioral model.

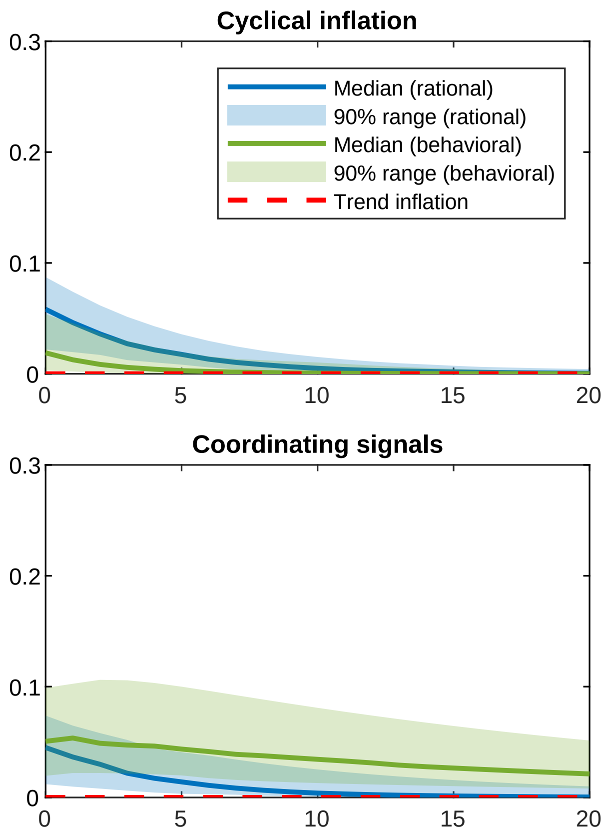

Figure 4: **Propagation of shocks to expectations.** Impulse response functions of long-run inflation expectations of every forecaster to a one-standard-deviation cyclical inflation shock (top left panel), trend inflation shock (top right panel), noise innovation to the coordinating signal (bottom left panel) and to every idiosyncratic signal (right panel). 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.

Figures 5 and 6 shows the estimation results where we replace our baseline trend-cycle

[40]This value allows for relatively quick changes in $\sigma_{c,t}$. Our results are robust to setting γ_c to a lower value of 0.1. The detailed results are available upon request.

model with a constant parameter trend-cycle model estimated over the sample since the beginning of the Great Moderation in 1984 (see also last row of Table 2). Based on such a simpler trend-cycle model we obtain the same type of misspecifications in the rational model and the behavioral model yields qualitatively similar estimation results regarding forecaster's overconfidence in private information and persistent expectations bias. We use a time-varying trend-cycle model as our baseline model since it is in line with a large empirical literature and especially in times such as the recent inflation episode a constant parameter trend-cycle model has important shortcomings.

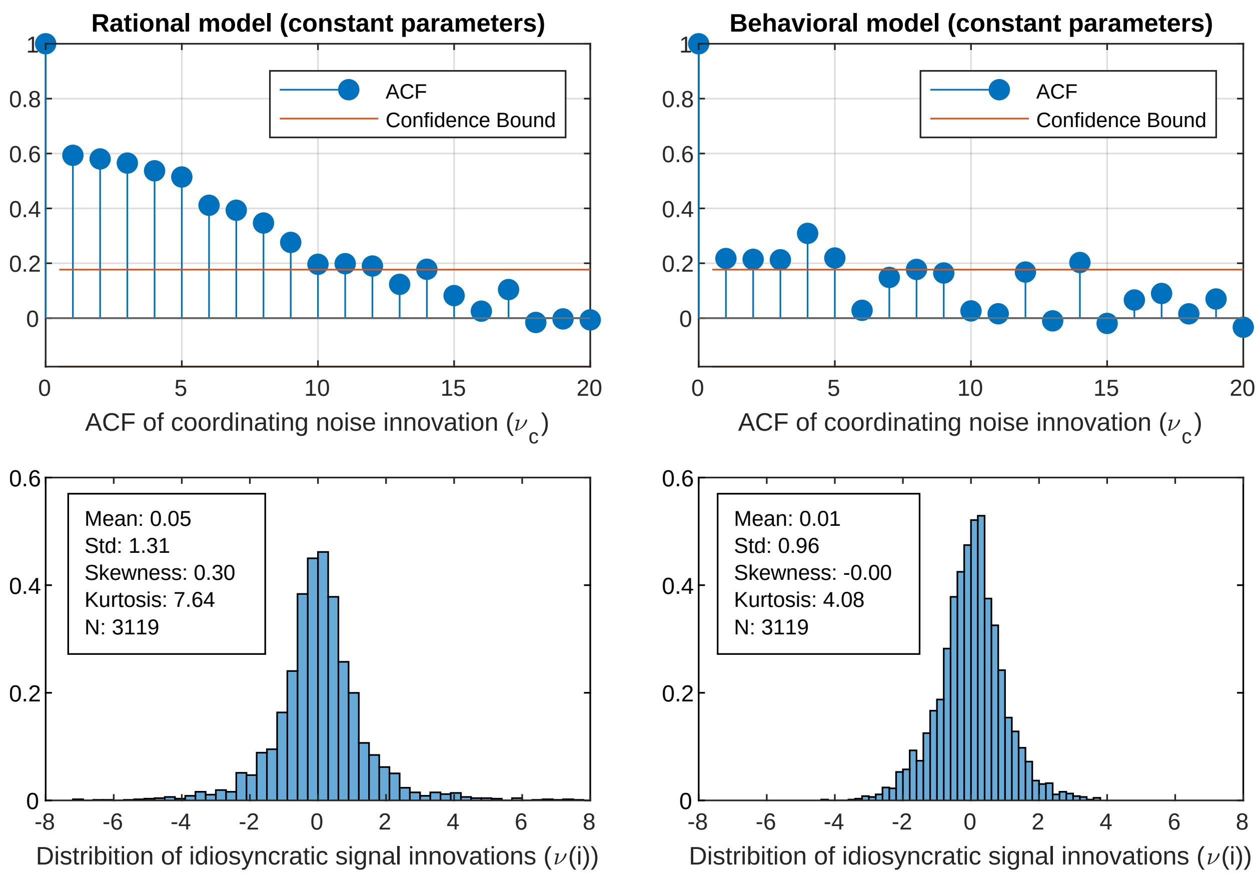

Figure 5: **Rational vs behavioral model: assessing the misspecification under inflation model with constant parameters estimated over sample from 1984.** The upper panels show the autocorrelogram of the estimated in-sample innovation to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The lower panels show the distribution of all the estimated in-sample innovations to the idiosyncratic signals over time and across forecasters. The estimated noise innovations are *filtered* estimates.

