Disagreement About the Term Structure of Inflation Expectations*

Hie Joo Ahn† Federal Reserve Board

Leland E. Farmer ‡ University of Virginia

June 25, 2025

Abstract

This paper develops a framework for analyzing individual forecasters' inflation expectations across different horizons. We show that inflation expectations are well captured by two factors, level and slope, and decompose expectations into contributions from long-term beliefs, public information, and private information. Our model uniquely captures heterogeneous responses to public information and how these responses amplify disagreement. Estimating the model with Survey of Professional Forecasters data, we find that in normal times, long-horizon disagreement is primarily driven by long-term beliefs, while short-horizon disagreement arises from private information. However, during economic downturns or periods of high uncertainty, heterogeneous responses to public information become the primary driver of disagreement across all horizons. The Fed's response to news reduces disagreement over public information, but not over private information or long-term beliefs. When disagreement about public information is high, monetary policy responses are delayed and a price puzzle emerges, underscoring the importance of clear monetary policy communication for anchoring inflation expectations. During the COVID-19 pandemic, disagreement over long-run inflation expectations based on public information was smaller and shorter-lived than during the Great Recession, indicating better-anchored long-run expectations.

JEL classification: E17, E31, E37, E52, E58, E65.

Keywords: Inflation Expectations, Term Structure, Disagreement, Monetary Policy

*We thank Hassan Afrouzi, Borağan Aruoba, Michael Bauer, Jens Christensen, Thomas Drechsel, Roger E. A. Farmer, Andrew Figura, Jim Hamilton, Edward Herbst, Elmar Mertens, Emi Nakamura, Jeremy Rudd, Jón Steinsson, Luminita Stevens, Eric Swanson, Allan Timmermann, Fabian Winkler, and Christian Wolf for helpful comments, and Travis Berge for the support of resources.

Disclaimer: The views expressed in this paper are those of the authors and do not necessarily reflect the views and policies of the Board of Governors or the Federal Reserve System.

†Federal Reserve Board of Governors, 20th Street and Constitution Avenue NW, Washington, DC 20551, U.S.A. Email: hiejoo.ahn@frb.gov

‡University of Virginia, Email: lef2u@virginia.edu

1 Introduction

Stable inflation expectations are essential for achieving the Federal Open Market Committee's (FOMC) dual mandate of price stability and full employment. Anchored expectations play a crucial role in measuring the slope of the Phillips curve (e.g., McLeay and Tenreyro, 2020; Hazell et al., 2022; Fitzgerald et al., 2024) and in shaping monetary policy design and effectiveness (e.g., Mertens and Williams, 2019; Clarida, 2021). While economists and policymakers traditionally focus on consensus forecasts to assess expectation anchoring, we argue that disagreement in inflation expectations provides an important metric. When expectations are well anchored, forecasters should largely agree on future inflation, resulting in low and stable disagreement. During the Great Recession and COVID-19 pandemic, consensus long-run inflation expectations remained low and stable, but disagreement fluctuated markedly, reaching levels comparable to or exceeding those of the early 1990s, a period of high and volatile inflation. Given that professional forecasters are typically better informed about macroeconomic news than other agents, varying disagreement levels may serve as an independent gauge of expectation anchoring, complementing consensus forecasts.

To assess the degree to which expectations are anchored and the role of monetary policy in anchoring them, we jointly consider disagreement across forecasting horizons and the sources of that disagreement. The source matters: disagreement stemming from public information about monetary policy can be reduced through policy communication and action, while disagreement stemming from private information or individual long-run beliefs is less influenced by policy. To formally analyze disagreement, we introduce a novel empirical model of the individual term structure of inflation expectations, allowing it to vary across forecasters and over time. We then decompose changes in disagreement across horizons into contributions from individual long-term beliefs, private information, and public information at each point in time. We empirically demonstrate that the public information component is closely tied to monetary policy, particularly the Fed's response to news, and that disagreement about public information significantly reduces monetary policy effectiveness and generates a price puzzle. Our approach sheds light on the public information channel through which monetary policy anchors expectations, identifies periods of weak anchoring, and highlights the importance of anchored expectations for the effectiveness of monetary policy.

Specifically, we extend the Nelson and Siegel (1987) term structure model to individual

inflation expectations in the Survey of Professional Forecasters (SPF). We model each forecaster's expectations using two factors: the level, representing long-run expectations, and the slope, capturing the difference between the nowcast and long-end forecasts. This individual-level term structure model recovers the full trajectory of inflation forecasts across all horizons using only two factors, despite incomplete data on forecasting horizons.

These two factors provide a comprehensive basis for assessing anchored inflation expectations. Previous studies have examined either long-run inflation expectation levels (e.g., Fisher et al., 2022) or the resilience of long-run expectations to inflationary shocks (e.g., Bundick and Smith, 2020; Coibion and Gorodnichenko, 2025a). The former approach, level anchoring, is silent on how recent inflationary shocks affect near- and medium-term expectations, which is important for designing and implementing monetary policy. The latter, path anchoring, does not address where long-term expectations ultimately settle. Our model jointly captures both dimensions and assesses anchoring along each, using both consensus expectations and disagreement.

We find that the cross-sectional mean of the level factor closely tracks consensus long-run inflation expectations and the inflation trend from Stock and Watson (2016). The level factor declined throughout the 1990s and remained low and relatively stable from 2000 onward, suggesting well-functioning level anchoring. The mean slope factor is stationary, effectively offsetting transitory inflation fluctuations and preventing their pass-through to long-run expectations. This reflects well-functioning path anchoring: regardless of short-term inflation movements, the estimated consensus slope adjusts to maintain long-term stability. During the Great Recession, both mean factors increased in the latter half of the downturn, reflecting modest de-anchoring of long-run expectations and heightened near-term deflation concerns. However, both factors declined afterward, indicating resilient inflation expectations. During COVID-19, mean level factor fluctuations were muted relative to the Great Recession, whereas the mean slope factor exhibited a pronounced swing, capturing the pandemic's unprecedented inflation volatility. Despite large recessionary shocks, consensus estimates suggest that both level- and path-anchoring have remained effective since 2000.

However, disagreement offers a more nuanced perspective. Disagreement about both factors increased substantially during the Great Recession, its aftermath, and the COVID-19 pandemic. In particular, disagreement about the two factors during the Great Recession and its aftermath, and disagreement about the slope during the pandemic, exceeded early 1990s levels when

inflation was considered nonstationary. Elevated disagreement about the level factor persisted for several years after the Great Recession, reflecting divergent views on post-recession inflation, the consequences of quantitative easing, and uncertainty about the slope of the Phillips curve (e.g., Forbes, 2008; Ball and Mazumder, 2011).[1] During COVID-19, long-run disagreement increased only moderately, while disagreement about the slope spiked during the run-up in inflation, exhibiting the sample's largest dispersion. This suggests professional forecasters' long-run expectations were better anchored during the pandemic than during the Great Recession, while near- to medium-term views diverged sharply over how much large inflationary shocks would affect inflation.[2] The stark contrast between consensus forecasts and disagreement illustrates how disagreement serves as an independent anchoring indicator.

We further identify the information sources of disagreement, which have direct implications for the effectiveness of monetary policy communication. Forecasters typically disagree because of differences in long-term beliefs and varying interpretations of public and private information (e.g., Coibion and Gorodnichenko, 2012; Farmer et al., 2021). If forecasters rely solely on long-term beliefs and private information, monetary policy, which is a form of public information, would have limited anchoring effects. Thus, identifying the portion of forecasts attributable to each source is critical for assessing how effectively monetary policy can anchor expectations.

We estimate a Bayesian dynamic factor model to capture each individual's level and slope factors, which has an unbalanced panel structure with missing observations inherited from the SPF. We decompose each forecaster's term structure into components attributable to individual fixed effects, common factors, and idiosyncratic components, and interpret them economically as individual long-term beliefs, public information, and private information. This allows us to break down the entire distribution of individual forecasts across horizons by information source at each point in time. Our noisy information model offers a realistic characterization of professional forecasters' behavior, building on the idea that inflation forecasts are based on projections of trend and transitory components (e.g., Stock and Watson, 2016) shaped by individual long-term beliefs, private information, and public news. Conceptually, forecasts

[1] Some held the view that inflation would run away because of the amount of liquidity provided monetary policy, while others expected the U.S. to experience sustained low inflation similar to Japan (Forbes, 2008). Relatedly, some argued that the Phillips curve was flat, while others argued that the Phillips curve was still steep with the correct empirical model (Ball and Mazumder, 2011).

[2] The improved anchoring of long-run inflation expectations may reflect the announcement of the 2 percent numerical inflation target in 2012. Bundick and Smith (2020) also provide evidence that the target likely contributed to anchoring expectations after its adoption.

of trend and transitory inflation components correspond to the level and slope factors in our term structure model based on the Beveridge and Nelson (1981) decomposition, while the three information sources align with the dynamic factor structure underlying expectation formation. This establishes a direct analogy between the empirical and structural models, providing novel insights into the forecast dispersion of professional forecasters.

A key feature of our model is that forecasters have heterogeneous responses to public news, which are captured by forecaster-specific loadings on the common component. The dispersion of individual loadings causes public news to become a significant source of time-varying disagreement, departing from previous dispersed information studies (e.g., Herbst and Winkler, 2021).[3] Our noisy information model shows that heterogeneity in responses to public information stems from forecaster-specific Kalman gains, reflecting differences in forecasters' models of inflation's trend and transitory components. By structurally and empirically decomposing each forecaster's expectations into contributions from three information sources, our model reveals how public information, including monetary policy and key data releases, shapes the evolving distribution of the term structure of inflation expectations. These insights offer valuable guidance for designing effective communication strategies and forward guidance.

We find distinct roles for each information source in disagreement across forecasting horizons. Long-term beliefs drive long-run disagreement, and private information drives short-run disagreement, both accounting for the bulk of disagreement in normal times. However, public information emerges as a key driver of heightened disagreement during recessions and elevated inflation uncertainty, primarily because of heterogeneous reactions to public news. This suggests that increased disagreement, combined with the prominent role of public information, signals an elevated risk of weakened expectation anchoring during turbulent times. The greater dispersion from public information could reflect both divergent interpretations of public news and increased attention to it.

This implies that effective monetary policy action and communication are critical for stabilizing inflation expectations during periods of heightened disagreement. We show that the Fed's response to news (Bauer and Swanson, 2023) reduces disagreement attributable to public information, but not that attributable to private information or long-term beliefs. This high-

[3] A few recent studies adopt similar approaches. Kim (2023) allows forecasters to respond differently to fundamental economic shocks by using distinct forecasting models. Similarly, Fisher et al. (2022) permit heterogeneous forecast revisions in response to a common signal, analogous to our public news, although they assume a shared model of inflation dynamics among forecasters.

lights monetary policy as an important component of public information that professional forecasters closely monitor.[4] Because of the increased role of public information in heightened disagreement, timely and effective monetary policy communication can help anchor expectations through the public information channel, particularly during recessions and periods of volatile inflation.

These findings illustrate a direct link between the effectiveness of monetary policy and the information sources of disagreement about future inflation. When public information is the primary source of disagreement, the economy's response to monetary policy shocks is delayed, and a price puzzle emerges. Conversely, when non-public information, such as private information and long-term beliefs, drive disagreement, monetary policy has rapid and statistically significant stabilizing effects. This result remains robust when including individual-level uncertainty and monetary policy uncertainty among other controls, suggesting that disagreement reflects more than just forecasters' perceived uncertainty. This indicates that firmly anchored inflation expectations are a prerequisite for effective monetary policy transmission. Our analysis has important implications for the literature on the effectiveness of monetary under disagreement (e.g., Andrade et al., 2019; Falck et al., 2021; Dong et al., 2024). While prior studies focus solely on heterogeneous beliefs or disagreement itself, we highlight the importance of identifying the information sources of disagreement when assessing expectation anchoring and its role in monetary policy transmission.

This paper contributes to the expectation formation literature in several ways. First, we are the first to extend Nelson and Siegel (1987)'s model to individual-level data by jointly estimating individual parameters using a Bayesian approach. We show that individual term-structure factors map directly to forecasts of inflation's trend and transitory components in a noisy information model, with direct implications for both level- and path-anchoring. Second, our Bayesian dynamic factor model decomposes distributional changes in forecasters' expectations across all horizons into components from three distinct information sources. Public information serves as the key channel through which monetary policy communication shapes and anchors expectations. The factor structure of the noisy information model reveals that heterogeneous responses to public information reflect differences in forecaster-specific Kalman gains and their interactions with fundamental shocks and information noise. We empirically evaluate the relative contributions of individual long-term beliefs, public information, and private information within

[4] Relatedly, Andrade et al. (2019) show that forward guidance reduces the degree of heterogeneous beliefs.

a unified framework, whereas previous studies typically examine these sources in isolation. Our analysis reveals that disagreement among heterogeneous responses to public information plays a central role in anchoring expectations. Finally, we show that public-information-driven disagreement about future inflation is a critical factor in the transmission of monetary policy and the emergence of the price puzzle.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature. Section 3 discusses the Survey of Professional Forecasters data. Section 4 introduces the individual-level parametric model for the term structure of inflation expectations. Section 5 presents the estimation results. Section 6 introduces a noisy information model that decomposes each forecaster's inflation projection into long-term belief, private information, and public information contributions. Section 7 shows how to decompose disagreement about inflation expectations into contributions from the three information sources. Section 8 analyzes how disagreement attributable to public and non-public information sources affect monetary policy effectiveness. Section 9 concludes.

2 Literature Review

Our paper contributes to multiple strands of literature concerning expectation formation and anchored inflation expectations for monetary policy. Recent studies focus on modeling individuals' inflation expectations (e.g.,Herbst and Winkler, 2021; Fisher et al., 2022), but none jointly model both short-term and long-term inflation expectations at the individual level with individual factors that correspond to the trend and transitory variations in inflation. Herbst and Winkler (2021) and Kim (2023) concentrate solely on short-term inflation expectations, while Fisher et al. (2022) focus on long-term inflation expectations, modeling forecasters' expectations using a trend-cycle decomposition based on SPF data. Meanwhile, research on the term structure of inflation expectations has largely centered on consensus paths (e.g., Aruoba, 2020), overlooking individual-level trajectories and variations across forecasting horizons. Crump et al. (2023) estimate a multivariate trend-cycle model, also considering the term structure of expectations based on consensus expectations from various surveys, rather than individual-level responses.

We use a Nelson-Siegel model to parsimoniously characterize the individual-level term structure of inflation expectations. Methodologically, our paper is closest to Herbst and Winkler (2021) and Fisher et al. (2022), as these studies statistically characterize individual-level expectations

and provide economic interpretations based on dispersed information models. In contrast to these studies, our model characterizes the complete term structure of inflation expectations by fitting the response of inflation expectations across all forecasting horizons. We also allow for heterogeneous responses to public information, unlike Herbst and Winkler (2021).[5] This enables us to comprehensively analyze the anchoring process of inflation expectations by studying not only long-run expectations but also the transition to the long run based on both the consensus and disagreement, as well as the roles of public and private information at each forecasting horizon.

Additionally, our paper contributes to the literature on noisy information and learning by providing a coherent empirical framework that disentangles the contributions of long-run beliefs, private information, and public information to distributional changes in inflation expectations across forecasting horizons. Previous studies have examined the roles of public and private information, long-run beliefs, or prior beliefs as sources of expectation formation (e.g., Patton and Timmermann, 2010; Andrade et al., 2016; Lahiri and Sheng, 2008; Maćkowiak et al., 2023; Herbst and Winkler, 2021; Farmer et al., 2021; Fofana and Reis, 2024).[6] While earlier studies have focused on one or a subset of these three information sources individually, we develop a unified statistical framework to evaluate their relative importance in shaping disagreement. To the best of our knowledge, this paper is the first to recover the role of long-run beliefs, public information, and private information, in driving disagreement across forecasting horizons using a formal statistical model.