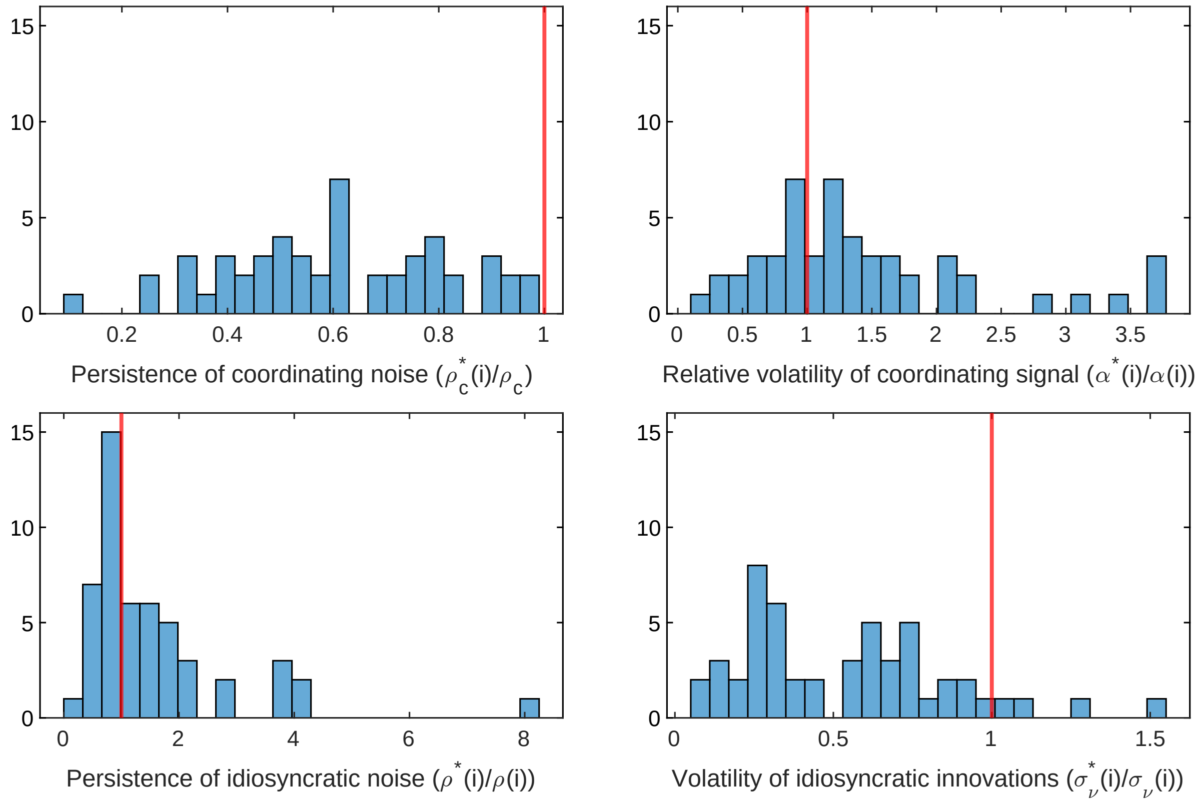

Figure 6: **Deviations from rationality with constant parameter trend-cycle model.** Ratio of the posterior mode of the behavioral parameters ($\rho_c^*(i)$, $\alpha^*(i)$, $\rho^*(i)$, $\sigma_\nu^*(i)$) to the posterior mode of the true value of the corresponding structural parameters (ρ_c, $\alpha(i)$, $\rho(i)$, 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.

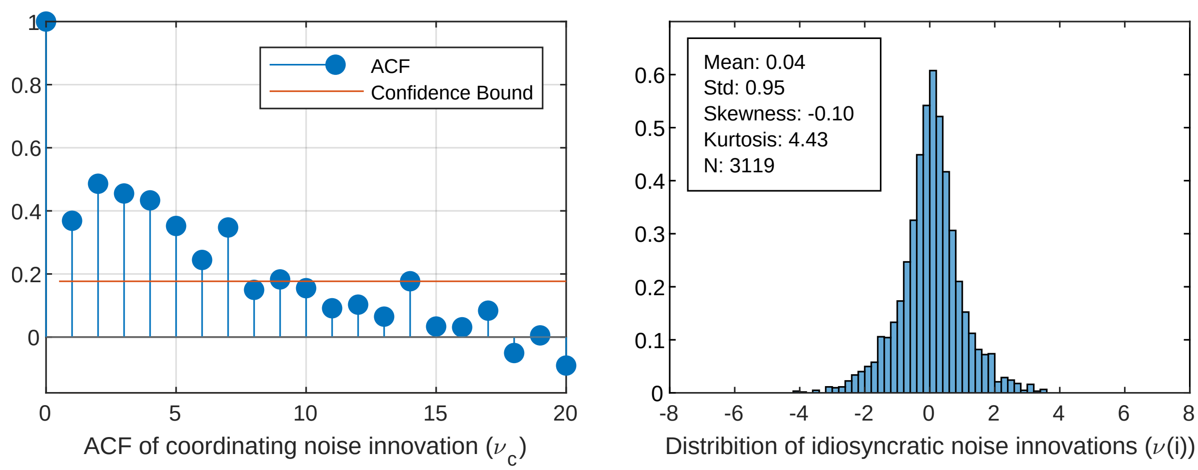

Figure 7: **Behavioral model only with overconfidence: assessing the misspecification.** The left panel shows the autocorrelogram of the estimated in-sample innovation to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The right panel shows the distribution of all the estimated in-sample innovations to the idiosyncratic signals over time and across forecasters. The estimated noise innovations are *filtered* estimates. 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 Robustness of misspecification of rational model

In the following we provide 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 both cases we use filtered estimates of the innovations. We focus on two aspects for the robustness: (i) allowing forecasters to learn about the parameters of the trend-cycle model in real-time and (ii) alternative model specifications for the estimation of trend inflation.

Real-time estimation of the trend-cycle model In our baseline estimation we assume forecasters know the parameters of the trend-cycle model which are estimated using information from the full sample. An alternative – arguably more realistic assumption – would be that forecasters estimate the parameters based on the information they have in real time. In Figure 8 we assume that forecasters do not know the parameters from our baseline trend-cycle model and instead estimate them in real-time. This means 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. Alternatively, we assume that forecasters estimate the parameters in real-time but based on the simpler constant parameter model included also in the last row of Table 2 and estimated since the beginning of the Great Moderation in 1984. The idea here is that parameter learning might be different in a time-varying setting relative to a constant parameter model which by construction is more backward-looking and hence slower in adjusting. The results for this analysis are shown in Figure 9. In the lhs panel we also use the constant parameter model to obtain the "true" estimate of the trend such that the parameter estimates between forecasters and the econometrician converge towards the end of the sample as forecasters have the same information set. In the rhs panel we assume that the "true" estimate of the trend comes from our baseline model such that there is also a role for potential discrepancies between constant and time-varying parameter models even with the same information set. In all three cases the rational model still exhibits the same two dimensions of misspecification.

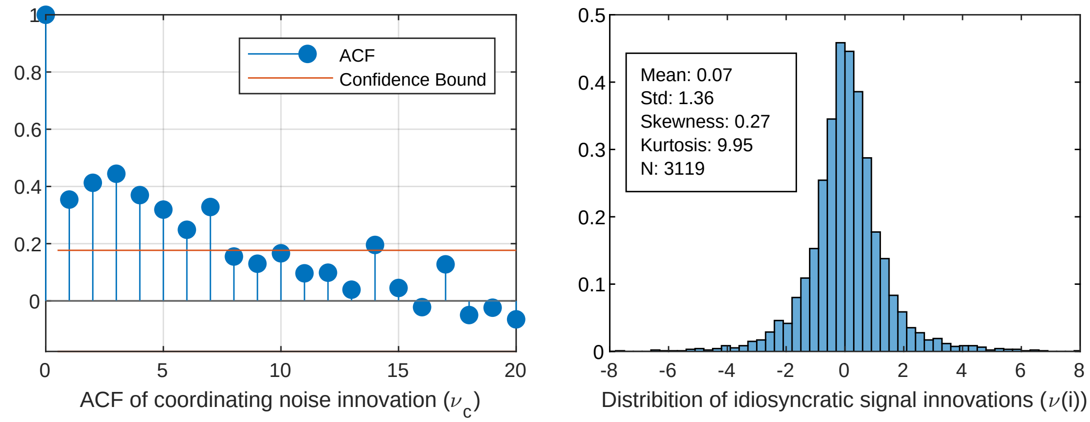

Figure 8: **Rational model: assessing the misspecification under real-time estimation of trend-cycle model.** The upper panels show the autocorrelogram of the estimated in-sample innovation to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The lower panels show the distribution of all the estimated in-sample innovations to the idiosyncratic signals over time and across forecasters. The estimated noise innovations are *filtered* estimates.