This paper also contributes to the literature on anchored inflation expectations (e.g., Bundick and Smith, 2020; Castillo-Martinez and Reis, 2024; Coibion and Gorodnichenko, 2025b). Reis (2020) develops a parsimonious structural model to characterize discrepancies in inflation expectations between financial markets and households, finding that inflation became increasingly unanchored from 2014 onward. Meanwhile, Bundick and Smith (2020) analyze the impact of monetary policy on inflation expectations, as measured by the Michigan Survey of Consumers

[5] Fisher et al. (2022) allow for heterogeneous responses to public information.

[6] For example, Lahiri and Sheng (2008) develop a theoretical Bayesian learning model in which experts' forecasts are shaped by their long-term beliefs and their interpretations of public information. This model accounts for the evolution of both within-forecaster variability and between-forecaster disagreement in GDP projections over various forecast horizons. Patton and Timmermann (2010) find that the key source of persistent disagreement stems from heterogeneity in priors and show that the differences in opinion move countercyclically. Farmer et al. (2021) document the importance of long-term beliefs in the formation of macroeconomic expectations by professional forecasters. The empirical findings of these studies align with ours. Namely, private information is the main source of short-term disagreement but long-term beliefs are the primary driver of long-run disagreement.

and TIPS break-even rates, finding that inflation expectations became more anchored following the announcement of an explicit inflation target in 2012. Bundick and Smith assess the degree of anchoring by examining the path from near-term expectations to long-run expectations only, focusing on whether short-run deviations revert to the long-run target. Coibion and Gorodnichenko (2025b) measure the degree of households' anchoring based on the association between short-run inflation expectations and long-run inflation expectations, similar in spirit to Bundick and Smith. Fisher et al. (2022) also studies anchored inflation expectations, but their main focus is on long-run expectations. Unlike these studies, we first demonstrate that disagreement can serve as an independent metric for assessing anchored expectations. Second, our model jointly characterizes the transition from short-term to long-term expectations and the level of long-run expectations, thereby capturing both path and level anchoring. Finally, our model uncovers the role of three distinct information sources in shaping both the path and long-run level of expectations, identifying which sources enhance or hinder the anchoring process.

Lastly, this paper relates to recent studies on the interaction between monetary policy and disagreement among economic agents (e.g., Andrade et al., 2016; Glas and Hartmann, 2016; Ehrmann et al., 2019).[7] Ehrmann et al. (2019) show that long-horizon time-contingent and state-contingent forward guidance effectively reduces disagreement, while short-horizon time-contingent forward guidance does not. Falck et al. (2021) reveal that a price puzzle arises in states of high disagreement but disappears in states of low disagreement using state-dependent local projections. Dong et al. (2024) empirically demonstrate that inflation disagreement weakens the power of forward guidance and conventional monetary policy, providing a structural model in which households hold heterogeneous beliefs about the inflation target of the central bank. Our research differs from recent studies by further identifying the sources of disagreement and explicitly showing that disagreement attributable to public information is the component that delays the effects of monetary policy and leads to a price puzzle.

[7] Andrade et al. (2016) provide empirical evidence for heterogeneous beliefs about forward guidance and analyze the effect of monetary policy in a new Keynesian model. Glas and Hartmann (2016) distinguish individual inflation uncertainty and disagreement between forecasters, showing that overall disagreement increases during contractionary policy periods.

3 Data: The Survey of Professional Forecasters

This section discusses inflation expectations data from the Survey of Professional Forecasters.

3.1 Notation

First, we define some notation that we will use throughout the rest of the paper. Denote the price level at time t by P_t (in our case this will refer to the consumer price index). Let $\pi_{s\to t}$ be the continuously compounded inflation rate between time s and time t:

$$\pi_{s\to t} \equiv log(P_t) - log(P_s). \tag{1}$$

Throughout we will work with continuously compounded inflation rates because of their time-additive properties. Define the forecast, made at time t, of the inflation rate between times r and s, as $\pi_{r\to s|t}$. Finally, let $q_{a-1}(t)$ denote the final quarter of the year prior to the year time t is in.

3.2 Definition of Forecasted Quantities

We collect data on CPI inflation forecasts from the SPF, conducted by the Federal Reserve Bank of Philadelphia. The survey is sent out in the first month of each quarter and responses are collected around the middle of the quarter, e.g. mid-February in Q1. Survey participants are asked to forecast the average quarterly level of the CPI (or transformations of this quantity) at various horizons. The SPF CPI Inflation Forecasts can be broken into 4 categories:

- 1-period backcasts to 4-quarter ahead forecasts of annualized quarter-over-quarter CPI inflation :

  $$100 \times \left[\left(\frac{P_{t+h}}{P_{t+h-1}}\right)^4 - 1\right]$$

  for $h = -1,\ldots,4$.

- 1 to 3-year ahead forecasts of Q4 over Q4 CPI inflation:

  $$100 \times \left[\frac{P_{q_{a-1}(t)+4j}}{P_{q_{a-1}(t)+4(j-1)}} - 1\right]$$

  for $j = 1,\ldots,3$.

- Forecasts of average Q4 over Q4 CPI inflation over the next 5 years:

$$100 \times \left[\left(\prod_{j=1}^{5} \frac{P_{q_{a-1}(t)+4j}}{P_{q_{a-1}(t)+4(j-1)}} \right)^{\frac{1}{5}} - 1 \right]$$

- Forecasts of average Q4 over Q4 CPI inflation over the next 10 years:

$$100 \times \left[\left(\prod_{j=1}^{10} \frac{P_{q_{a-1}(t)+4j}}{P_{q_{a-1}(t)+4(j-1)}} \right)^{\frac{1}{10}} - 1 \right]$$

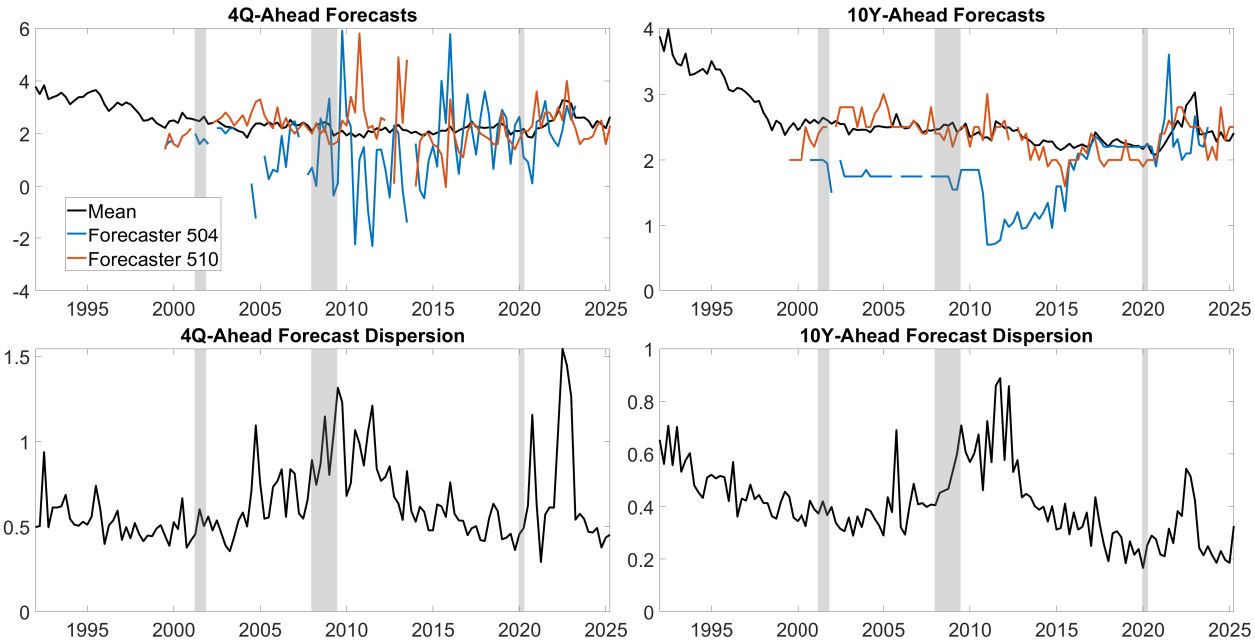
The first type of forecast is what's known as a *fixed horizon* forecast and the other three types are known as *fixed event* forecasts. We assume that these forecasts correspond to forecasts of continuously compounded inflation so that they line up with our model specification directly.[8] That is, we assume

$$\begin{aligned}
100 \times \mathbb{E}_t \left[\left(\frac{P_{t+h}}{P_{t+h-1}} \right)^4 - 1 \right] &\approx 400 \times \pi_{t+h-1 \to t+h|t} \\
100 \times \mathbb{E}_t \left[\frac{P_{q_{a-1}(t)+4j}}{P_{q_{a-1}(t)+4(j-1)}} - 1 \right] &\approx 100 \times \pi_{q_{a-1}(t)+4(j-1) \to q_{a-1}(t)+4j|t} \\
100 \times \mathbb{E}_t \left[\left(\prod_{j=1}^{5} \frac{P_{q_{a-1}(t)+4j}}{P_{q_{a-1}(t)+4(j-1)}} \right)^{\frac{1}{5}} - 1 \right] &\approx 20 \times \pi_{q_{a-1}(t) \to q_{a-1}(t)+19|t} \\
100 \times \mathbb{E}_t \left[\left(\prod_{j=1}^{10} \frac{P_{q_{a-1}(t)+4j}}{P_{q_{a-1}(t)+4(j-1)}} \right)^{\frac{1}{10}} - 1 \right] &\approx 10 \times \pi_{q_{a-1}(t) \to q_{a-1}(t)+39|t}.
\end{aligned}$$

Our sample begins in 1991Q4, which is the first time that forecasts of average Q4 over Q4 inflation over the subsequent ten years become available, and runs through 2025Q1. This is necessary to be able to estimate the long-end of the term structure. Three-year ahead forecasts of Q4 of Q4 inflation and forecasts of average Q4 over Q4 inflation over the subsequent five years first become available in 2005Q3. We restrict our sample to forecasters who report nowcasts to four-quarter ahead forecasts and either a 5-year or 10-year average forecast in at least one quarter to ensure that we can identify long-run forecasts.

[8] This turns out to be a relatively innocuous assumption, see Aruoba (2020) for a discussion.

Figure 1: SPF FORECAST SUMMARY STATISTICS AND EXAMPLES



Notes: The black lines in the upper two panels show the average of inflation forecasts 1 quarter (left panel) and 10 years (right panel) ahead. The colored lines are the forecasts of two forecasters whose IDs are 504 and 510 in the survey. The bottom two panels report the cross-sectional standard deviation of the forecasts.

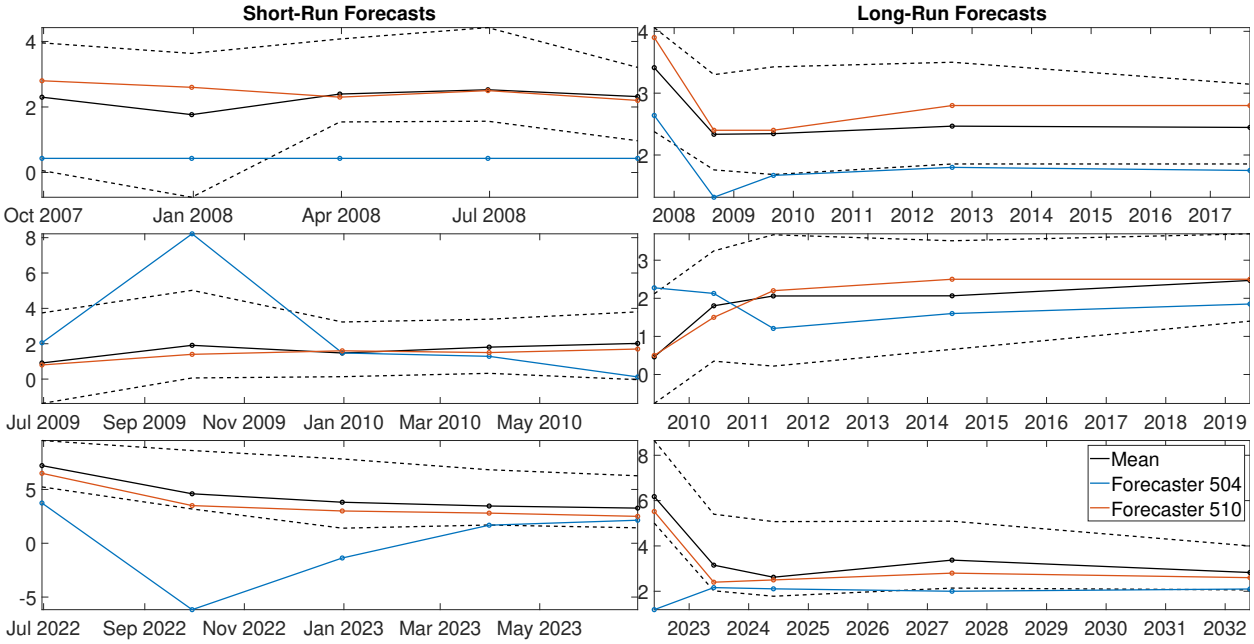
Sources: SPF and authors' calculation.

3.3 Properties of SPF CPI Forecasts

There are 172 unique forecasters in the data set during our sample period. In any given quarter, there are between 28 and 53 forecasters who report a forecast and a median of 37. Forecasters remain in the data set for between 1 and 112 quarters with a median tenure of 14 quarters (3 and a half years).

Figure 1 displays the distribution of 4-quarter ahead forecasts and 10-year ahead forecasts. The upper panels display the mean (which we will refer to as consensus) of the forecasts and the projections of two individual forecasters. The upper-left panel shows 4-quarter ahead forecasts and the upper-right panel reports 10-year ahead forecasts. The gray bars represent NBER recessions. In the first ten to fifteen years of the sample there is a downward trend in both short- and long-run inflation forecasts from around 4% to 2%. Short-run forecasts tend to exhibit more time series volatility than long-run forecasts as they react more strongly to transitory shocks. The bottom panels also show the cross-sectional standard deviation (which we will refer to as dispersion or disagreement) of projections four quarters ahead and 10 years ahead. Short-run forecasts typically exhibit higher dispersion than long-run forecasts, with a notable exception

Figure 2: THREE TERM STRUCTURES OF OBSERVED SPF FORECASTS



Notes: The left panel displays short-horizon forecasts, while the right panel displays long-horizon forecasts. The top panels present a snapshot of SPF forecasts for 2007Q3, the middle panels show the snapshot for 2009Q2, and the bottom panels display the snapshot for 2022Q3. The black lines capture consensus forecasts. The colored lines capture the forecasts of two forecasters whose IDs are 504 and 510 in the survey. The dashed lines are the forecasts at the 5th and 95th percentiles of the distribution.

Sources: SPF and authors' calculation.

being the early 1990s.

In addition to Figure 1 which captures the aggregate properties of the SPF forecasts, we also wish to explore how the expectations of the individual forecasters evolve over time. To this end, Figure 2 plots the reported forecasts over all horizons of two different forecasters for three different time periods: 1) 2007Q3, 2) 2009Q2 and 3) 2022Q3. The solid black line is the consensus, and the dashed lines are the 5th and 95th percentiles of the distribution. The three dates reported in the figure help illustrate the wide variety of shapes the term structure of inflation expectations can take and how different an individual forecaster can be relative to the consensus. The levels of forecasts, as well as the trajectories over the forecasting horizons, are all different.