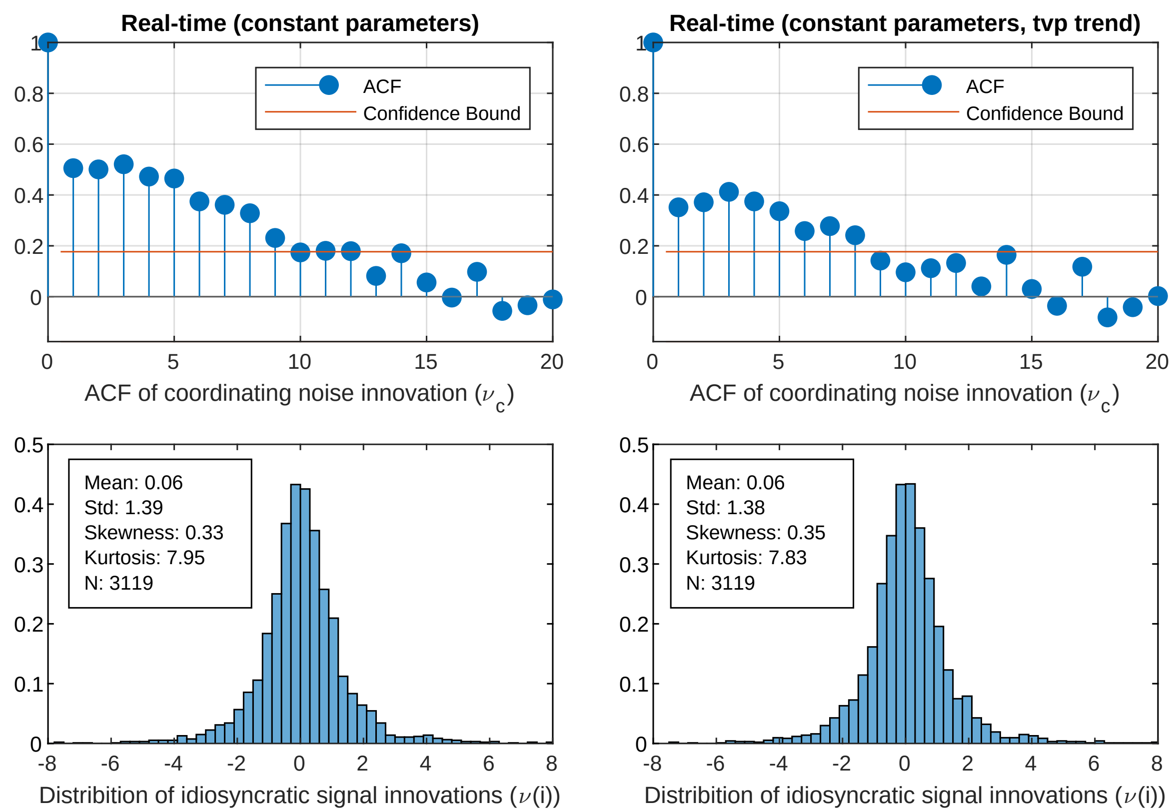

Figure 9: **Rational model: assessing the misspecification under real-time estimation with constant parameters over sample from 1984.** The upper panels show the autocorrelogram of the estimated in-sample innovation to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The lower panels show the distribution of all the estimated in-sample innovations to the idiosyncratic signals over time and across forecasters. The estimated noise innovations are *filtered* estimates. Both figures correspond to the rational model and forecasters using a constant parameter trend-cycle model estimated in real-time with sample start in 1984. In the lhs figure the “true” trend is also based on a constant parameter trend-cycle model but estimated over the full sample since 1984. The rhs figure instead assumes that the “true” trend is from our baseline trend-cycle model with time-varying parameters.

Alternative model specifications for the estimation of trend inflation In our baseline model the trend comes from an estimated trend cycle model with time-varying parameters. Instead, in the following figures we focus on some of the alternative model specifications also presented in the second panel of Table 2. In Figure 10 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 11, we estimate the trend using core CPI inflation instead of headline CPI inflation. In the right panel of Figure 11, we estimate the trend using our baseline model but assuming constant parameters and only start the estimation at the beginning of the Great Moderation in 1984. In all four 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.

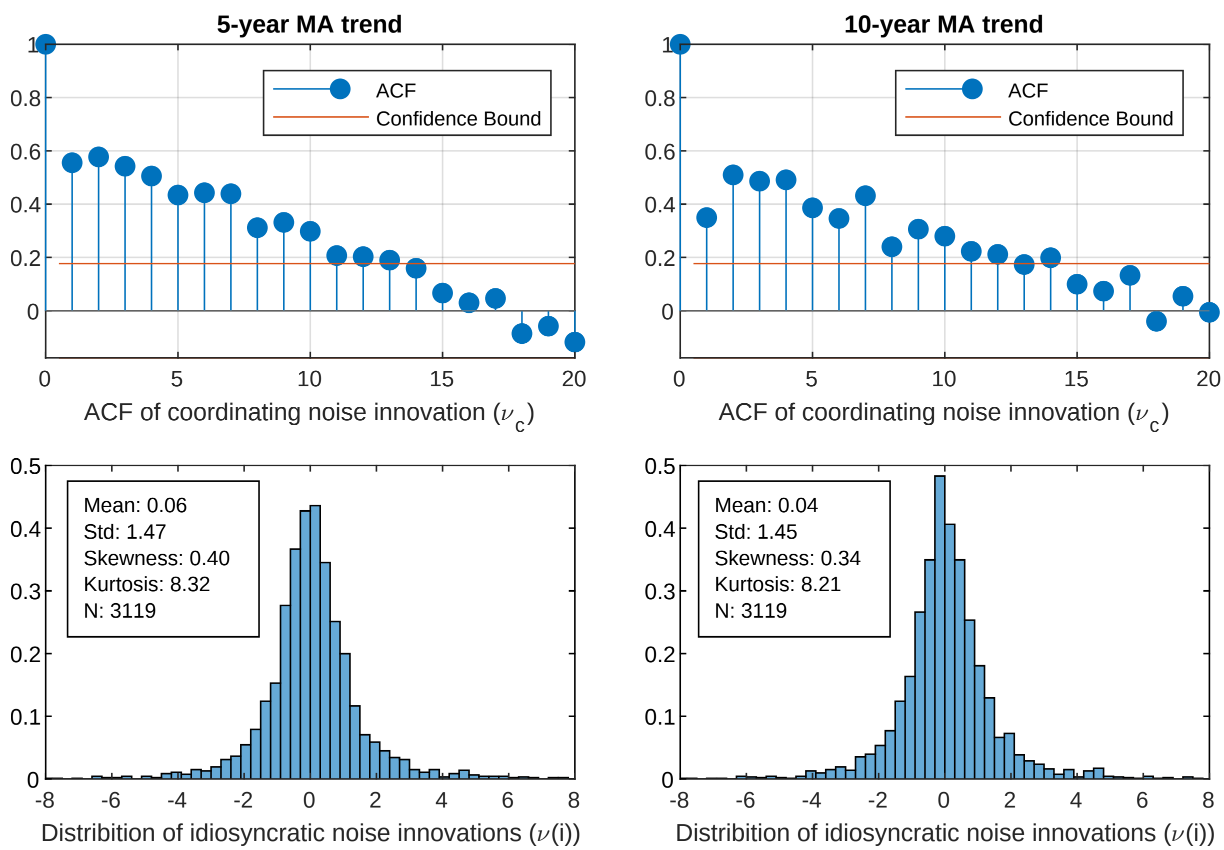

Figure 10: **Rational model: assessing the misspecification under moving average trend.** The upper panels show the autocorrelogram of the estimated in-sample innovation to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The lower panels show the distribution of all the estimated in-sample innovations to the idiosyncratic signals over time and across forecasters. The estimated noise innovations are *filtered* estimates. 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.