Following Patton and Timmermann (2010), we only keep forecasters who submit 12 or more forecasts. This allows us to have a higher degree of confidence in our dynamic factor model decomposition. However, information loss from this treatment is not substantial, as it only lowers the number of forecasters in a given quarter by a few people on average.

4 An Individual Term Structure of Inflation Forecasts

In this section, we specify and estimate a model to recover the complete path of inflation forecasts over a 10-year horizon at each point in time, for all forecasters.

4.1 Model

Following Aruoba (2020), we employ a Nelson-Siegel model for the term structure of inflation expectations:

$$\pi_{it\to t+h|t} = L_{it} - \left(\frac{1-e^{-\lambda_i h}}{\lambda_i h}\right) S_{it}, \tag{2}$$

where L_{it} and S_{it} are forecaster-specific level and slope factors, and the λ_i are forecaster-specific shape parameters. Given this representation, the forecast of inflation between any two horizons h_1 and h_2 is given by

$$\pi_{it+h_1\to t+h_2|t} = L_{it} - \left(\frac{e^{-\lambda_i h_1} - e^{-\lambda_i h_2}}{\lambda_i (h_2 - h_1)}\right) S_{it}.$$

Following Diebold et al. (2008), we specify the following decomposition for the factors

$$L_{it} = \alpha_{L,i} + \beta_{L,i} L_t + \varepsilon_{L,it} \tag{3}$$

$$S_{it} = \alpha_{S,i} + \beta_{S,i} S_t + \varepsilon_{S,it} \tag{4}$$

where L_t and , S_t are level and slope factors which are common to all forecasters, $\alpha_{L,i}$, and $\alpha_{S,i}$ are forecaster-specific constant terms which govern the overall level of the individual factors, $\beta_{L,i}$ and $\beta_{S,i}$ are forecaster-specific loadings on the common factors, and $\varepsilon_{L,it}$ and $\varepsilon_{S,it}$ capture the purely idiosyncratic components of the individual factors.[9]

Our goal is to specify a model which can capture rich heterogeneity across forecasters while remaining parsimonious and interpretable. For this reason, we omit the curvature factor. We

[9] In this decomposition, we assume that the loadings are time-invariant. In the dynamic factor model, what is identified is the common component (Dempster et al., 1977), which is the product of the loading and the time-varying factor. This common component captures forecaster i's heterogeneous reactions to common shocks. Therefore, in principle, our model does not preclude time-varying loadings. More importantly, the common component enters our measures of disagreement, meaning its contribution—which we interpret as public information in the following section—incorporates potential time-variations in the loadings. We could allow for slow time variation in the loading parameters, but this would require identifying assumptions about the dynamic process of both the loading and the factor to estimate the time-varying loading and factor explicitly. In this paper, we follow the most standard approach in the dynamic factor model literature—constant loadings and time-varying factors.

find that curvature plays a limited role in the observed individual-level trajectories of inflation expectations.[10] We also assume $\lambda_i = \lambda$. Since λ_i primarily determines the peak of the curvature, this simplification does not have material effects in the estimation given the absence of the curvature factor.

We assume that the common factors follow independent AR(1) processes:

$$\begin{aligned} L_t &= \rho_L L_{t-1} + u_{L,t} \\ S_t &= \rho_S S_{t-1} + u_{S,t}. \end{aligned} \tag{5}$$

Since the scale of the common factors and the factor loadings are not separately identified, we normalize the shocks to the common factors $u_{L,t}$ and $u_{S,t}$ to have unit variance, and we assume the shocks are uncorrelated.[11]

$$\begin{bmatrix} u_{L,t} \\ u_{S,t} \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right) \tag{6}$$

In addition, we assume that the idiosyncratic components evolve according to AR(1) processes which are independent across forecasters.[12]

$$\begin{aligned} \varepsilon_{L,it} &= \rho_{L,i} \varepsilon_{L,it-1} + u_{L,it} \\ \varepsilon_{S,it} &= \rho_{S,i} \varepsilon_{S,it-1} + u_{S,it} \end{aligned} \tag{7}$$

We also assume that the covariance matrix is diagonal, so that the factors evolve independently for each forecaster:

$$\begin{bmatrix} u_{L,it} \\ u_{S,it} \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma^2_{L,i} & 0 \\ 0 & \sigma^2_{S,i} \end{bmatrix} \right). \tag{8}$$

[10] We consider a model specification which includes the curvature factor and AR(3) processes for the factors. Our empirical results remain robust. These results are available upon request. The curvature component is intended to capture nonlinearity across forecasting horizons. However, since the slope component is modeled with an exponential function, it already accounts for some nonlinearity. Including the curvature component neither improves model fit nor materially alters the results.

[11] This treatment is a standard approach in the literature on dynamic factor models (e.g., Diebold et al., 2008).

[12] The dynamics of the common and idiosyncratic components can be generalized to follow VARs with additional lags at minimal computational cost. This makes little difference in the estimated factors and loadings which is why we stick with AR(1) processes in our baseline specification.

For parsimony, we make the further simplifying assumptions that $\rho_{L,i}, \rho_{S,i}, \sigma^2_{L,i}$, and $\sigma^2_{S,i}$ are the same for all forecasters i and equal to $\bar{\rho}_L, \bar{\rho}_S, \sigma^2_L$, and σ^2_S respectively.[13]

Since the model is specified for continuously compounded inflation at the quarterly frequency, mapping the model predictions to observed SPF forecasts is straightforward. Our model can be cast as a linear Gaussian state space model of the form

$$\mathbf{x}_t = \mathbf{F}\mathbf{x}_{t-1} + \mathbf{u}_t, \quad \mathbf{u}_t \sim N(\mathbf{0}, \mathbf{Q}) \tag{9}$$

$$\mathbf{y}_t = \mu_y + \mathbf{H}\mathbf{x}_t + \mathbf{v}_t, \quad \mathbf{v}_t \sim N(\mathbf{0}, \mathbf{R}) \tag{10}$$

where $\mathbf{x}_t$ is a vector of states containing the common and idiosyncratic factors, $\mathbf{F}$ captures the dynamics of the states over time, μ_y is a vector of forecaster-specific fixed effects, $\mathbf{H}$ details the mapping from the states to the observed forecasts, and $\mathbf{Q}$ and $\mathbf{R}$ are the covariance matrices of the innovations to the state equation and the measurement errors respectively. The details of the state-space representation are presented in Appendix A.

4.2 Estimation

Our baseline model has a total of 431 parameters consisting of

- Forecaster-specific means $\{\alpha_{L,i}, \alpha_{S,i}\}_{i=1}^n$
- Forecaster-specific factor loadings $\{\beta_{L,i}, \beta_{S,i}\}_{i=1}^n$
- Factor autocorrelation parameters $\rho_L, \rho_S, \bar{\rho}_L$, and $\bar{\rho}_S$
- Idiosyncratic factor conditional variances σ^2_L and σ^2_S
- Shape parameter λ
- Measurement error variances $\sigma^2_{v,1}, \ldots, \sigma^2_{v,20}$.[14]

[13] See Appendix F.1 for more discussion on the modeling assumptions and their implications. Our simplifying assumptions do not imply that the realized idiosyncratic components are identical; rather, the estimated idiosyncratic components can still differ substantially across individuals. Even with these simplifying assumptions, we have 431 parameters to estimate. While it is in principle possible to allow these parameters to vary across individuals, this more flexible approach would add more than 200 additional parameters, increasing the uncertainty of the estimates. Therefore, we believe our current approach strikes a good balance between flexibility and parsimony.

[14] See Appendix A.2.1 for more details on the measurement equation. In summary, we use one-quarter to four-quarter ahead fixed-horizon forecasts, along with two-year forward, three-year forward, five-year average, and ten-year average fixed-event forecasts. For each quarter of the year, the fixed-event forecasts cover different forecast

The parameter vector is denoted as

$$\boldsymbol{\theta} = [\alpha_1^L, \ldots, \alpha_n^S, \beta_1^L, \ldots, \beta_n^S, \rho_L, \rho_S, \bar{\rho}_L, \bar{\rho}_S, \sigma_L^2, \sigma_S^2, \lambda, \sigma_{v,1}^2, \ldots, \sigma_{v,20}^2]'.$$

Since our model is a linear Gaussian state-space model, we employ the Kalman filter to conduct inference on the latent variables and form the likelihood function. We estimate the model with a Gibbs sampler, detailed in Appendix B.

5 Estimation Results

This section reports and discusses our results from the estimation of the term-structure model.

5.1 Parameter Estimates

Table 1 reports the median, 5th, and 95th percentiles of the posterior distributions for our model parameters. In the case of the forecaster fixed-effects and factor loadings $\alpha_{L,i}, \alpha_{S,i}, \beta_{L,i}$, and $\beta_{S,i}$, we report the median, 5th, and 95th percentiles across posterior draws of the average value across forecasters. The average value of $\alpha_{L,i}$ is consistent with the Fed's 2% inflation anchor, since CPI inflation is known to be slightly higher—by about half a percentage point on average—than core PCE inflation, which the Fed explicitly targets. On average, the term structure of inflation expectations is upward sloping as indicated by a positive value of $\alpha_{S,i}$. Both the common and idiosyncratic factors are estimated to be highly persistent, with autocorrelation coefficients of between 0.72 and 0.94 at the quarterly frequency. Finally, measurement error standard deviations are estimated to be between 10 basis points (for 10-year average expectations in the first quarter of the calendar year) and 67 basis points (for one-quarter ahead expectations).

5.2 Consensus

Figure 3 plots the median of the posterior distribution for the smoothed common factors, $L_{t|T}$ and $S_{t|T}$, along with 95 percent credible intervals. Note that both series are normalized so that their conditional variances are 1. The top panel plots the smoothed common level factor. The estimate closely tracks variations in long-run inflation expectations. The estimate exhibits a

horizons, resulting in four distinct measurements. Consequently, the measurement model comprises twenty equations: four from the fixed-horizon forecasts and sixteen from the fixed-event forecasts.

Table 1: POSTERIOR PARAMETER DISTRIBUTION STATISTICS

Parameter	Median	95% CI	Parameter	Median	95% CI
$\alpha_{L,i}$	2.3583	[2.2497, 2.4629]	$\sigma_{\nu,6}$	0.2558	[0.2421, 0.2706]
$\alpha_{S,i}$	0.1793	[0.0286, 0.3079]	$\sigma_{\nu,7}$	0.2132	[0.1992, 0.2279]
$\beta_{L,i}$	0.1702	[0.1344, 0.2143]	$\sigma_{\nu,8}$	0.0963	[0.0798, 0.1136]
$\beta_{S,i}$	0.4034	[0.3417, 0.4741]	$\sigma_{\nu,9}$	0.3000	[0.2813, 0.3206]
ρ_L	0.9129	[0.8179, 0.9866]	$\sigma_{\nu,10}$	0.3204	[0.3011, 0.3416]
ρ_S	0.8206	[0.6766, 0.9519]	$\sigma_{\nu,11}$	0.2950	[0.2770, 0.3150]
$\bar{\rho}_L$	0.9432	[0.9234, 0.9613]	$\sigma_{\nu,12}$	0.3271	[0.3085, 0.3473]
$\bar{\rho}_S$	0.7214	[0.6856, 0.7563]	$\sigma_{\nu,13}$	0.1811	[0.1686, 0.1945]
σ_L	0.1185	[0.1105, 0.1268]	$\sigma_{\nu,14}$	0.1992	[0.1862, 0.2134]
σ_S	0.4319	[0.4143, 0.4499]	$\sigma_{\nu,15}$	0.2295	[0.2149, 0.2451]
λ	0.1674	[0.1602, 0.1745]	$\sigma_{\nu,16}$	0.2559	[0.2408, 0.2724]
$\sigma_{\nu,1}$	0.6651	[0.6497, 0.6809]	$\sigma_{\nu,17}$	0.1900	[0.1787, 0.2021]
$\sigma_{\nu,2}$	0.4660	[0.4553, 0.4771]	$\sigma_{\nu,18}$	0.1875	[0.1760, 0.1995]
$\sigma_{\nu,3}$	0.4379	[0.4280, 0.4480]	$\sigma_{\nu,19}$	0.1740	[0.1627, 0.1859]
$\sigma_{\nu,4}$	0.4374	[0.4277, 0.4474]	$\sigma_{\nu,20}$	0.2147	[0.2024, 0.2276]
$\sigma_{\nu,5}$	0.2656	[0.2523, 0.2800]			

Notes: The “Median" column reports the posterior median of the corresponding parameter and the “95% CI" column reports the parameter’s 95-percent credible interval. For the parameters $\alpha_{L,i}$, $\alpha_{S,i}$, $\beta_{L,i}$, and $\beta_{S,i}$, the table reports statistics for the average value across forecasters.

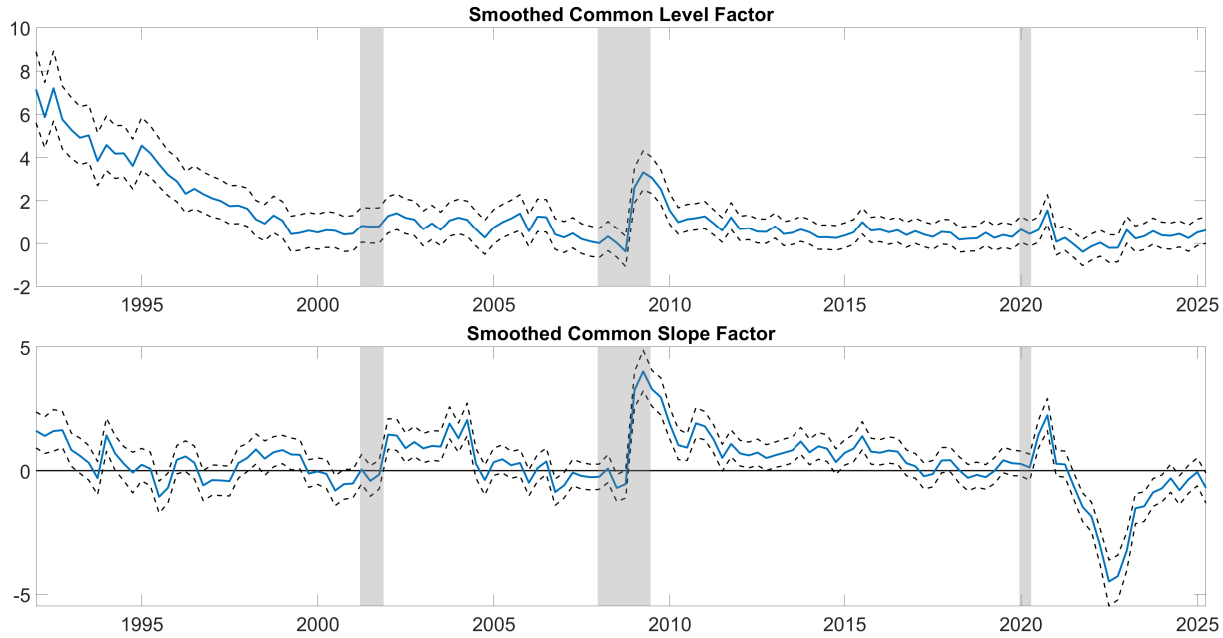
Sources: Authors’ calculation

sharp downward trend in the 1990s and then stays low and relatively stable from 2000 onward. This observation suggests that long-run inflation expectations remain stable and well-anchored, indicating effective level anchoring.