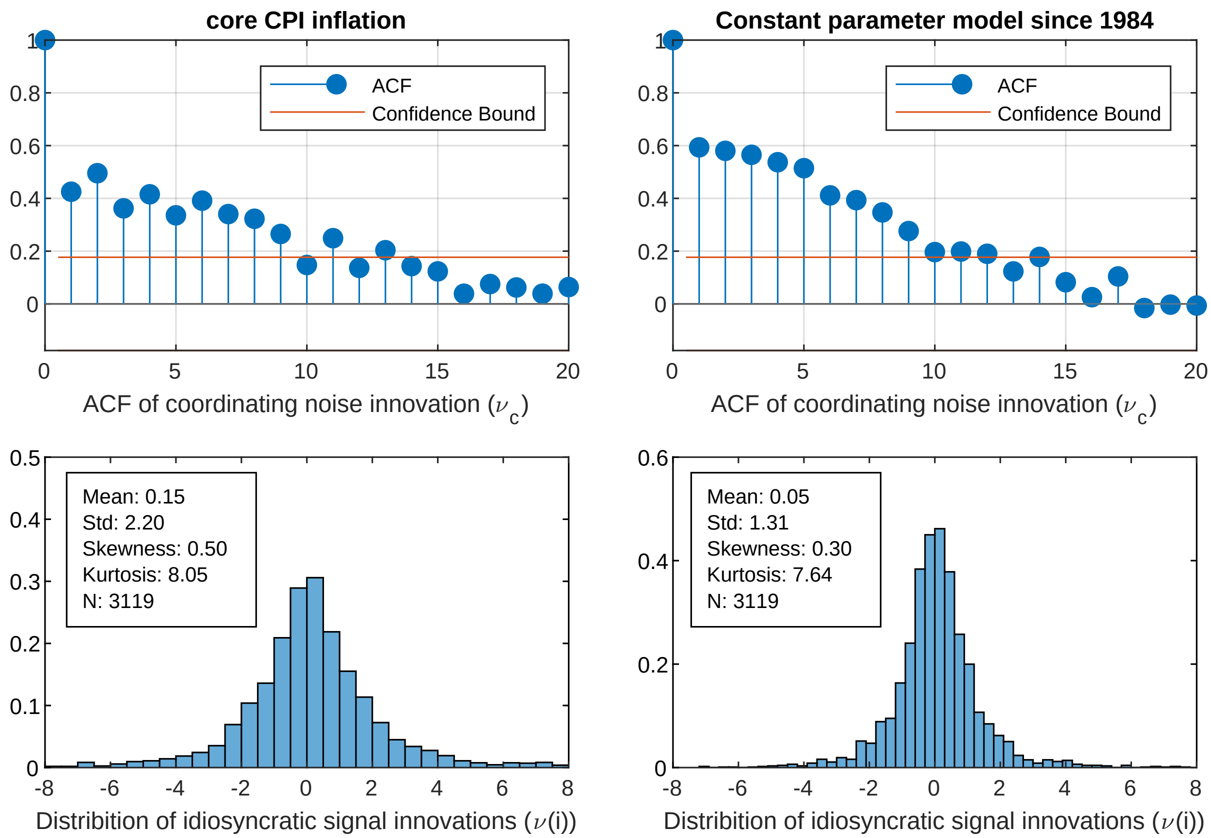

Figure 11: **Rational model: assessing the misspecification under core inflation and constant parameter trend-cycle model (estimated from 1984).** The upper panels show the autocorrelogram of the estimated in-sample innovation to the coordinating signals. The horizontal red lines are the 95-percent confidence bands for statistical significance. The lower panels show the distribution of all the estimated in-sample innovations to the idiosyncratic signals over time and across forecasters. The estimated noise innovations are *filtered* estimates. The lhs figure corresponds to the rational model with core CPI inflation instead of headline CPI inflation. The rhs figure is based on a trend-cycle model with constant parameters estimated only over the sample from 1984.

G A joint estimation approach

In the following we describe an alternative estimation approach that would allow for a joint estimation of all model parameters from the trend-cycle model and forecasters' signal extraction problems. For computational reasons, we focus on a constant parameter trend-cycle model (see last row of Table 2). In Figures 5 and 6 we have shown that such a model yields robust findings compared to our baseline model. In terms of notation let's define the parameter vector of the trend-cycle model as θ^1 and the parameter vector of the behavioral model as θ^2. We use Bayesian maximum likelihood estimation based on the following steps:

(i) Conditional on θ^1, use the Kalman filter/smoother on the trend-cycle model to obtain the likelihood $f(y_{CPI}|\theta^1)$ and an estimate of the trend and cyclical component of CPI inflation.

(ii) Impose the smoothed estimates of the trend and cyclical component of CPI inflation as observables within the data matrix Y (see lhs of equation (16).

(iii) Conditional on θ^1 and θ^2, use the Kalman filter on the full behavioral forecaster model to get the likelihood f(Y|θ^1,θ^2)

(iv) Define joint log-posterior as sum of the two log-likelihoods and the sum of all the log-prior probability distributions evaluated at θ^1 and θ^2, respectively.

(v) Optimize over this objective function to find the posterior mode for θ^1 and θ^2.

It is important to note that this approach still uses the smoothed estimate of trend inflation but its estimate is also influenced by the expectations data since all parameters are optimized based on the joint likelihood.

H Relationship to diagnostic expectations

Our paper documents two features of survey forecasts on long-run inflation expectations: overreaction to news and persistent forecast errors. A common approach in the literature to rationalize overreaction is diagnostic expectations (DE; Bordalo et al. (2018)). This section shows that in a simple one-signal Gaussian environment, one can establish a mapping between our misperception parameter ϕ (overconfidence) and the diagnostic parameter θ under some conditions. With multiple signals, the equivalence breaks unless the forecaster is symmetrically overconfident across all signals. We further show that DE — like the rational benchmark — can generate persistent forecast errors in a one-signal environment or in multi-signal environments where noise persistence is symmetric across signals. When persistence differs across signals, however, both DE and rational forecast errors die out quickly, whereas our misperception mechanism (persistent expectations bias) continues to deliver persistent forecast errors.

H.1 Overconfidence: one-signal model

Data generating process Assume for simplicity there is only one forecaster and the data generating process is given by[41]

$$x_t = x_{t-1} + u_t, \qquad u_t \sim \mathcal{N}(0, \sigma_u^2), \tag{26}$$

[41] All qualitative conclusions below are unchanged if x_t follows a stable AR(1) process $x_t = \alpha x_{t-1} + u_t$ with $|\alpha| < 1$ rather than a random walk; the closed-form expressions for the stationary-regime $\bar{P}$ and the $\phi \leftrightarrow \theta$ mapping change but the structure of the equivalence does not.

and the signal is given by

$$s_t = x_t + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2). \tag{27}$$

Rational benchmark The rational forecaster correctly perceives σ_u and σ_ε. The Kalman recursion is

$$\begin{aligned}
x^{\text{rat}}_{t|t-1} &= x^{\text{rat}}_{t-1|t-1} && (28)\\
P^{\text{rat}}_{t|t-1} &= P^{\text{rat}}_{t-1|t-1} + \sigma_u^2 && (29)\\
F^{\text{rat}}_{t|t-1} &= P^{\text{rat}}_{t|t-1} + \sigma_\varepsilon^2 && (30)\\
x^{\text{rat}}_{t|t} &= x^{\text{rat}}_{t|t-1} + P^{\text{rat}}_{t|t-1}\, F^{\text{rat},-1}_{t|t-1} \left(s_t - x^{\text{rat}}_{t|t-1}\right) && (31)\\
&= x^{\text{rat}}_{t|t-1} + \frac{P^{\text{rat}}_{t-1|t-1} + \sigma_u^2}{P^{\text{rat}}_{t-1|t-1} + \sigma_u^2 + \sigma_\varepsilon^2} \left(s_t - x^{\text{rat}}_{t|t-1}\right). && (32)
\end{aligned}$$

Because the state is a random walk, $x^{\text{rat}}_{t|t}$ itself does not converge; the forecast-error variance $P^{\text{rat}}_{t|t}$ does, however, converging to

$$\bar{P} = \frac{-\sigma_u^2 + \sqrt{\sigma_u^4 + 4\sigma_u^2\sigma_\varepsilon^2}}{2}. \tag{33}$$

We refer to this as the *steady state* of the Kalman filter: the state x_t itself remains non-stationary, but the filter's one-step-ahead forecast-error variance—and hence its Kalman gain—stabilize.[42] The steady-state rational gain is

$$\bar{K} = \frac{\bar{P} + \sigma_u^2}{\bar{P} + \sigma_u^2 + \sigma_\varepsilon^2}. \tag{34}$$

Both cognitive deviations we study below can be viewed as perturbations of this gain: overconfidence distorts it indirectly by misrepresenting σ_ε, while diagnostic expectations distort it directly by scaling it by $(1 + \theta)$.

Overconfidence The forecaster misperceives the signal precision: she perceives the signal-noise volatility to be $\sigma_\varepsilon^* = \phi\sigma_\varepsilon$, with $\phi < 1$ corresponding to overconfidence (the signal is perceived as more informative than it truly is) and $\phi > 1$ to underconfidence.[43] We will refer

[42] We emphasize this to avoid a possible source of confusion: "steady state" here refers to convergence of the filter's covariance recursion (the Riccati equation), not of the state variable x_t, which continues to follow a random walk.

[43] This formalization of overconfidence dates back to Daniel et al. (1998) in the context of private signals. Scheinkman and Xiong (2003) apply the same formal device to public signals. The interpretation of ϕ as a cognitive parameter does not require the signal to be private.

to this as the "overconfidence" model (with superscript $*$) throughout, in contrast to the "diagnostic" (DE) model introduced below. The prediction step is identical to the rational one; the update step replaces σ_ε^2 with $\sigma_\varepsilon^2\phi^2$ in the gain:

$$x^*_{t|t} \quad = \quad x^*_{t|t-1} + \frac{P^*_{t-1|t-1} + \sigma_u^2}{P^*_{t-1|t-1} + \sigma_u^2 + \sigma_\varepsilon^2\phi^2}\left(s_t - x^*_{t|t-1}\right). \tag{35}$$

Because the misperception enters the variance-updating equation, $P^*_{t|t}$ converges to a steady-state value that differs from the rational one,

$$\bar{P}^* = \frac{-\sigma_u^2 + \sqrt{\sigma_u^4 + 4\sigma_u^2\sigma_\varepsilon^2\phi^2}}{2}. \tag{36}$$

If $\phi < 1$ we have $\bar{P}^* < \bar{P}$; if $\phi > 1$ we have $\bar{P}^* > \bar{P}$. The steady-state overconfidence gain is

$$\bar{K}^* = \frac{\bar{P}^* + \sigma_u^2}{\bar{P}^* + \sigma_u^2 + \sigma_\varepsilon^2\phi^2}. \tag{37}$$

Diagnostic expectations Under DE, the forecaster correctly perceives σ_ε but distorts the posterior by over-weighting states that have become more representative in light of the news. For Gaussian posteriors, this is equivalent to scaling the rational Kalman gain by $(1+\theta)$, where $\theta \geq 0$ parameterizes the strength of the diagnostic distortion. The prediction step is identical to the rational one; the update step is

$$x^{\mathrm{DE}}_{t|t} \quad = \quad x^{\mathrm{DE}}_{t|t-1} + (1+\theta)\frac{P^{\mathrm{DE}}_{t-1|t-1} + \sigma_u^2}{P^{\mathrm{DE}}_{t-1|t-1} + \sigma_u^2 + \sigma_\varepsilon^2}\left(s_t - x^{\mathrm{DE}}_{t|t-1}\right). \tag{38}$$

The diagnostic parameter does not enter the variance-updating equation, so the DE forecast-error variance coincides with the rational one: $P^{\mathrm{DE}}_{t|t}$ converges to $\bar{P}$ as given in equation (33). The steady-state DE gain is therefore

$$\bar{K}^{\mathrm{DE}} = (1+\theta)\,\bar{K} = (1+\theta)\frac{\bar{P} + \sigma_u^2}{\bar{P} + \sigma_u^2 + \sigma_\varepsilon^2}. \tag{39}$$

Equivalence We ask whether there exists a mapping from ϕ to θ such that the two models produce identical belief paths $\{x_{t|t}\}$ for every realization of the innovations $\{u_t, \varepsilon_t\}$. At the filter's steady state the answer is yes: both filters reduce to a constant-gain updating rule on the innovation $s_t - x_{t-1|t-1}$, so matching the steady-state Kalman gains $\bar{K}^* = \bar{K}^{\mathrm{DE}}$ ensures that the responses of $x_{t|t}$ to a state shock u_t and to a signal-noise shock ε_t coincide simultaneously across the two models.

Solving $\bar{K}^* = \bar{K}^{\mathrm{DE}}$ for θ yields

$$\theta \;=\; \frac{\bar{P}^*\sigma_\varepsilon^2 - \bar{P}\sigma_\varepsilon^2\phi^2 + \sigma_\varepsilon^2\sigma_u^2(1-\phi^2)}{(\bar{P}+\sigma_u^2)(\bar{P}^*+\sigma_u^2+\sigma_\varepsilon^2\phi^2)}. \tag{40}$$

A value $\phi > 1$ corresponds to underconfidence and delivers $\theta < 0$; $\phi < 1$ corresponds to overconfidence and delivers $\theta > 0$. For $\phi = 1$, $\bar{P}^* = \bar{P}$ and $\theta = 0$, recovering the rational model.

Discussion Two caveats bound the scope of this result. First, the mapping delivers exact equivalence only when the filter is initialized at the steady state. Off the steady state, the two filters follow different Riccati recursions—$P^*_{t|t}$ depends on $\sigma_\varepsilon^2\phi^2$ while $P^{\mathrm{DE}}_{t|t}$ depends on σ_ε^2—so their Kalman gains differ during the transient. Matching them period by period would require a time-varying θ_t. Second, equation (40) is derived under constant σ_u and σ_ε. In our baseline model the volatility of trend inflation is time-varying, which corresponds to time-variation in σ_u here. The filter then never reaches a constant steady state, so exact equivalence again requires θ to be time-varying; equivalence with a constant θ becomes an approximation, reasonable as long as parameter variation is modest.

Numerical verification Using equation (40) in equation (38), we simulate the overconfidence, diagnostic, and rational filters in response to one-standard-deviation shocks to the signal noise ε and to the state innovation u. Each filter is initialized at its own steady-state forecast-error variance: $\bar{P}$ for the rational and diagnostic filters, $\bar{P}^*$ for the overconfidence filter. We set the parameters as follows: $\sigma_u = 0.25$, $\sigma_\varepsilon = 0.5$, $\phi = 0.75$, and θ implied by equation (40). Figure 12 illustrates that both the overconfidence and the diagnostic models lead to an overreaction relative to the rational benchmark, and that their responses are identical under both shocks (the overconfidence and diagnostic lines overlap). This confirms that the two models produce identical belief paths.

H.2 Overconfidence: multiple signals

The mapping established in Section H.1 is tight because the forecaster has only one signal, so the overconfidence wedge has a single degree of freedom—an inflated Kalman gain—that DE can replicate. With multiple signals this is no longer automatic: overconfidence can be asymmetric across signals, whereas DE applies a common scalar distortion to all signal weights. We make the point in a static Bayesian updating problem for transparency and confirm numerically at the end of the subsection that it carries over to the dynamic filter.

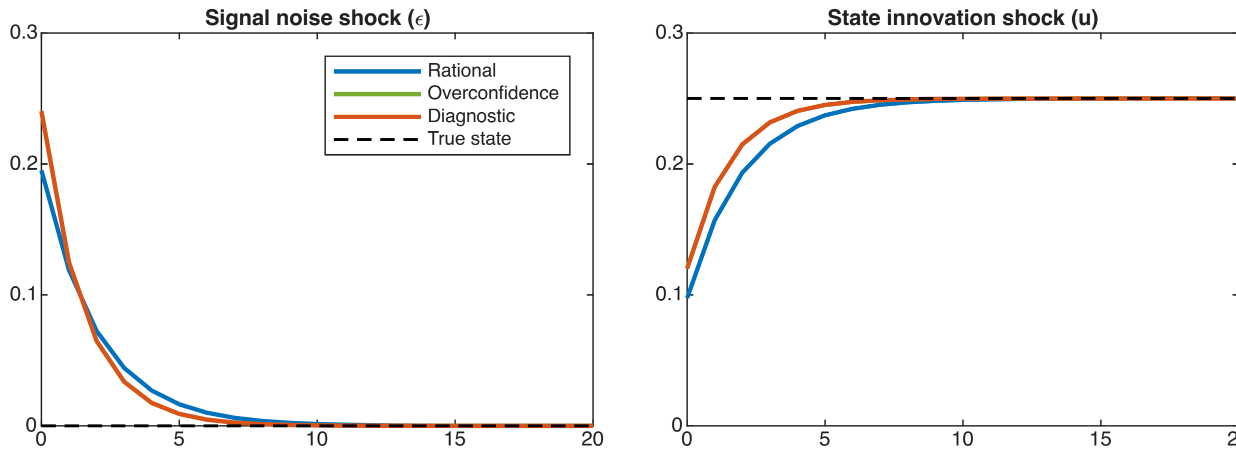

Figure 12: **Overconfidence and diagnostic expectations: one-signal model.** Impulse response functions of beliefs about x_t to a one-standard-deviation shock to the signal noise ε (left panel) and to the state innovation u (right panel), under the rational model, the overconfidence model with $\phi = 0.75$, and the diagnostic expectations model with θ implied by equation (40). The black dashed line shows the true response of x_t. Parameters are set to $\sigma_u = 0.25$, $\sigma_\varepsilon = 0.5$. Each filter is initialized at its own steady-state forecast-error variance ($\bar{P}$ for the rational and diagnostic filters, $\bar{P}^*$ for the overconfidence filter). Shocks hit in period 0. The overconfidence and diagnostic lines overlap in both panels.

Setup The forecaster has Gaussian prior $x \sim \mathcal{N}(0, \sigma_x^2)$ and observes two signals,

$$s_i = x + \varepsilon_i, \qquad \varepsilon_i \sim \mathcal{N}(0, \sigma_{\varepsilon_i}^2), \qquad i \in \{1, 2\}, \tag{41}$$

with $\varepsilon_1 \perp \varepsilon_2$. Write $\tau_0 = 1/\sigma_x^2$ for the prior precision and $\tau_i = 1/\sigma_{\varepsilon_i}^2$ for signal precisions.

Rational benchmark The rational forecaster correctly perceives σ_x and $(\sigma_{\varepsilon_1}, \sigma_{\varepsilon_2})$. Standard Bayesian updating yields a Gaussian posterior with mean

$$\mathbb{E}[x \mid s_1, s_2] = \frac{\tau_1}{\tau_0 + \tau_1 + \tau_2}\, s_1 + \frac{\tau_2}{\tau_0 + \tau_1 + \tau_2}\, s_2. \tag{42}$$

The rational posterior mean is a precision-weighted average of the signals, with the prior centered at zero contributing no shift but absorbing a fraction $\tau_0/(\tau_0 + \tau_1 + \tau_2)$ of the weight.