The bottom panel plots the smoothed common slope factor. The slope is typically positive, indicating that the term structure of inflation expectations is upward sloping on average. It is important to note that the slope factor is stationary, largely reflecting transitory fluctuations in realized inflation. The slope factor offsets transitory fluctuations in realized inflation and prevents them from passing through to long-run expectations. Despite short-term inflation fluctuations, the estimated consensus slope adjusts to preserve stability at the long end, indicating effective path anchoring.

While both factors remained relatively stable during the 2001 recession, they exhibited notable swings during the Great Recession and the pandemic recession. During the Great Recession, both the mean level and slope factors rose in the latter half of the downturn, reflecting modest de-anchoring of long-run expectations – partly in response to quantitative easing – and near-term deflation concerns following the sharp decline in inflation at the recession’s

Figure 3: SMOOTHED COMMON FACTORS



Notes: The upper panel plots the posterior median of the smoothed common level factor (blue line) along with 95 percent credible intervals (black dashed line). The bottom panel plots the posterior median of the smoothed common slope factor (blue line) along with the 95 percent credible intervals (black dashed line). The shaded areas denote NBER recessions.

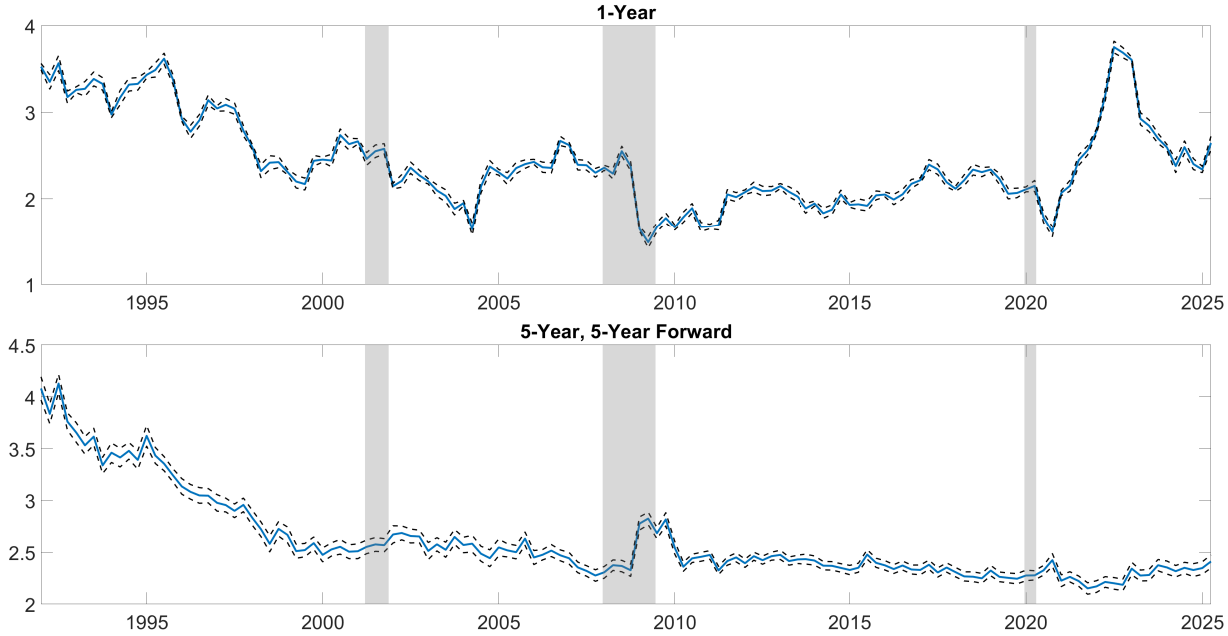
Sources: Authors' calculation

onset.[15] After the recession, however, both factors returned to the prerecession levels, which seems to suggest resilience of inflation expectations across forecasting horizons after a large shock. Notably, the COVID-19 pandemic is interpreted as an episode of better anchored long-run inflation expectations, as the consensus level did not move noticeably after the pandemic recession. Meanwhile, the mean slope factor exhibited a pronounced swing, capturing the sharp decline and subsequent surge in inflation during the pandemic period. According to consensus estimates, the COVID-19 shock is largely transitory, and hence forecasters expect inflation to eventually normalize and return to its pre-pandemic long-run level. Overall, consensus estimates suggest that both level and path anchoring have remained effective, with long-run inflation expectations being resilient despite the large recessionary shocks.

Figure 4 shows the median and 95-percent credible intervals of the forecast distribution mean (across forecasters) at the one-year (top panel) and 5-year, 5-year forward horizons (bottom panel), as implied by the model estimates. We choose these two forecasting horizons, as they are

[15]The modest rise in the level provides one explanation about why the Great Recession was characterized by missing deflation. Despite the drop in price inflation, forecasters expected that the demand-boosting stimulus would eventually raise trend inflation, likely preventing a persistent decline in price inflation.

Figure 4: SMOOTHED CONSENSUS FORECASTS



Notes: This figure plots the posterior median of the average inflation forecast posterior distribution (blue line) recovered from the individual-level term structure model along with 95-percent credible intervals (black dashed line). The two panels correspond to different forecasting horizons, 1-year and 5-year, 5-year forward, respectively. The shaded areas denote NBER recessions.

Sources: Authors' calculation

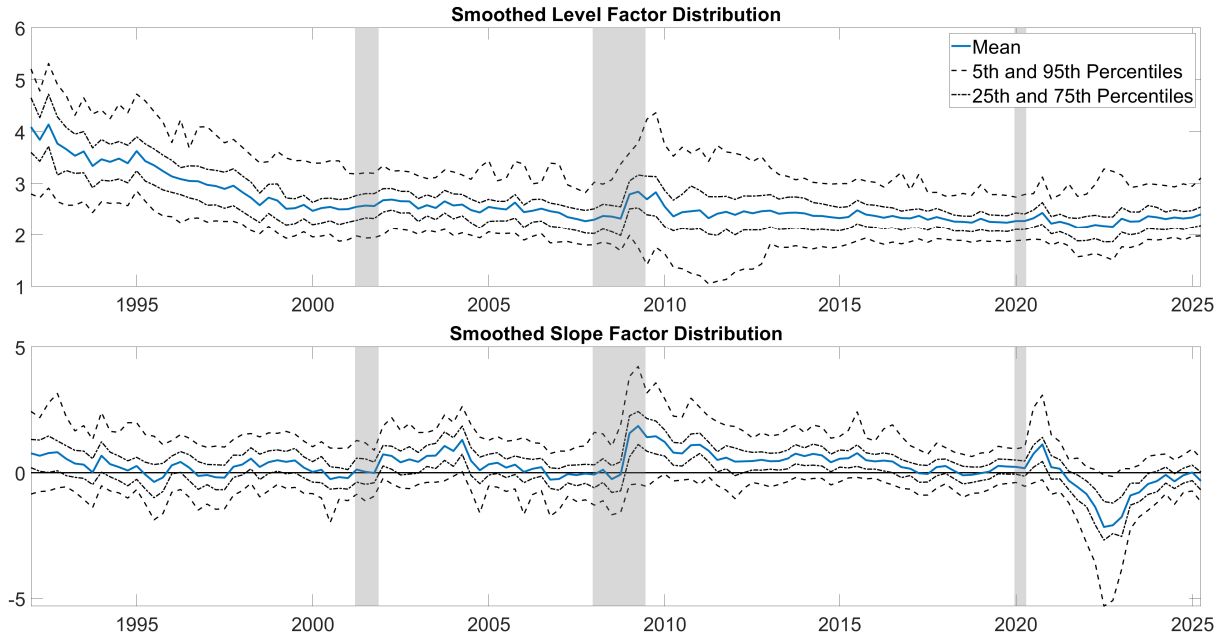
the most commonly cited short-term and long-term forecasts, although our model recovers the forecast at any horizon. Differences in the variation across these forecasts illustrate how shocks to realized inflation propagate and dissipate over different time horizons. Short-term forecasts closely follow the consensus inflation nowcast and, by extension, realized inflation. However, these near-term fluctuations largely fade in forecasts beyond the 5-year horizon, leaving long-run inflation expectations near the inflation trend, as captured by the common level factor. Although 5-year, 5-year forward inflation expectations rose slightly in the second half of the Great Recession, the increase was modest relative to the levels observed in the 1990s.

5.3 Dispersion: A Proxy for Disagreement

This section discusses the distribution of individual term-structure components and forecasts, with a particular emphasis on dispersion, as our primary interest lies in understanding disagreement and how it differs from the consensus.

Figure 5 plots the distributions of forecasters' smoothed level and slope factors. Specifically, we plot posterior medians of the 5th, 25th, 50th, 75th, and 95th percentiles of forecasts across

Figure 5: SMOOTHED FACTOR DISTRIBUTIONS



Notes: The figure shows the cross-sectional distributions of the individual level factors (upper panel) and individual slope factors (bottom panel). The solid blue line is the posterior median of the median factor across forecasters. The dashed-dotted lines depict the posterior medians of the 25th and 75th percentiles. The dashed lines depict the posterior medians of the 5th and 95th percentiles. The shaded areas denote NBER recessions.

Sources: Authors' calculation

forecasters. The factors are in units of annualized percentage points. The top panel displays the distribution of individual smoothed level factors, and the bottom panel displays that of individual smoothed slope factors.

The dispersion offers a view that is somewhat different from what the consensus estimates suggest. Although consensus long-run inflation expectations remain low and stable after 2000, the dispersion of estimates changes substantially over time. In particular, the dispersion seen during and after the Great Recession is larger than that seen in the early 1990s when consensus long-run inflation expectations were around four percent and trending down, and hence long-run inflation expectations were evaluated to be weakly anchored. The increased dispersion of the level factor persisted for several years after the Great Recession. Similarly, the dispersion of slope factor widened to levels comparable to the early 1990s and remained elevated for a few years after the Great Recession. This suggests that forecasters disagreed on how much of the inflation decline during the Great Recession was transitory and where inflation would ultimately stabilize, signaling an unanchoring of long-run expectations.

Indeed, during this period of a weakened inflation anchor, views diverged on where inflation

would settle once the effects of the Great Recession and expansionary monetary policy had subsided. Some expected that inflation would run away to a high level because of aggressive and prolonged quantitative easing, while others predicted that inflation would remain low like in Japan (e.g., Forbes, 2008). Economists and forecasters also disagreed about the slope of the Phillips curve, with some believing that the Phillips correlation was essentially zero, and others believing that the tradeoff between real activity and inflation remained significant.[16] From 2013 onward, the disagreement about the long run narrowed dramatically, reflecting the lift-off of quantitative easing to stabilize inflation.

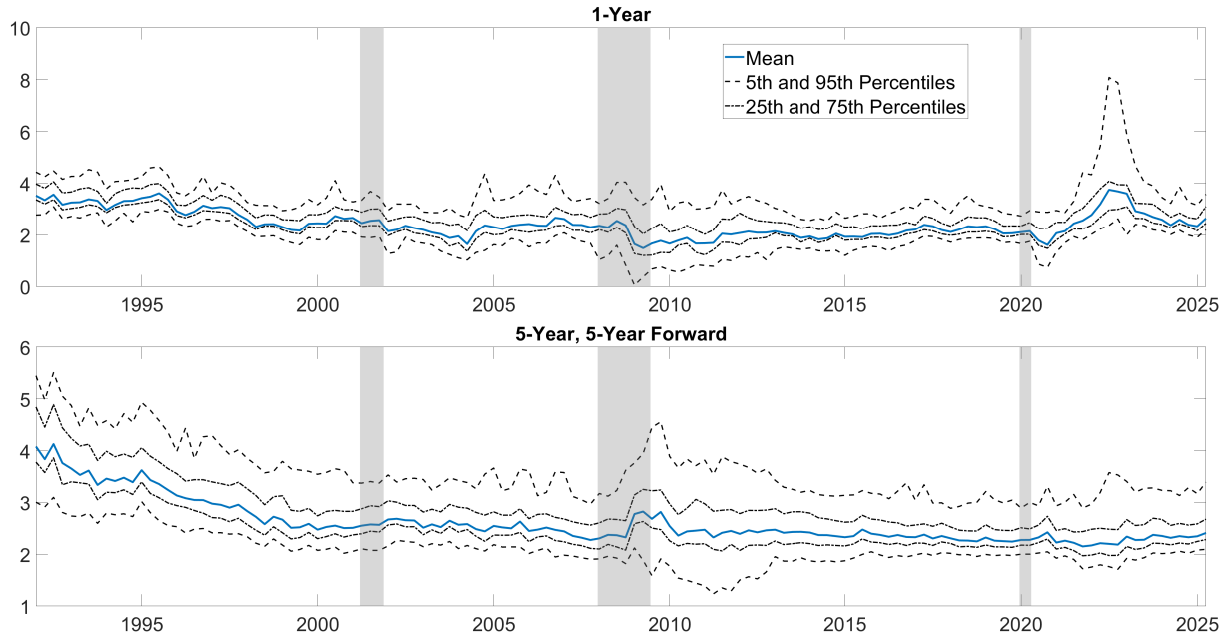
Disagreement during the COVID-19 pandemic followed a somewhat different pattern than in the Great Recession. The dispersion of the level factor rose only modestly, reaching levels comparable to those seen during post-2000 expansions. This suggests that long-run inflation expectations remained more firmly anchored than in the aftermath of the Great Recession. Meanwhile, the dispersion of the slope spiked at the onset of the pandemic and then reached a record-high level during the run-up in inflation. This suggests that forecasters held divergent views on the effects of shocks, such as COVID-19 and the resulting stimulus policies, on inflation in the near and medium term, while their long-run expectations largely remained as aligned as in normal times. Consistent with the consensus evidence, disagreement in the level and slope suggests that long-run inflation expectations were better anchored during the pandemic than during the Great Recession and its aftermath. However, the extent and persistence of disagreement point to somewhat weaker anchoring than what is implied by consensus estimates.[17]

The distribution of individual term-structure factors maps directly into the distribution of inflation forecasts. Figure 6 shows the distribution of inflation forecasts at the 1-year horizon (top panel) and the 5-year, 5-year forward horizon (bottom panel). The dispersion of long-term inflation forecasts corresponds closely to that of the level factor, while the dispersion of short-term forecasts mirrors that of the slope factor. Notably, the distribution of long-term forecasts is positively skewed, exhibiting greater right skewness than observed during the Great Recession. Similarly, short-term inflation forecasts display increased positive skew. This suggests that some forecasters anticipated significantly higher inflation in both the short- and long-term,

[16] This disagreement, reflecting weakly anchored inflation expectations, is also evident in academic research on the puzzling inflation dynamics during the Great Recession and subsequent recovery. See, for example, Ball and Mazumder (2011), Constâncio (2015), and Friedrich (2016).

[17] The improved anchoring of long-run inflation expectations may reflect the Federal Reserve's adoption of a 2 percent numerical inflation target in 2012. Bundick and Smith (2020) also provide evidence that this target helped anchor expectations following its introduction.

Figure 6: DISTRIBUTION OF FORECASTS



Notes: The figure shows the cross-sectional distribution of individual inflation forecasts at two different forecast horizons. The solid blue line is the posterior median of the mean forecast across forecasters. The dotted lines depict the posterior medians of the 25th and 75th percentiles. The dashed lines depict the posterior medians of the 5th and 95th percentiles. The shaded areas denote NBER recessions.