Overconfidence The forecaster misperceives the noise volatility of each signal: she perceives $\sigma_{\varepsilon_i}^* = \phi_i \sigma_{\varepsilon_i}$ for $i \in \{1, 2\}$, with $\phi_i < 1$ corresponding to overconfidence about signal i. Replacing τ_i by the perceived precision $\tau_i^* = \tau_i/\phi_i^2$ in equation (42) yields

$$\mathbb{E}^{\phi_1, \phi_2}[x \mid s_1, s_2] = \frac{\tau_1/\phi_1^2}{\tau_0 + \tau_1/\phi_1^2 + \tau_2/\phi_2^2}\, s_1 + \frac{\tau_2/\phi_2^2}{\tau_0 + \tau_1/\phi_1^2 + \tau_2/\phi_2^2}\, s_2. \tag{43}$$

The key feature of this expression is that the overconfidence wedges ϕ_1, ϕ_2 carry signal-specific indices: nothing in the setup forces $\phi_1 = \phi_2$.

Diagnostic expectations Under DE, the forecaster correctly perceives all noise variances but scales each of the rational posterior weights by $(1+\theta)$ (see also the online appendix of Bordalo et al., 2020):

$$\mathbb{E}^{\theta}[x \mid s_1, s_2] = (1+\theta)\frac{\tau_1}{\tau_0+\tau_1+\tau_2}\, s_1 + (1+\theta)\frac{\tau_2}{\tau_0+\tau_1+\tau_2}\, s_2. \tag{44}$$

When the mapping holds Matching equations (43) and (44) signal by signal gives

$$\frac{\tau_i/\phi_i^2}{\tau_0+\tau_1/\phi_1^2+\tau_2/\phi_2^2} = (1+\theta)\frac{\tau_i}{\tau_0+\tau_1+\tau_2}, \qquad i = 1, 2. \tag{45}$$

Taking the ratio of the two conditions yields $1/\phi_1^2 = 1/\phi_2^2$, so $\phi_1 = \phi_2 \equiv \phi$: the forecaster must be *symmetrically* overconfident about both signals. Setting $\phi_1 = \phi_2$ then collapses equation (43) to the single-signal form with combined precision $\tau_1 + \tau_2$, and the one-signal mapping of Section H.1 delivers θ.

When the mapping fails If overconfidence is asymmetric, $\phi_1 \neq \phi_2$, no scalar θ solves equation (45) exactly. To see this directly, suppose $\phi_1 < \phi_2$—i.e., the forecaster is more overconfident about s_1 than s_2. Then the weight on s_1 in equation (43) exceeds its rational counterpart while the weight on s_2 falls below. Increased confidence in s_1 inflates its weight *and* mechanically crowds out reliance on s_2 in the Bayesian combination. DE, by contrast, scales *both* weights in tandem: no $(1+\theta)$ can simultaneously raise the weight on s_1 and lower the weight on s_2 relative to rational. The two belief distortions are therefore formally distinct whenever overconfidence varies across signals; how large the quantitative difference is between the two in practice is ultimately an empirical question.

Implications for our estimated model In our own model, forecasters observe three signals—inflation, a coordinating signal, and an idiosyncratic signal—and the estimated misperceptions vary across them. Forecasters are asymmetrically overconfident. By the argument above, no constant-θ DE formulation can exactly reproduce the parameterization our estimation identifies in the data.

Numerical illustration To confirm that the asymmetric-overconfidence argument carries over from the static problem to the filtering environment of Section H.1, we extend the one-signal Kalman filter to two signals and simulate the asymmetric case. Signal-noise volatilities are equal, $\sigma_{\varepsilon_1} = \sigma_{\varepsilon_2} = 0.5$, and overconfidence is applied only to the first signal ($\phi_1 < 1$, $\phi_2 = 1$). The diagnostic parameter θ is chosen so that the on-impact response of overconfidence and DE to a shock to ε_1 coincide. Figure 13 shows responses of beliefs to shocks to ε_1, ε_2, and u. Under a shock to ε_1 (left panel), overconfidence and DE coincide on impact by construction but diverge thereafter. Under a shock to ε_2 (middle panel), the two diverge in opposite directions relative to the rational benchmark: the overconfident forecaster overweights s_1 and thereby crowds out reliance on s_2, so the overconfidence response lies *below* rational; DE scales the weight on s_2 up by $(1 + \theta)$, so the DE response lies *above* rational. Under a shock to u (right panel), DE's uniform up-scaling delivers a larger total gain than overconfidence's selective inflation of τ_1^*, so the DE response exceeds the overconfidence response. No constant θ can exactly match all three responses simultaneously: this is the dimensional mismatch identified above, surviving in the dynamic filter.

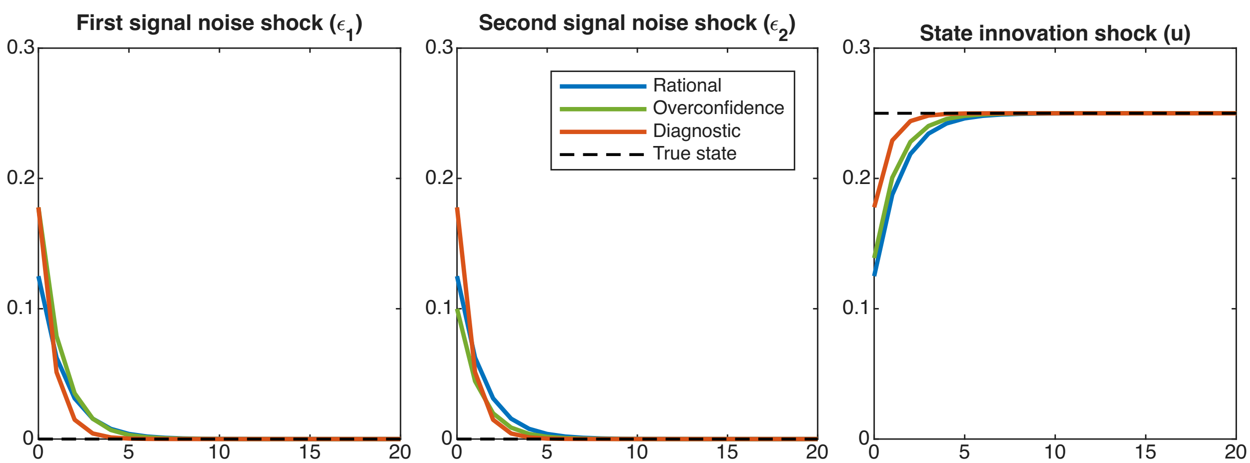

Figure 13: **Overconfidence and diagnostic expectations: two-signal dynamic filter.** Impulse response functions of beliefs about x_t to one-standard-deviation shocks to the first signal noise ε_1 (left panel), the second signal noise ε_2 (middle panel), and the state innovation u (right panel), in a two-signal extension of the Section H.1 setup, under the rational model, an overconfidence model with asymmetric overconfidence ($\phi_1 < 1$, $\phi_2 = 1$), and a diagnostic expectations model. The black dashed line shows the true response of x_t. Parameters are set to $\sigma_u = 0.25$, $\sigma_{\varepsilon_1} = \sigma_{\varepsilon_2} = 0.5$; θ is chosen so that the overconfidence and diagnostic responses coincide on impact in the left panel. Each filter is initialized at its own steady-state forecast-error variance. Shocks hit in period 0.

H.3 Persistent expectations bias

Sections H.1 and H.2 considered misperception of signal *precision* (overconfidence). We now consider misperception of signal *persistence*: the noise is autoregressive, and the forecaster correctly perceives its variance but believes it dies out faster than it actually does.

Data generating process The state and signal are now

$$
\begin{aligned}
x_t &= x_{t-1} + u_t, \qquad u_t \sim \mathcal{N}(0, \sigma_u^2), && (46)\\
v_t &= \rho_v \, v_{t-1} + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2), && (47)\\
s_t &= x_t + v_t, && (48)
\end{aligned}
$$

with $|\rho_v| < 1$ and innovations mutually and serially independent.

Rational benchmark The state is now $z_t = (x_t, v_t)'$, with

$$
z_t = A z_{t-1} + \begin{pmatrix} u_t \\ \varepsilon_t \end{pmatrix}, \qquad s_t = C z_t, \tag{49}
$$

where $A = \text{diag}(1, \rho_v)$ and $C = (1, 1)$. The rational forecaster runs the Kalman filter on this two-dimensional system; the posterior mean $\hat{z}_{t|t}$ converges to a steady-state recursion with constant gain vector $\bar{K} \in \mathbb{R}^2$, and the belief about the fundamental is $\hat{x}_{t|t} = e_1' \hat{z}_{t|t}$ where $e_1 = (1, 0)'$.