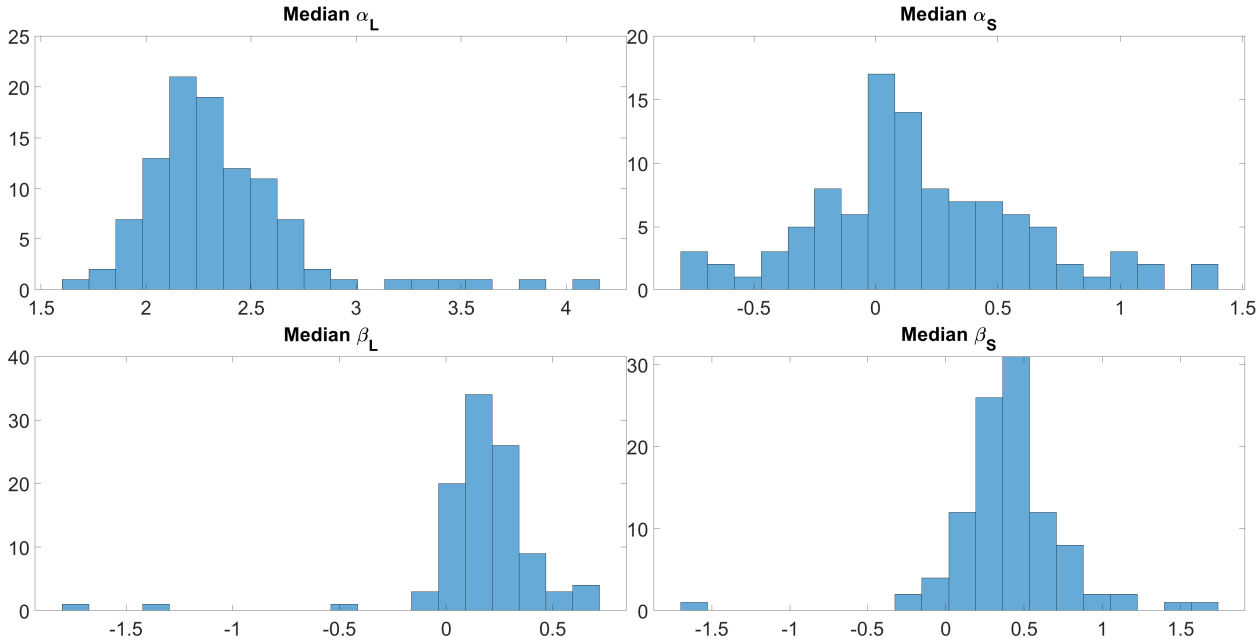
Sources: Authors' calculation

reflecting an elevated perception of upside inflation risk. This upside riskm, despite lower overall disagreement compared to the Great Recession, distinguishes the pandemic period from the Great Recession.

We also examine the standard deviation (disagreement) and skewness of the forecasting distribution in Appendix D. Consistent with our observation in Figure 6, disagreement increased dramatically over the course of the Great Recession and stayed elevated for a few years after the end of the recession. During this time, disagreement is much larger than in the 1990s across forecasting horizons. In addition, skewness also increased during the pandemic and reached its highest levels in the sample period. Both disagreement and skewness declined after 2022 but remained elevated relative to pre-pandemic levels at the end of 2023. This observation suggests that higher moments of the cross-sectional distribution carry information about the degree of anchoring in inflation expectations quite different from what is captured by the consensus.

The stark contrast between consensus forecasts and disagreement suggests that disagreement serves as an independent indicator of the degree of inflation anchoring. In subsequent sections, we examine the drivers of disagreement.

Figure 7: DISTRIBUTIONS OF POSTERIOR MEDIANS: CONSTANTS AND FACTOR LOADINGS



Notes: The figure shows the cross-sectional distribution of posterior median fixed effets and factor loadings for both the level and slope factors.

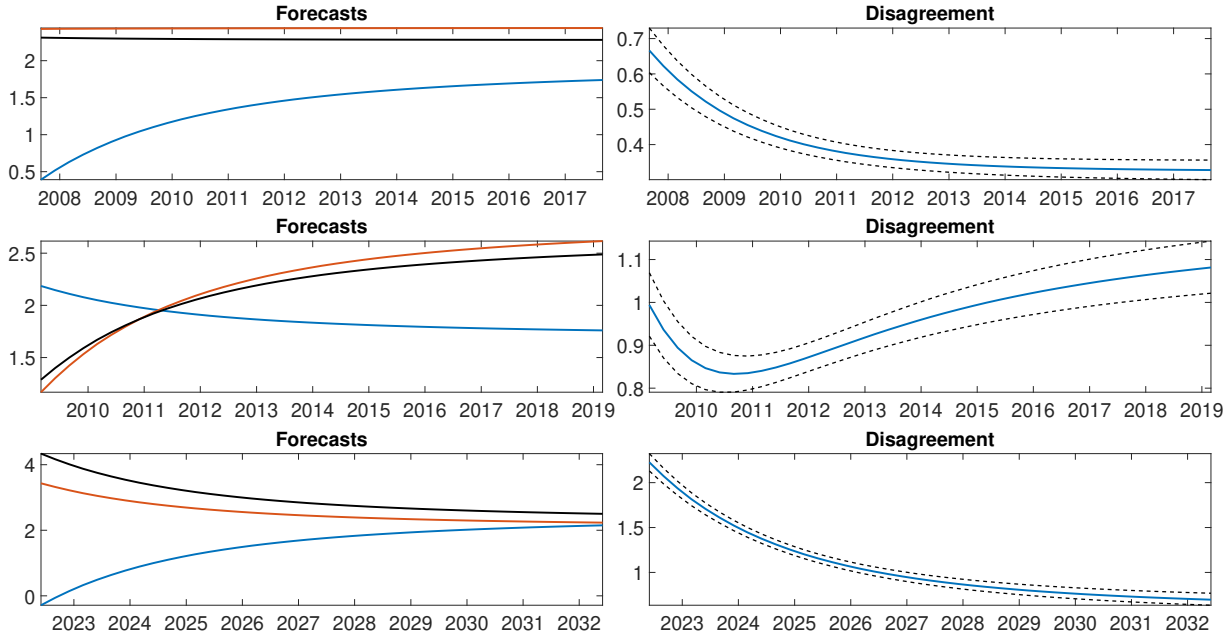
Sources: Authors' calculation

5.4 Forecaster-Specific Parameters and Dynamics

In this section, we examine the sources of forecast dispersion in our statistical model, focusing on individual fixed effects and factor loadings. Figure 7 displays the cross-sectional distribution of $\alpha_{L,i}$ and $\alpha_{S,i}$ (individual fixed effects) and $\beta_{L,i}$ and $\beta_{S,i}$ (loadings on common term structure factors), showing the posterior medians across forecasters. The upper panels reveal substantial disagreement about the long-run means of the level and slope factors. Although most forecasters expect the long-term mean level to be between 2% and 2.5%, a significant share anticipates it to exceed 2.5% (upper left panel). For the slope, the majority of forecasters have a long-run mean close to zero, though the distribution shows a large dispersion ranging from -0.7 to 1.4 (upper right panel).

The bottom panels show the distributions of factor loadings, revealing substantial dispersion and indicating notable heterogeneity in responses to common factors. Specifically, the loadings for the level factor predominantly range between 0 and 0.6, though a small number of forecasters exhibit negative loadings. The slope factor shows even greater dispersion, spanning from 0 to 1.5, with a few forecasters who also have negative loadings. While most forecasters adjust their expectations in the same direction, a small subset reacts in the opposite direction. The

Figure 8: TERM STRUCTURE OF INFLATION EXPECTATIONS AT THREE DATES



Notes: The left panels display estimated posterior medians of the term structure of inflation expectations for the consensus (black line) and two individual forecasters (forecaster 504 in blue and forecaster 510 in red). The right panels display posterior medians of the term structure of disagreement (solid blue line) along with 90% credible intervals (dashed black lines) over the corresponding forecast horizon. The three rows correspond to the dates 2007Q3, 2009Q2, and 2022Q2, respectively.

Sources: Authors' calculation

heterogeneity in fixed effects and factor loadings highlights that even professional forecasters hold markedly different beliefs about long-run inflation expectations and the near-term path toward that long-run level, and they respond quite differently to news related to these two components.

Next, we examine the individual term structures of inflation expectations. Our parsimonious model flexibly captures diverse patterns of disagreement across horizons, often diverging from the consensus. Figure 8 shows the term structure of inflation forecasts of two individual forecasters along with the consensus (left panels) and the disagreement across forecasting horizons (right panels). First, we consider 2007Q3 (upper panels). This was the final survey conducted before the start of the Great Recession. The consensus term structure is relatively flat, with the nowcast and 10-year ahead inflation expectations being around 2.25%. Meanwhile, the two forecasters disagree about the nowcasts and the direction of slopes, as shown in the flat term structure of forecaster 510 (red line) and the upward sloping term structure of forecaster 504 (blue line). Overall, disagreement is largest in the short run at 70 basis points but decreases over the forecasting horizon to about 30 basis points, yielding a downward-sloping term structure of

disagreement (upper right panel).

A very different pattern is observed in 2009Q2 – the peak of the Great Recession. The term structure is sharply upward-sloping for forecaster 510 and the consensus, while it is moderately downward-sloping for forecast 504. The disagreement is high in the short run at 1%, smallest at about the one-year horizon at 85 basis points, and then increases again in the long run to about 1.1% (middle right panel). This non-linearity signals higher uncertainty in the immediate future and in the long run than in the short to medium run.

Finally, we consider 2022Q2 (bottom panels). Inflation expectations are decreasing for forecaster 510 and the consensus, which they are sharply increasing for forecaster 504. The disagreement is largest in the current quarter but reduces over the forecasting horizon, producing a downward-sloping term structure of disagreement (lower right panel). Although the term structure of disagreement is similar to that before the Great Recession, the magnitude of disagreement is three times as large as that seen in 2007 in the short run.

6 A Noisy Information Model of Disagreement

This section introduces a noisy information model with heterogeneous forecasters to obtain insight into the sources of disagreement about expected inflation. Reflecting the forecasting practices of professionals, each forecaster predicts inflation by separately projecting its trend and transitory component. We then map the structural model to our dynamic factor model to provide economic interpretations for the statistical model estimates. Each forecaster's inflation forecast is decomposed into an individual fixed effect, a common component, and an idiosyncratic component. The fixed effect captures long-term beliefs, the common component captures responses to and interpretations of public information and fundamental shocks, and the idiosyncratic component captures private information.

6.1 Derivation of Belief Dynamics

Following previous literature on inflation forecasting often cited by professional forecasters (e.g., Stock and Watson, 2007, 2016), inflation is decomposed into its trend and transitory components

as follows:

$$\pi_t = \tau_t + c_t \tag{11}$$

$$\tau_t = \tau_{t-1} + \sigma_\tau \varepsilon_{\tau,t} \qquad \varepsilon_{\tau,t} \overset{\text{iid}}{\sim} N(0,1) \tag{12}$$

$$c_t = \rho c_{t-1} + \sigma_c \varepsilon_{c,t} \qquad \varepsilon_{c,t} \overset{\text{iid}}{\sim} N(0,1) \tag{13}$$

where τ_t is a stochastic trend component which follows a random walk and c_t is a transitory component which follows an AR(1) process.[18] ρ is the autocorrelation coefficient of c_t, and σ_τ and σ_c are the volatility of the fundamental shocks to the trend and transitory components, respectively. According to the trend-cycle decomposition of Beveridge and Nelson (1981), the trend is defined as the long-end (or infinite-horizon) forecast of a variable, making it analogous to the level factor of the term structure. The gap between inflation and its trend, the transitory component, corresponds to (the negative of) the slope factor.

Define the state vector of the system as $\boldsymbol{x}_t \coloneqq [\tau_t, c_t]'$ and the shock vector as $\boldsymbol{\varepsilon}_t \coloneqq [\varepsilon_{\tau,t}, \varepsilon_{c,t}]'$. The dynamics of the states can be written in vector form as

$$\boldsymbol{x}_t = \boldsymbol{A}\boldsymbol{x}_{t-1} + \boldsymbol{\Sigma}\boldsymbol{\varepsilon}_t. \tag{14}$$

We allow agents to misperceive the true parameters of the DGP so that agent i's perceived DGP is

$$\boldsymbol{x}_t = \boldsymbol{A}_i\boldsymbol{x}_{t-1} + \boldsymbol{\Sigma}_i\boldsymbol{\varepsilon}_t. \tag{15}$$

Agents receive signals, $\boldsymbol{s}_{it}$, about the trend and transitory components separately[19], which are each contaminated with public and private noise shocks:

$$\boldsymbol{s}_{it} = \boldsymbol{x}_t + \boldsymbol{\Omega}_i\boldsymbol{w}_{it} + \boldsymbol{\Theta}\boldsymbol{z}_t \tag{16}$$

[18] Studies like Stock and Watson (2007), Chan et al. (2013), and Mertens (2016) employed a similar approach to model inflation dynamics. This approach has also been adopted by many professional forecasters. See, for example, Federal Reserve Bank of New York (2023).

[19] Assuming that agents receive independent signals about the trend and transitory components implies that their beliefs about the two components evolve independently. While we make this assumption for simplicity, it aligns with standard inflation forecasting models in which the trend and transitory components are modeled separately, with shocks to each assumed to be independent (e.g., Stock and Watson, 2007; Chan et al., 2013). Relatedly, the literature on the Phillips curve (for example, Hazell et al., 2022) links the transitory component of inflation to economic slack, conditional on the inflation trend controlled with time fixed effects. Allowing agents to receive noisy signals about overall inflation directly would generate general VAR dynamics of beliefs about the states, without affecting our main quantitative results.

where

$$\boldsymbol{\Omega}_i \sim \begin{bmatrix} \omega_{\tau,i} & 0 \\ 0 & \omega_{c,i} \end{bmatrix}, \quad \boldsymbol{\Theta} \sim \begin{bmatrix} \theta_\tau & 0 \\ 0 & \theta_c \end{bmatrix}$$

and

$$\boldsymbol{w}_t \coloneqq \begin{bmatrix} w_{\tau,it} \\ w_{c,it} \end{bmatrix} \overset{\text{iid}}{\sim} N(\mathbf{0}, \boldsymbol{I}), \quad \boldsymbol{z}_t \coloneqq \begin{bmatrix} z_{\tau,t} \\ z_{c,t} \end{bmatrix} \overset{\text{iid}}{\sim} N(\mathbf{0}, \boldsymbol{I}).$$

$\boldsymbol{w}_{it}$ is a shock which is idiosyncratic to each forecaster (private information), while $\boldsymbol{z}_t$ is a shock that is common across all forecasters (public information). For simplicity, all forecasters are assumed to agree on the standard deviations of the common shocks, captured by the matrix $\boldsymbol{\Theta}$.[20]

Let $\mathbb{F}_i^*$ denote the expectations operator of a fully rational Bayesian agent with the same parameter beliefs as agent i. If agents were fully rational Bayesians, they would form forecasts of the state vector in the following way:

$$\mathbb{F}_{it}^*[\boldsymbol{x}_{t+1}] = \boldsymbol{A}_i \left\{ \boldsymbol{G}_i \boldsymbol{s}_{it} + (\boldsymbol{I} - \boldsymbol{G}_i) \mathbb{F}_{it-1}^*[\boldsymbol{x}_t] \right\}, \tag{17}$$

where $\boldsymbol{G}_i$ is the steady-state Kalman gain matrix of forecaster i. We will instead work with nowcasts rather than one-step ahead forecasts, which are given by;

$$\mathbb{F}_{it}^*[\boldsymbol{x}_t] = \boldsymbol{G}_i \boldsymbol{s}_{it} + (\boldsymbol{I} - \boldsymbol{G}_i) \boldsymbol{A}_i \mathbb{F}_{it-1}^*[\boldsymbol{x}_{t-1}]. \tag{18}$$

Finally, we allow agents to have long-term beliefs about $\boldsymbol{x}_t$ which lead to permanent deviations from Bayesian updating, denoted by $\boldsymbol{b}_i$. Agent i's expectation operator is denoted by $\mathbb{F}_i$ as opposed to $\mathbb{F}_i^*$ and leads to the belief updating equation:

$$\mathbb{F}_{it}[\boldsymbol{x}_t] = \boldsymbol{b}_i + \boldsymbol{G}_i \boldsymbol{s}_{it} + (\boldsymbol{I} - \boldsymbol{G}_i) \boldsymbol{A}_i \mathbb{F}_{it-1}^*[\boldsymbol{x}_{t-1}]. \tag{19}$$

Next, we work with the Bayesian belief updating equation to obtain a representation of beliefs as a function of the infinite history of realized states, public noise shocks, and private noise shocks. Recursively substituting in the expression for lagged nowcasts we get:

$$\mathbb{F}_{it}^*[\boldsymbol{x}_t] = \sum_{k=0}^{\infty} \left[(\boldsymbol{I} - \boldsymbol{G}_i) \boldsymbol{A}_i \right]^k \boldsymbol{G}_i \boldsymbol{s}_{it-k}$$

[20] This assumption allows for a cleaner understanding of what a public noise shock is, although it is not necessary for our main results.

$$= \sum_{k=0}^{\infty} \left[(\boldsymbol{I} - \boldsymbol{G}_i) \boldsymbol{A}_i \right]^k \boldsymbol{G}_i \boldsymbol{x}_{t-k} + \sum_{k=0}^{\infty} \left[(\boldsymbol{I} - \boldsymbol{G}_i) \boldsymbol{A}_i \right]^k \boldsymbol{G}_i \boldsymbol{\Omega}_i \boldsymbol{w}_{it-k}$$
$$+ \sum_{k=0}^{\infty} \left[(\boldsymbol{I} - \boldsymbol{G}_i) \boldsymbol{A}_i \right]^k \boldsymbol{G}_i \boldsymbol{\Theta} \boldsymbol{z}_{t-k}.$$

Thus, the Bayesian belief can be written as follows:

$$\mathbb{F}^*_{it}[\boldsymbol{x}_t] = \sum_{k=0}^{\infty} \boldsymbol{\lambda}_{i,k} \boldsymbol{\varepsilon}_{t-k} + \boldsymbol{u}_{it} + \boldsymbol{v}_{it}, \tag{20}$$

where

$$\begin{aligned}
\boldsymbol{\lambda}_{i,k} &:= \sum_{j=0}^{k} \left[(\boldsymbol{I} - \boldsymbol{G}_i) \boldsymbol{A}_i \right]^k \boldsymbol{G}_i \boldsymbol{A}^{k-j} \boldsymbol{B} \\
\boldsymbol{v}_{it} &:= \sum_{k=0}^{\infty} \left[(\boldsymbol{I} - \boldsymbol{G}_i) \boldsymbol{A}_i \right]^k \boldsymbol{G}_i \boldsymbol{\Omega}_i \boldsymbol{w}_{it-k} \\
&= \boldsymbol{\Phi}_i \boldsymbol{v}_{it-1} + \boldsymbol{\Psi}_i \boldsymbol{w}_{it} \\
\boldsymbol{\Phi}_i &:= (\boldsymbol{I} - \boldsymbol{G}_i) \boldsymbol{A}_i \\
\boldsymbol{\Psi}_i &:= \boldsymbol{G}_i \boldsymbol{\Omega}_i \\
\boldsymbol{u}_{it} &:= \sum_{k=0}^{\infty} \left[(\boldsymbol{I} - \boldsymbol{G}_i) \boldsymbol{A}_i \right]^k \boldsymbol{G}_i \boldsymbol{\Theta} \boldsymbol{z}_{t-k} \\
&= \boldsymbol{\Phi}_i \boldsymbol{u}_{it-1} + \boldsymbol{\Xi}_i \boldsymbol{z}_t \\
\boldsymbol{\Xi}_i &:= \boldsymbol{G}_i \boldsymbol{\Theta}.
\end{aligned}$$

Accounting for the agent's long-term beliefs, we decomposition of their nowcast into three distinct components using (20):

$$\mathbb{F}_{it}[\boldsymbol{x}_t] = \underbrace{\boldsymbol{b}_i}_{\text{long-term belief}} + \underbrace{\sum_{k=0}^{\infty} \boldsymbol{\lambda}_{i,k} \boldsymbol{\varepsilon}_{t-k} + \boldsymbol{u}_{it}}_{\text{common component}} + \underbrace{\boldsymbol{v}_{it}}_{\text{idiosyncratic component}}. \tag{21}$$

This decomposition is consistent with the statistical decomposition of the common and idiosyncratic components, as identified from the covariance matrix of the demeaned or standardized data (Dempster et al., 1977).

Equation (21) allows us to rewrite the beliefs about the trend and transitory component in a

way that clearly maps into equations (3) and (4):

$$\mathbb{F}_{it}[\tau_t] = b_i^\tau + \mathbb{F}_{it}^C[\tau_t] + \mathbb{F}_{it}^I[\tau_t] \tag{22}$$

$$\mathbb{F}_{it}[c_t] = b_i^c + \mathbb{F}_{it}^C[c_t] + \mathbb{F}_{it}^I[c_t], \tag{23}$$

where b_i^τ and b_i^c are forecaster i's long-term beliefs about inflation's trend and transitory components, $\mathbb{F}_{it}^C[\tau_t]$ and $\mathbb{F}_{it}^C[c_t]$ are the forecaster's belief about the common components of τ_t and c_t, and $\mathbb{F}_{it}^I[\tau_t]$ and $\mathbb{F}_{it}^I[c_t]$ are the forecaster's beliefs about the idiosyncratic components of τ_t and c_t. Note that equation (21) indicates that the common component captures heterogeneous interpretations of fundamental shocks and public information, while the idiosyncratic component reflects private information. In our model, these differences in interpretation arise from the distinct forecasting models used by each forecaster, as shown in equation (15).

6.2 Relationship to the Statistical Model

Next, we establish a mapping between the empirical dynamic factor structure of the individual level and the slope factors of our statistical model (equations (3) and (4)) and the dynamic factor structure of the expectations of individual forecasters about the inflation trend and transitory components derived from the noisy information model (equations (22) and (23)). Proposition 1 states that there is an approximate one-to-one mapping between the parameters of the statistical dynamic factor model and those of the noisy information model.

Proposition 1. *Let g_τ and g_c denote the Kalman gains associated with filtering the trend and transitory components from the noisy information model under the true parameters, and let $\rho_c = \mathbb{E}_i[\rho_i]$. Assume that the Kalman gains of agents are correct on average, so that $\mathbb{E}_i[g_{i,\tau}] = g_\tau$ and $\mathbb{E}_i[g_{i,c}] = g_c$. Further, assume that Kalman gains are not too dispersed across agents so that $Var(g_{i,\tau})$ and $Var(g_{i,c})$ are relatively small. Then, we have the following approximate mapping from the noisy information model to the statistical model*

$$\alpha_i^L = b_i^\tau \qquad \alpha_i^S = b_i^c$$

$$\beta_i^L = g_{i,\tau}\left(\sigma_\tau^2 + \theta_\tau^2\right)^{1/2} \qquad \beta_i^S = g_{i,c}\left(\sigma_c^2 + \theta_c^2\right)^{1/2}$$

$$\bar{\rho}_L = 1 \qquad \bar{\rho}_S = \frac{\sigma_c^2\left[\rho - g_c(\rho - \rho_c)\right] + \theta_c^2(1 - g_c)\rho_c}{\sigma_c^2 + \theta_c^2}$$

$$\rho_{L,i} = 1 - g_{i,\tau} \qquad \rho_{S,i} = (1 - g_{i,c})\rho$$

$$\sigma^2_{L,i} = \omega^2_{i,\tau} \qquad \sigma^2_{S,i} = \omega^2_{i,c}$$
$$\lambda_i = -\log \rho_i.$$

See Appendix C for the proof.[21]

A direct corollary of this proposition is that we can recover the Kalman gains of each forecaster associated with the trend and transitory components by taking the ratio of their estimated loadings to the average loading.

Corollary 1. *Forecaster i's factor loadings for the trend, relative to the corresponding cross-sectional average, correspond to the forecaster's Kalman gain for the trend prediction relative to the corresponding cross-sectional average. The same holds for the transitory component.*

$$\frac{\beta^L_i}{\bar{\beta}^L} = \frac{g_{i,\tau}}{g_\tau} \qquad \frac{\beta^S_i}{\bar{\beta}^S} = \frac{g_{i,c}}{g_c}.$$

This corollary tells us that cross-sectional heterogeneity in Kalman gains accounts for the dispersion in factor loadings that capture heterogeneous responses to public information. Note that heterogeneity in Kalman gain essentially reflects cross-sectional differences in forecasting models.

Next, we demonstrate how heterogeneous interpretations of fundamental shocks and public information create forecast dispersion and drive the time-varying share of the common component in disagreement. To illustrate this point, we rewrite the nowcast of $\boldsymbol{x}_t$ in equation (21) in the following way:

$$\mathbb{F}_{it}[\boldsymbol{x}_t] = \underbrace{\boldsymbol{b}_i}_{\text{long-term belief}} + \underbrace{\boldsymbol{\lambda}_{i,0}\boldsymbol{\varepsilon}_t + \boldsymbol{\Xi}_i \boldsymbol{z}_t + \boldsymbol{\Psi}_i \boldsymbol{w}_{it}}_{\text{contemporaneous contribution}} + \underbrace{\sum_{k=1}^{\infty} \boldsymbol{\lambda}_{i,k}\boldsymbol{\varepsilon}_{t-k} + \boldsymbol{\Phi}_i\left(\boldsymbol{u}_{it-1} + \boldsymbol{v}_{it-1}\right)}_{\text{lagged contribution}}. \tag{24}$$

Equation (24) illustrates the mechanism through which heteroegeneous responses to common news generate a time-varying share of public information in disagreement about future

[21] Note that we allow forecasters' models for the trend and transitory components to differ across individuals, as reflected in $\rho_{L,i}$, $\rho_{S,i}$, $\sigma_{L,i}$, and $\sigma_{S,i}$. However, in our empirical implementation, we assume these parameters are common across individuals to reduce the number of parameters to estimate. This does not imply that individuals' realized forecasts of trend and transitory components are identical or exhibit the same ex-post dynamics. Although estimating over 200 additional parameters is feasible by extending nowcasting techniques developed for large datasets (e.g., Banbura and Modugno (2014)) within our Bayesian framework, we adopt the pooled specification given the uncertainty associated with such a high-dimensional estimation.

inflation.

Corollary 2. *According to equation (24), differences in how agents perceive the persistence of the cycle ρ_i, the conditional variances of fundamental shocks $\sigma^2_{i,c}$ and $\sigma^2_{i,\tau}$, and the variances of private signals $\omega^2_{i,c}$ and $\omega^2_{i,\tau}$ lead to heterogeneous loadings on the public signal via Ξ_i, the extent to which forecaster i takes signal about the fundamental from public information. Consequently, dispersion in Ξ_i across agents leads to time-varying disagreement driven by realizations of the public signals $\boldsymbol{z}_t$.*

7 Empirical Decomposition of Disagreement

Given the decomposition based on the noisy information model in Section 6, we analyze how individual long-term beliefs, heterogeneous responses to public information, and private information, each contribute to individual inflation forecasts. Section 7.1 presents the decomposition of individual forecasts. Section 7.2 provides a new measure for the sensitivity of disagreement to each source of information. Section 7.3 discusses the empirical estimates of the decomposition.

7.1 Information Sources of Disagreement

We proceed in two steps. First, we decompose forecaster i's inflation forecast at each forecasting horizon into the three components outlined above. Second, we decompose the cross-sectional variance into contributions from the three sources at each point in time.

Equations (3) and (4) decompose individual level and slope factors into the three information sources. Here, we give economic interpretations to each component using the level factor as an example.

$$L_{it} = \underbrace{\alpha_{L,i}}_{\text{long-term belief}} + \underbrace{\beta_{L,i} L_t}_{\text{public info.}} + \underbrace{\varepsilon_{L,it}}_{\text{private info.}} \tag{25}$$

The level factor of forecaster i, L_{it}, is decomposed into portions representing *long-term beliefs* ($\alpha_{L,i}$), *public information* ($\beta_{L,i} L_t$), and *private information* ($\varepsilon_{L,it}$). Note that $\alpha_{L,i}$ is estimated with individual fixed effects, and $\beta_{L,i} L_t$ and $\varepsilon_{L,it}$ are the common component and the idiosyncratic component, respectively, in the dynamic factor model for the level factor.

Define the long-term belief component of L_{it} to be $L_i^{ltb} = \alpha_{L,i}$, the common component of L_{it} to be $L_{it}^{pub} = \beta_{L,i} L_t$, and the idiosyncratic component of L_{it} to be $L_{it}^{priv} = \varepsilon_{L,it}$. We then

rewrite L_{it} as

$$L_{it} = L_i^{ltb} + L_{it}^{pub} + L_{it}^{priv}. \tag{26}$$

We give similar economic interpretations to each component of S_{it} in Equation (4).[22] Thus, S_{it} is rewritten as:

$$S_{it} = S_i^{ltb} + S_{it}^{pub} + S_{it}^{priv}. \tag{27}$$

Using the decompositions in Equations (26) and (27), forecaster i's h-quarter-ahead inflation forecast at time t, $\pi_{it\to t+h|t}$, is expressed as the sum of three components representing long-run beliefs ($\pi_{it\to t+h|t}^{ltb}$), public information ($\pi_{it\to t+h|t}^{pub}$), and private information ($\pi_{it\to t+h|t}^{priv}$):

$$\begin{aligned}
\pi_{it\to t+h|t} &= L_{it} - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right) S_{it} \\
&= \underbrace{L_i^{ltb} - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right) S_i^{ltb}}_{\text{long-term belief}} + \underbrace{L_{it}^{pub} - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right) S_{it}^{pub}}_{\text{public information}} + \underbrace{L_{it}^{priv} - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right) S_{it}^{priv}}_{\text{private information}} \\
&= \pi_{it\to t+h|t}^{ltb} + \pi_{it\to t+h|t}^{pub} + \pi_{it\to t+h|t}^{priv}.
\end{aligned} \tag{28}$$

7.2 Disagreement Shares

Using equation (28), we compute the fraction of the overall variance of $\pi_{it\to t+h|t}$ explained by each information source, which we call *the disagreement shares*. Conceptually, our disagreement shares resemble the idea of a *beta* in finance, which gauges the sensitivity of stocks to a common factor. We construct the information share θ for disagreement using an approach proposed in Fujita and Ramey (2009).[23] Fujita and Ramey show that when a variable is expressed as the sum of different sub-components, the sum of the θs, defined as the ratio of the covariance between each component and the total to the variance of the total, is equal to one.[24]

[22] Analogous to the level factor, $\alpha_{S,i}$ is estimated with the individual fixed effects, and $\beta_{S,i}S_t$ and $\varepsilon_{S,it}$ are the common component and the idiosyncratic component, respectively, in the dynamic factor model for the slope factor.

[23] Fujita and Ramey (2009) provide a decomposition for the contribution of inflows to unemployment and that of outflows from unemployment to the variance of the unemployment rate. Note that the θ estimates from Fujita and Ramey (2009) can technically be negative or go above 1.