Persistent expectations bias The forecaster correctly perceives σ_u and σ_ε but misperceives the persistence of the noise: she believes $\rho_v^* = \phi_\rho \rho_v$ with $\phi_\rho < 1$.[44] As in Section H.1, the superscript $*$ denotes the forecaster's perceived parameters and resulting filter objects. Running the Kalman filter under the perceived DGP delivers a steady-state gain vector $\bar{K}^* \neq \bar{K}$ and posterior means $\hat{x}^*_{t|t} = e_1' \hat{z}^*_{t|t}$. The filter's transition matrix is also distorted ($A^* = \text{diag}(1, \phi_\rho \rho_v)$), so the prediction step itself differs from the rational one—not only the gain.

Diagnostic expectations Under DE, the forecaster correctly perceives the data generating process, including ρ_v. The DE distortion scales the rational Kalman gain by $(1 + \theta)$:

$$
\hat{z}^{\text{DE}}_{t|t} = \hat{z}^{\text{DE}}_{t|t-1} + (1 + \theta) \bar{K} \left(s_t - C \hat{z}^{\text{DE}}_{t|t-1} \right). \tag{50}
$$

The transition matrix A used in the prediction step is the rational one, so DE distorts only the gain whereas persistence misperception distorts both the gain and the dynamics; the two mechanisms can therefore generate different IRFs even when the impact response is matched.

[44] $\phi_\rho < 1$ corresponds to underestimating persistence: the forecaster expects past noise to die out faster than it does. The opposite case $\phi_\rho > 1$ is treated symmetrically; we focus on $\phi_\rho < 1$ because it is the empirically more relevant case for explaining persistent forecast biases.

Persistence in the IRFs Both DE and the persistent-expectations-bias filter can generate persistent IRFs of $\hat{x}_{t|t}$ following a signal-noise shock—and hence persistent forecast errors, since the true state is unchanged—but the underlying mechanisms differ. The persistence of the IRF is governed by the eigenvalues of $(I - \bar{K}C)A$, the matrix that governs how beliefs evolve from period to period under the Kalman filter. For the DE filter, the diagnostic distortion rescales the response to innovations but does not change the underlying belief dynamics. Across a wide range of parameter combinations, ρ_v is the dominant driver of medium-horizon IRF persistence: both rational and DE responses decay at essentially the same rate, with θ primarily affecting the initial impact response (overreaction) without reshaping the IRF.

The persistent-expectations-bias filter generates additional persistence through a different mechanism. Because the forecaster believes the noise dies out faster than it actually does, her filter's prediction $\hat{s}^*_{t|t-1} = CA^*\hat{z}^*_{t-1|t-1}$ understates the realized signal s_t, and her filter's innovation contains an excess term proportional to $\rho_v(1 - \phi_\rho)\, v_{t-1}$: a portion of past noise is mistaken for movements in the true state. Her filter applies the gain $\bar{K}^*$ to this innovation, so the gain on x converts the misattributed noise into a sustained revision of beliefs about the fundamental. The bias accumulates as long as the noise process v_{t-1} remains non-zero, producing a hump-shaped IRF.

Numerical illustration Figure 14 illustrates the responses for $\sigma_u = 0.25$, $\sigma_\varepsilon = 0.5$, $\rho_v = 0.9$, $\phi_\rho = 0.75$ (perceived persistence $\rho^*_v \approx 0.6$), and $\theta = 0.5$. All three models generate persistent IRFs of beliefs about x_t, but the persistent-expectations-bias model differs from rational and DE in the shape of those IRFs. Under a shock to the persistent signal noise (left panel), the bias model generates a pronounced hump, while rational and DE cluster together in a smoother decay.[45] Under a shock to the state innovation (right panel), the persistent-expectations-bias forecaster's beliefs converge to the truth more quickly than rational and DE—the same misattribution mechanism, working in the forecaster's favor when the persistent signal movement actually reflects a fundamental shift.

Extension to two signals The persistent-noise environment extends naturally to two signals:

$$v_t = \rho_v\, v_{t-1} + \varepsilon_{1,t}, \qquad \varepsilon_{1,t} \sim \mathcal{N}(0, \sigma^2_{\varepsilon_1}), \tag{51}$$

[45] The diagnostic IRF exhibits small early-period oscillations: the $(1 + \theta)$ scaling can push one eigenvalue of $(I - (1 + \theta)\bar{K}C)A$ slightly negative, producing alternating-sign decay before the dominant positive eigenvalue takes over; the same artifact appears in the multi-signal figure below. We have explored distant-memory DE following Bianchi et al. (2023), which yields smoother IRFs but leaves the qualitative conclusions about medium-horizon dynamics unchanged.

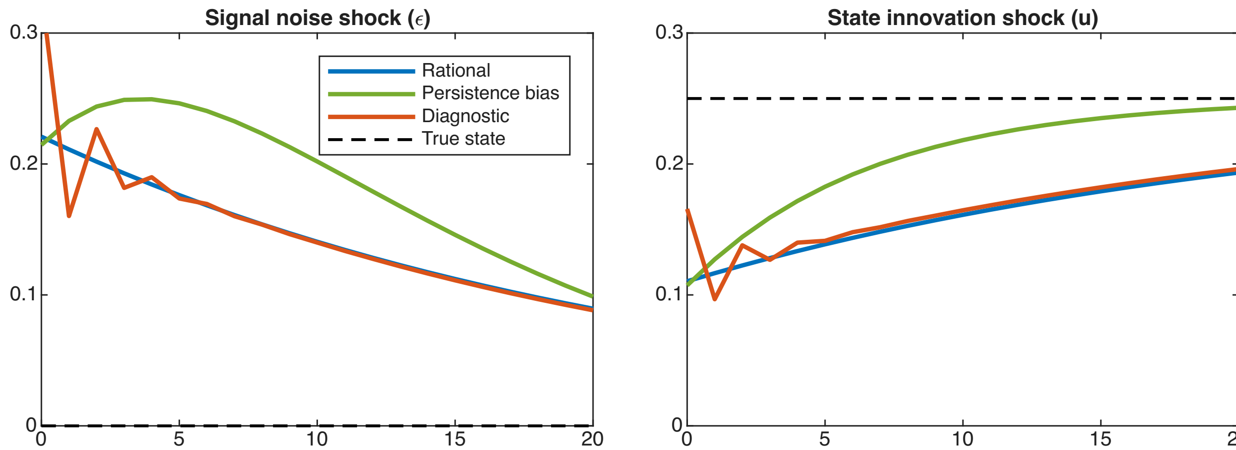

Figure 14: **Persistent expectations bias and diagnostic expectations: one signal.** Impulse response functions of beliefs about x_t to a one-standard-deviation shock to the signal noise (left panel) and to the state innovation u (right panel), under the rational model, a persistent-expectations-bias model with $\phi_\rho = 0.75$, and a diagnostic expectations model. Parameters: $\sigma_u = 0.25$, $\sigma_\varepsilon = 0.5$, $\rho_v = 0.9$. The diagnostic parameter is $\theta = 0.5$. The black dashed line shows the true response of x_t. Each filter is initialized at its own steady-state forecast-error covariance. Shocks hit in period 0.

$$w_t = \rho_w w_{t-1} + \varepsilon_{2,t}, \qquad \varepsilon_{2,t} \sim \mathcal{N}(0, \sigma^2_{\varepsilon_2}), \tag{52}$$

$$s_{1,t} = x_t + v_t, \qquad s_{2,t} = x_t + w_t, \tag{53}$$

with $|\rho_v|, |\rho_w| < 1$ and innovations mutually and serially independent. The forecaster correctly perceives ρ_w but misperceives ρ_v: $\rho_v^* = \phi_\rho \rho_v$.

The same eigenvalue logic applies in this multi-signal setting. The dominant eigenvalue of $(I - \bar{K}C)A$ depends on the joint noise-persistence structure: when both noise processes have similar persistence, signals are similarly confounding and the filter is slow to extract the state; when persistence is asymmetric (one ρ much smaller than the other), the less-persistent signal carries more information about the state and the filter extracts it more quickly. As in the one-signal environment, rational and DE behave similarly in the medium term; θ primarily affects the response in the first few periods rather than medium-term dynamics. In the asymmetric case, both rational and DE therefore generate less persistence.