[24] Alternatively, the variance of $\pi_{it\to t+h|t}$ can be decomposed into the sum of the variances of $\pi_{it\to t+h|t}^{ltb}$, $\pi_{it\to t+h|t}^{pub}$, and $\pi_{it\to t+h|t}^{priv}$, assuming that innovations to the common and idiosyncratic components are empirically uncorrelated in the dynamic factor model. In practice, however, realized shocks to these components may exhibit finite-sample comovement, introducing a wedge between the total variance and the sum of the component variances (this comovement tends to vanish in large samples). To address this, we follow the method proposed by Fujita

Note that the h-period ahead inflation forecast of individual i at time t is composed of the portion accounted for by long-term beliefs, public information, and private information. Restating our decomposition of inflation projections we have that:

$$\pi_{it\to t+h|t} = \pi^{ltb}_{it\to t+h|t} + \pi^{pub}_{it\to t+h|t} + \pi^{priv}_{it\to t+h|t}. \tag{29}$$

Taking the covariance of both sides with $\pi_{it\to t+h|t}$ and dividing through by the variance of $\pi_{it\to t+h|t}$, we obtain the following expression:

$$1 = \theta^{ltb}_{h,t} + \theta^{pub}_{h,t} + \theta^{priv}_{h,t}, \tag{30}$$

where

$$\begin{aligned}
\theta^{ltb}_{h,t} &= \frac{\text{Cov}_i\left(\pi_{it\to t+h|t}, \pi^{ltb}_{it\to t+h|t}\right)}{\text{Var}_i(\pi_{it\to t+h|t})} \\
\theta^{pub}_{h,t} &= \frac{\text{Cov}_i\left(\pi_{it\to t+h|t}, \pi^{pub}_{it\to t+h|t}\right)}{\text{Var}_i(\pi_{it\to t+h|t})} \\
\theta^{priv}_{h,t} &= \frac{\text{Cov}_i\left(\pi_{it\to t+h|t}, \pi^{priv}_{it\to t+h|t}\right)}{\text{Var}_i(\pi_{it\to t+h|t})}.
\end{aligned}$$

We denote the disagreement shares of individual long-term belief, public information, and private information for the h-period-ahead forecast formed at time t by $\theta^{ltb}_{h,t}$, $\theta^{pub}_{h,t}$, and $\theta^{priv}_{h,t}$, respectively.

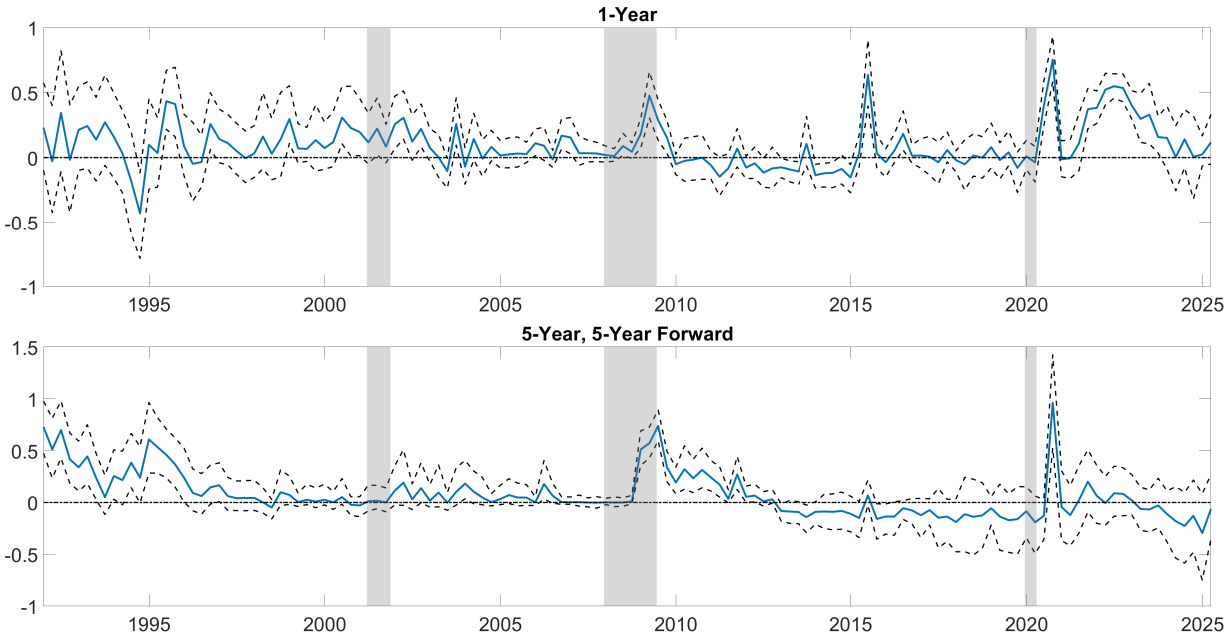
7.3 Decomposition Results

We compute the contributions of the three information sources to disagreement in 1-year-ahead (short-run) and 5-year, 5-year forward (long-run) forecasts.[25] We make a few key observations. First, private information is the primary source of short-run disagreement, explaining approximately 60 percent of the dispersion. Second, individual long-term beliefs are the main source of disagreement for long-run inflation forecasts, accounting for about 60 percent of the dispersion. Finally, public information contributes the smallest share of disagreement across all forecasting horizons, making up around 10 percent for short-run and 7 percent for long-run forecasts on

and Ramey (2009).

[25] Note that our model permits this decomposition at any forecasting horizon.

Figure 9: FORECAST VARIANCE SHARE OF COMMON COMPONENT



Notes: This figure plots the forecast variance share of the common component, $\theta^{pub}_{h,t}$, for two forecasting horizons: 1-year ahead in the top panel, and 5-year, 5-year forward in the bottom panel. The blue line is the posterior median and the dashed black lines are the 5th and 95th percentiles of the posterior.

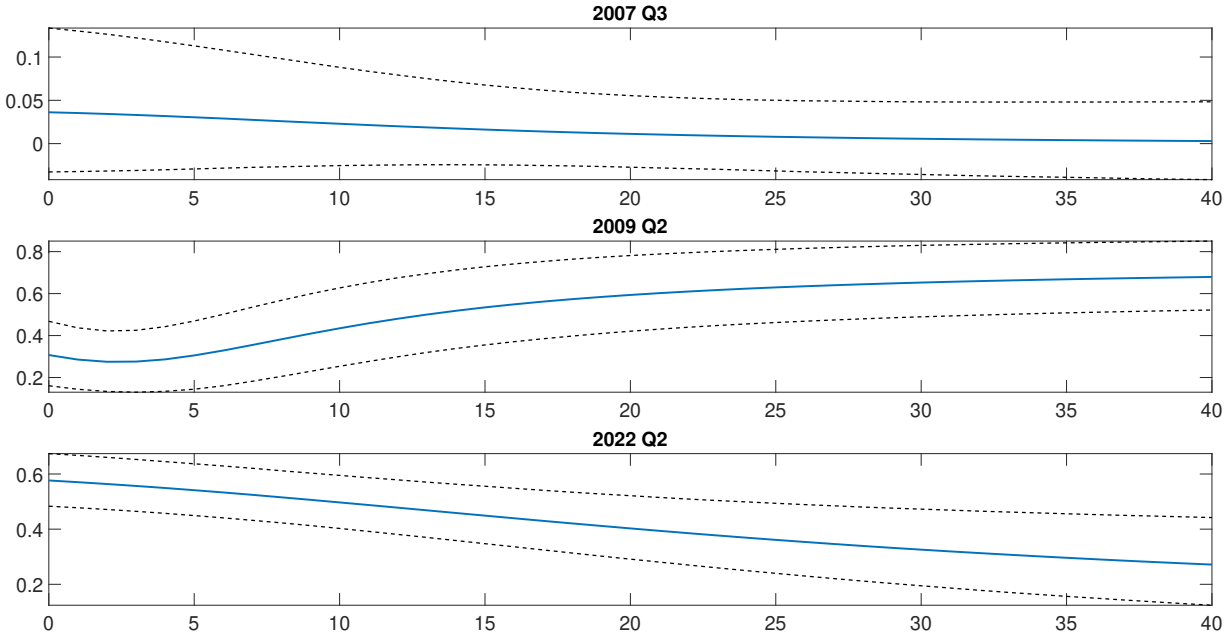
Sources: Authors' calculation

average.

However, as shown in Figure 9, the disagreement share of public information changes dramatically over time. The top panel shows the share of public information in the variance of one-year-ahead inflation forecasts ($\theta^{pub}_{1,t}$), while the bottom panel shows the same for five-year five-year forward forecasts ($\theta^{pub}_{5-10,t}$). Typically it fluctuates at values under 10 percent for both short- and long-horizon forecasts. However, the share exhibits strong countercyclicality, increasing in times of large economic shocks across forecasting horizons. The long-run public information share rises significantly during three major episodes: the early and mid-1990s, the Great Recession, and the COVID-19 pandemic. For the short-term horizon, it also increases in the mid-2010s following the plunge in oil prices, and steps up again in 2021 and 2022 amid the post-pandemic inflation surge. During the pandemic, the share of disagreement attributable to public information spikes to approximately 75 percent for short-horizon forecasts and 96 percent for long-horizon forecasts. This observation suggests that forecasters pay more attention to public information during economic downturns and periods of high inflation volatility, but translate the public information into their forecasts in different ways.

Furthermore, the share of disagreement attributable to public information varies dramatically

Figure 10: TERM STRUCTURE OF VARIANCE SHARES AT THREE DATES



Notes: The figure shows the fraction of overall disagreement about inflation expectations driven by heterogeneous responses to common information over a ten-year forecast horizon at three different points in time as measured by the θ measure proposed in Fujita and Ramey (2009). The top panel reports the estimates as of 2007:Q3, the middle panel as of 2009:Q2, and the bottom panel as of 2022:Q3. The shaded areas denote NBER recessions.

Sources: Authors' calculation

across forecasting horizons over time. Figure 10 shows the share of public information in disagreement across forecasting horizons for three particular periods. In 2007Q3, just before the Great Recession, the role of public information in disagreement is low, near 2%, across all forecasting horizons. In contrast, at the height of the Great Recession in 2009Q2, public information plays a much larger role. The share is larger for long-run forecasts than for short-run forecasts (middle panel), about 70% compared to 30%. This suggests that long-run inflation expectations were less firmly anchored than short-run expectations.

The COVID-19 pandemic is unique in that public information increased forecast disagreement across all horizons, but much more so in the short term (60%) than in the long term (30%). This pattern is the exact opposite of what occurred during the Great Recession. It suggests that long-run inflation expectations were better anchored during the pandemic, while short-term expectations were less so. Faced with unprecedented shocks and policy responses, forecasters paid close attention to public information. However, the extraordinary nature of inflation developments led them to interpret public news very differently when forming short-term forecasts. Despite heightened uncertainty and disagreement about the near-term inflation outlook,

forecasters appear to have had more aligned views about long-run expectations.

The role of public information in shaping long-run disagreement during periods of economic turmoil suggests that timely monetary policy action and communication can help anchor long-horizon expectations. In line with this, Bundick and Smith (2020) show that long-run inflation expectations became more firmly anchored following the FOMC's announcement of a numerical inflation target in 2012. Their findings offer a potential explanation for why forecasters disagreed less overall – and especially less about public information regarding long-run inflation – during the pandemic compared to the Great Recession.

8 Implications for Monetary Policy

In this section, we examine the relationship between the sources of disagreement and the effectiveness of monetary policy, focusing particularly on whether monetary policy communication helps anchor inflation expectations through the public information channel. Section 8.1 analyzes the extent to which the news component of monetary policy reduces disagreement attributable to public information. Section 8.2 explores how disagreement about public information influences the transmission of monetary policy shocks.

8.1 Effect of the Fed's Response to News on Disagreement

So far, we have implicitly assumed that public information includes monetary policy and hence monetary policy can anchor inflation expectations through disagreement about public information. We now examine whether this is the case empirically. For this analysis, we consider a local projection (LP) with an externally identified shock. We use the Fed's response to economic news from Bauer and Swanson (2022), as the externally identified news shock. This news component of monetary policy surprises, reflecting the Fed's interpretation of recent data releases, is measured as the difference between high-frequency monetary policy surprises and the orthogonalized monetary policy shock of Bauer and Swanson (2022).[26] If forecasters pay attention to the Fed's

[26]To illustrate the Fed's response to news, in September 2024, Chair Powell noted that payroll employment data for July and August had likely been overstated and was expected to be revised downward (FOMC press conference, September 18, 2024). This communication about labor market conditions helped reduce market uncertainty regarding the state of the economy and future policy, thereby influencing financial conditions. For example, on September 19, Fox Business reported: 'Fed's Powell: policymakers noted "artificially high" jobs data, revisions in rate cut decision.' On the same day, Morgan Stanley commented: 'U.S. equities were initially higher in the afternoon following the Fed's announcement on September 18, as stocks made an intra-day all-time high, while

responses to data releases, this news component should reduce disagreement about public information among forecasters.

Let y^p_{t+h} denote the disagreement about long-run inflation attributable to public information, measured by the standard deviation of 5-year, 5-year forward inflation forecasts attributable to public information. Similarly, let y^n_{t+h} denote the disagreement attributable to non-public information that includes both private information and individual long-term beliefs, and we measure y^n_{t+h} with the square root of the sum of the variances of 5-year, 5-year forward inflation forecasts attributable to both private information and individual long-term beliefs. We focus on this long-term forecasting horizon to examine the anchoring of long-run inflation expectations.

The local projection model is specified as follows:

$$y^j_{t+h} = \alpha^j_h + \beta^j_h z_t + \Gamma^j_h \mathbf{X}_{t-1} + e^j_{t+h} \qquad \text{for} \quad j \in [p, n] \quad h = 0, 1, \cdots, H, \tag{31}$$

where α^j_h is a constant, z_t is Fed's response to news, β^j_h captures the magnitude of pass-trough of the shock h quarters after impact. For z_t, we consider negative and positive components separately together with the raw estimate. The notation $\mathbf{X}_{t-1}$ refers to a set of macroeconomic control variables, all lagged by one period. The controls include four lags of the two-year Treasury yield, the first differenced log of industrial production (IP), the first differenced log of the consumer price index (CPI),the unemployment rate, the excess bond premium from Gilchrist and Zakrajšek (2012), the level of disagreement attributable to public information and that attributable to non-public information. We include two lags of z_t in the controls. The notation, e^j_{t+h}, is an error term. Note that the parameter draws from the Gibbs sampler are used to construct y^p_{t+h} and y^n_{t+h}. For each draw, we compute the corresponding impulse response. We then obtain posterior distributions of the impulse responses. The sample period is 1991:Q4-2019:Q4.[27]

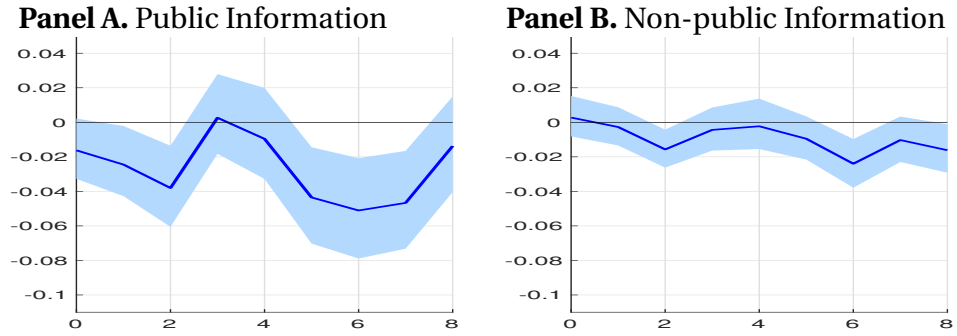
Figure 11 displays the responses of two disagreement measures to the Fed's response to news. Panels A and B present the results using the raw estimates of the news component. Overall, the Fed's response to news significantly reduces disagreement about public information, although its effects on disagreement about non-public information are muted and largely statistically insignificant. This result is mainly driven by negative Fed responses to news (Panels C and

bond yields—especially at the short end—moved lower.