In the model with persistent expectations bias the persistence comes from cumulative misattribution rather than from the filter's eigenvalues, and the misattribution is signal-specific—present for the misperceived process and absent for the correctly perceived one. The mechanism therefore continues to deliver hump-shaped, persistent IRFs even when the other signal's noise is transitory. In our model, the signal structure we estimate features exactly this asymmetry: the inflation signal carries transitory noise while the coordinating and idiosyncratic signals are more persistent. Misperception of signal-noise persistence can therefore generate the medium-horizon forecast-error persistence documented in Section 5 and Table 1. Unlike rational and DE, the model with persistent expectations bias delivers persistent IRFs in both symmetric and asymmetric configurations.

Figure 15 reports the responses to one-standard-deviation shocks to ε_1, ε_2, and u in two configurations: a *symmetric* case ($\rho_v = \rho_w = 0.9$, top row) and an *asymmetric* case ($\rho_v = 0.9$, $\rho_w = 0.5$, bottom row). Under a shock to ε_1, rational and DE IRFs are persistent in the symmetric case but die out quickly in the asymmetric case. The model with persistent expectations bias, by contrast, generates IRFs to ε_1 shocks that are very persistent and remain away from zero at $t = 20$ in both configurations.

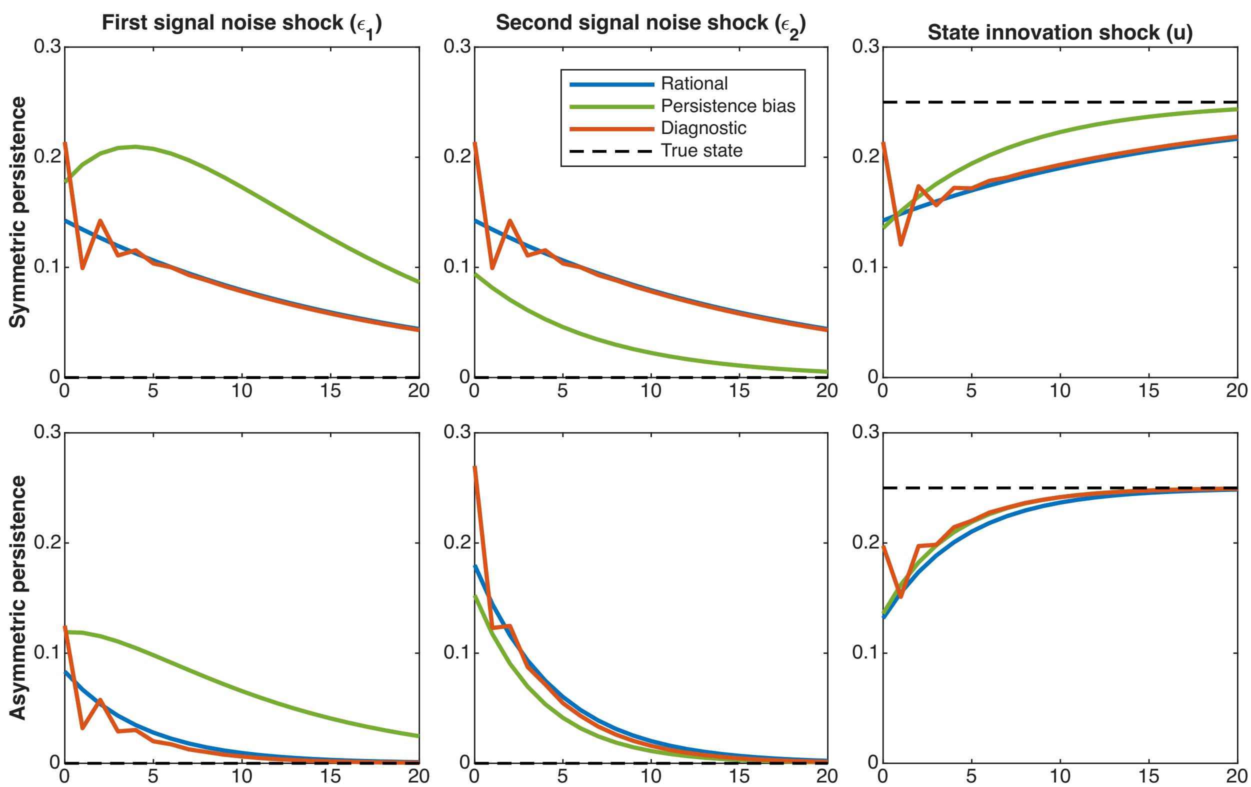

Figure 15: **Persistent expectations bias and diagnostic expectations: two signals.** Impulse response functions of beliefs about x_t to one-standard-deviation shocks to the first signal noise ε_1 (left), the second signal noise ε_2 (middle), and the state innovation u (right). Top row: symmetric noise persistence ($\rho_v = \rho_w = 0.9$). Bottom row: asymmetric noise persistence ($\rho_v = 0.9$, $\rho_w = 0.5$). Persistence misperception is applied only to the first signal ($\rho_v^* = \phi_\rho \rho_v$ with $\phi_\rho = 0.75$, ρ_w correctly perceived). Parameters: $\sigma_u = 0.25$, $\sigma_{\varepsilon_1} = \sigma_{\varepsilon_2} = 0.5$. The diagnostic parameter is $\theta = 0.5$. The black dashed line shows the true response of x_t. Each filter is initialized at its own steady-state forecast-error covariance. Shocks hit in period 0.

I 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 17 bps to be consistent with CPI inflation being on average 17 bps higher than PCE inflation. Finally, we apply a 5-quarter centered 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 2.17 percent accordingly. Note that this 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 2022Q4-2027Q4 by applying the Kalman smoother to the state-space model as defined in equations (15)-(16) 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. As a start, we assume that in the long run, i.e. after 2026Q4, inflation is equal to average expectations. This will be relaxed in step (iv).
- Trend inflation: We set trend inflation equal to the initial guess from step (i)
- I.i.d 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.

(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).[46] 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.

(iv) **Taking into account the contribution of persistent expectations bias**: Based on the previous steps we can obtain estimates of all innovations and derive the path of average expectations that would materialize assuming zero innovations from 2022Q4 forwards. This path reflects the contribution of the persistent expectations bias. Then, we assume that in the long-run, i.e. after 2026Q4, inflation is equal to the level of this path from 2027Q4. We repeat steps (ii) and (iii) under this assumption to obtain the anchoring-compatible inflation path.

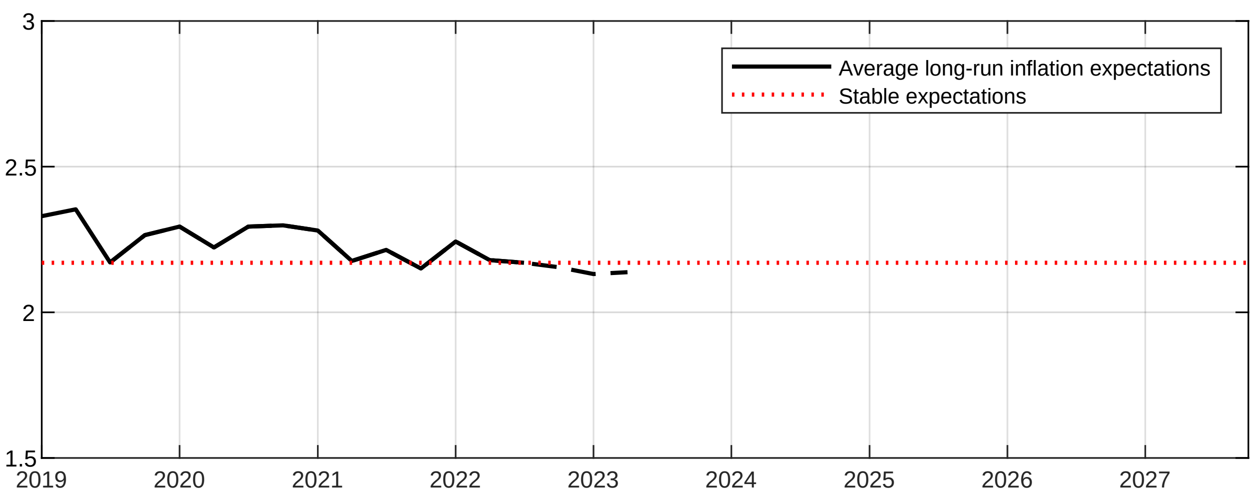

Figure 16: **Average long-run inflation expectations.** The black solid line shows average long-run inflation expectations until 2022Q4 when we start the exercise. The dashed black line shows the realizations of average SPF long-run inflation expectations observed since 2022Q4. The red dotted line corresponds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

[46]For 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.