[27]The sample period of Bauer and Swanson's orthogonalized shock ends in February 2020. For this reason, we end the sample period of this local projection analysis in 2019:Q4.

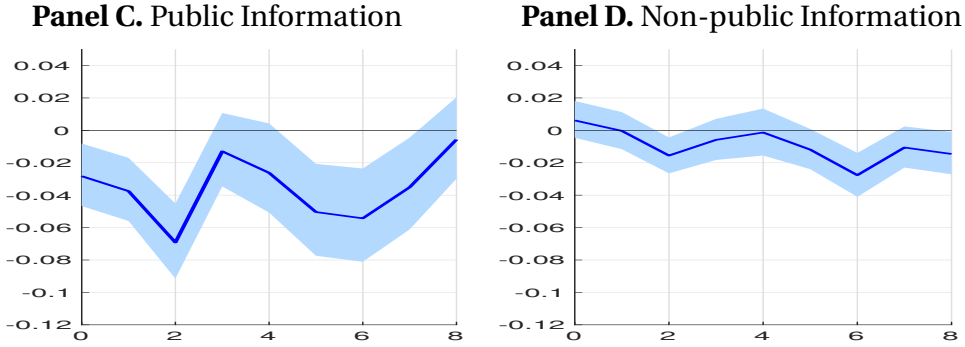
Figure 11: PROPAGATION OF FED'S REACTIONS TO NEWS

(Total Fed's response to news)

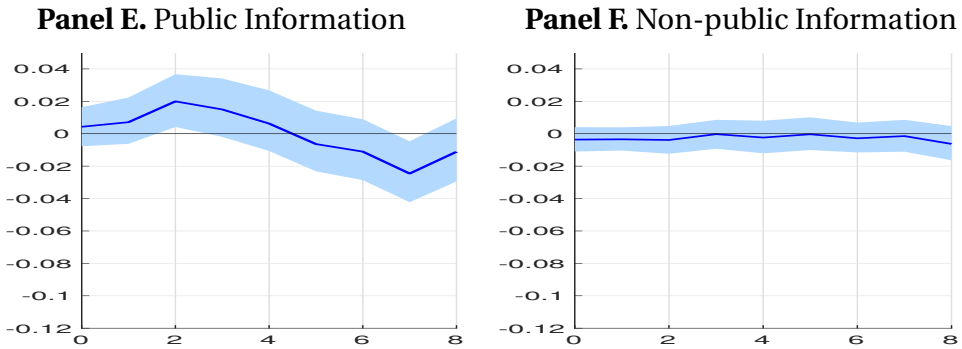


s

(Negative Fed's response to news)



(Positive Fed's response to news)



Notes: The figure shows the responses of disagreement about five-year five-year forward inflation expectations attributable to public information and non-public information following the Fed's response to news from Bauer and Swanson (2023), estimated from the linear model in Equation (31). HAC standard errors are reported. The shaded area indicate the 95% posterior intervals.

Source: Authors' calculation

D)– that is, instances where FOMC members interpret recent economic data releases as having negative implications for economic conditions. In contrast, positive Fed responses to news have small positive effects on disagreement about public information for up to four quarters after the shock, but these effects are largely statistically insignificant. Both negative and positive Fed responses to news have negligible effects on disagreement about non-public information.

This result suggests that forecasters' disagreement about public information is particularly responsive to negative Fed responses to news, and the Fed's communication about recent economic situations can help reduce disagreement about future inflation and anchor long-run inflation expectations in times of economic downturn and high inflation uncertainty. This observation also confirms that disagreement about public information is the component that can be reduced through monetary policy, highlighting the channel of heterogeneous responses to public information as the mechanism through which monetary policy anchors forecasters' inflation expectations.[28]

This result has three important implications. First, our empirical model effectively identifies the sources of information that contribute to disagreement, underscoring that the portion attributable to public information is the component that is reducible through effective monetary policy and is relevant for inflation anchoring. Second, the Fed's response to news plays a crucial role in shaping public information essential for inflation forecasting. In this sense, timely policy implementation and communication can help anchor economic agents' expectations by reducing their disagreement about future macroeconomic conditions. Finally, the share of disagreement driven by public information can serve as an auxiliary indicator of the degree to which inflation expectations are anchored.

8.2 Disagreement and Monetary Policy Effectiveness

Next, we investigate the extent to which disagreement about public information affects the transmission of monetary policy.

For this analysis, we estimate a nonlinear local projection model and compare the transmission of monetary policy shocks when the main source of disagreement about future inflation

[28] We consider several robustness checks. First, we replace the disagreement attributable to non-public information with that attributable solely to private information, and the effects of the Fed's response to news on this component are close to zero and are not statistically significant. Second, we include a measure of individual uncertainty from Binder (2017) to distinguish disagreement from uncertainty. The results remains robust, as shown in Figure E5 of Appendix E.

is public information versus non-public information. We consider two regimes: in Regime 1, non-public information is the source of disagreement, while in Regime 2, public information is the source. These regimes are distinguished by the disagreement share of public information, as introduced in Section 7, and we focus on 5-year, 5-year forward forecasts.

Consider the following nonlinear local projection model.

$$y^j_{t+h} = \alpha^j_h + \beta^j_{1,h}(1 - s_{t-1})z_t + \beta^j_{2,h}s_{t-1}z_t + \Gamma^j_h \mathbf{X}_{t-1} + e^j_{t+h} \quad \text{for} \quad j \in [p, n], \quad h = 0, 1, \cdots, H. \quad (32)$$

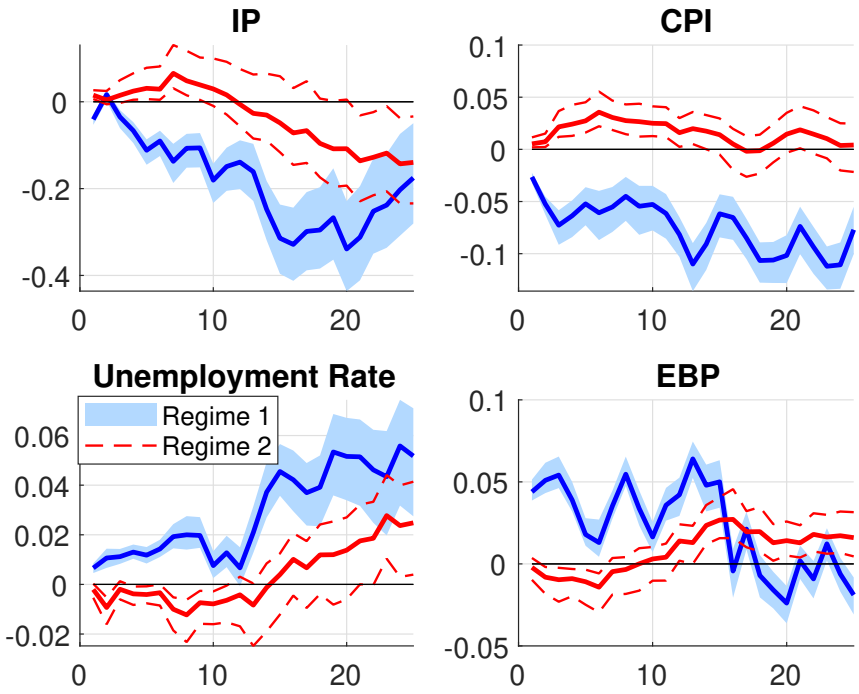
As the dependent variable (y^j_{t+h}), we consider the growth rate of industrial production, CPI inflation, the unemployment rate, and the excess bond premium. Second, we use the orthogonalized monetary policy shock from Bauer and Swanson (2022) as the externally identified shock. In this model, s_{t-1} is an indicator of Regime 2 and is measured using $\theta^{pub}_{5-10,t-1}$ from Equation (30). The indicator of Regime 1 is thus $(1 - s_{t-1})$, measured with $(1 - \theta^{pub}_{5-10,t-1})$. Since the disagreement share is estimated quarterly, we assign the same quarterly value to the three months of the corresponding quarter. Note that we use a one-quarter lag of s_t to avoid contemporaneous feedback from policy actions (Auerbach and Gorodnichenko, 2013). The parameter $\beta^j_{1,h}$ captures the magnitude of pass-through when disagreement is driven by private information (Regime 1), while $\beta^j_{2,h}$ reflects the magnitude when public information is the source of disagreement (Regime 2). The parameter draws from the Gibbs sampler are used to estimate $\theta^{pub}_{5-10,t-1}$ and hence s_{t-1}, which are then used to compute impulse responses. Finally, to maintain consistency with previous studies, our local projection model is specified at a monthly frequency. We consider 12 lags for the macroeconomic controls and two lags for the monetary policy shocks.[29] The sample period is from October 1991 to December 2019.

Figure 12 presents the impulse responses for the two regimes. We report the effects of a one-standard-deviation innovation in the orthogonalized monetary policy shock. In Panel A, the responses of regimes 1 and 2 are scaled by 0.9 and 0.1, respectively, reflecting the average disagreement shares of non-public information and public information.[30] For the responses of IP growth and CPI inflation, we report the cumulative effects of the monetary policy shock,

[29] As a robustness check, we include the uncertainty measure from Binder (2017), monetary policy uncertainty from Husted et al. (2020), and consensus inflation expectations from the SPF for the next year as additional macro controls. For consensus expectations, we use the data as they are instead of using the model-implied consensus to avoid potential issues specific to our model. The estimation results are robust to the inclusion of additional macro controls, and the results are reported in Appendix Appendix E.

[30] This approach ensures that the sum of the two impulse responses closely approximates the average responses.

Figure 12: PROPAGATION OF MONETARY POLICY SHOCKS OF THE TWO REGIMES (5-10 YEAR AHEAD)



Notes: The figure reports the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2023). The figure shows the impulse responses of Regimes 1 and 2 scaled by 0.9 and 0.1, respectively. The blue line shows the responses when non-public information is the source of disagreement (Regime 1), while the red line represents the responses when public information is the source of disagreement (Regime 2). The upper left figure shows the cumulative response of percent changes in industrial production; the upper right figure shows the cumulative responses of percent changes in the CPI; the bottom left figure shows the responses of the unemployment rate; and the bottom right figure displays the response of the excess bond premium (EBP). The shaded areas and dashed lines represent the 95% posterior intervals.

Sources: Authors' calculation

providing estimates that directly reflect changes in levels (Stock and Watson, 2018). Significant differences in the effects of monetary policy are observed between the two regimes. In regime 1 (blue lines), where non-public information is the source of disagreement, contractionary monetary policy has rapid and statistically significant effects on macroeconomic variables. In contrast, in regime 2 (magenta lines), where public information is the source of disagreement, the overall monetary policy effects become weaker and macroeconomic variables show delayed responses. Notably, a price puzzle emerges in the second regime.

Our empirical findings have important implications for the effectiveness of monetary policy. First, anchored inflation expectations are crucial for the effectiveness of monetary policy. When forecasters disagree about public information, including the stance of monetary policy, the stabilizing effects of monetary policy weaken significantly. Second, our empirical result also offers a

new perspective on the source of the price puzzle. Our result suggests that heterogeneous responses to public information can be an important channel driving the price puzzle. This implies that the price puzzle may be closely related to factors affecting anchored inflation expectations, including the effectiveness of monetary policy communication or the credibility of the central bank.[31] Overall, our empirical results suggest that well-anchored inflation expectations and clear monetary policy communication enhance policy effectiveness by reducing disagreement about future inflation attributable to public information.

9 Conclusion

This paper develops a novel empirical framework to analyze disagreement about inflation expectations among professional forecasters, capturing both long-term beliefs and time-varying responses to public and private information across various forecasting horizons. By extending the Nelson-Siegel model to individual-level Survey of Professional Forecasters (SPF) data, we identify two critical dimensions, level and slope, that shape the individual term structure of inflation expectations. Our Bayesian dynamic factor model decomposes these expectations into components driven by the three distinct information sources, uncovering why forecasters disagree at each point in time and across forecasting horizons and the extent to which monetary policy influences forecasters' inflation expectations.

Our analysis reveals that in normal times, long-run disagreement primarily reflects divergent long-term beliefs, whereas short-run disagreement largely arises from private information. However, during periods of economic turbulence, such as the Great Recession and the COVID-19 pandemic, heterogeneous responses to public information become a dominant source of disagreement across all horizons. Crucially, we find that this public-information driven disagreement is an important determinant of the effectiveness of monetary policy, delaying policy transmission and generating a price puzzle. In contrast, disagreement driven by private information or individual long-term beliefs has negligible effects on the transmission of monetary policy.

[31] Our finding aligns with Dong et al. (2024), who report that households' disagreement, as measured in the Michigan survey, weakens the effectiveness of monetary policy. Falck et al. (2021) also find that the price puzzle becomes more pronounced when professional forecasters disagree, based on SPF data. While our findings are consistent with theirs, we extend the analysis by highlighting the role of the information source behind the disagreement. In particular, the public information share of disagreement rises with overall disagreement, helping to reconcile previous empirical findings despite differences in data and methodology.

We show that the Fed's monetary policy communication is particularly effective during economic downturns and episodes of heightened inflation uncertainty. The Fed can reduce disagreement stemming from public information, thereby reinforcing the anchoring of inflation expectations and facilitating the achievement of its economic stabilization goals.

Overall, our approach underscores the importance of distinguishing between information sources of disagreement in inflation expectations. Policymakers should pay close attention not only to consensus forecasts but also to the structure and sources of disagreement among forecasters to assess the strength of the inflation anchor. Understanding these dynamics can provide essential insights for better monetary policy design, improved communication strategies, and more robust expectation anchoring in future economic crises.

References

Andrade, P., Crump, R. K., Eusepi, S., and Moench, E. (2016). Fundamental disagreement. *Journal of Monetary Economics*, 83:106–128.

Andrade, P., Gaballo, G., Mengus, E., and Mojon, B. (2019). Forward guidance and heterogeneous beliefs. *American Economic Journal: Macroeconomics*, 11(3):1–29.

Aruoba, S. B. (2020). Term structures of inflation expectations and real interest rates. *Journal of Business & Economic Statistics*, 38(3):542–553.

Auerbach, A. J. and Gorodnichenko, Y. (2013). Output spillovers from fiscal policy. *American Economic Review*, 103(3):141–46.

Ball, L. and Mazumder, S. (2011). Inflation Dynamics and the Great Recession. *Brookings Papers on Economic Activity*, 42(1 (Spring):337–405.

Banbura, M. and Modugno, M. (2014). Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data. *Journal of Applied Econometrics*, 29(1):133–160.

Bauer, M. D. and Swanson, E. T. (2022). A Reassessment of Monetary Policy Surprises and High-Frequency Identification. In *NBER Macroeconomics Annual 2022, volume 37*, NBER Chapters, pages 87–155. National Bureau of Economic Research, Inc.

Bauer, M. D. and Swanson, E. T. (2023). An alternative explanation for the "fed information effect". *American Economic Review*, 113(3):664–700.

Beveridge, S. and Nelson, C. (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'. *Journal of Monetary Economics*, 7(2):151–174.

Binder, C. C. (2017). Measuring uncertainty based on rounding: New method and application to inflation expectations. *Journal of Monetary Economics*, 90(C):1–12.

Bundick, B. and Smith, A. L. (2020). Did the Federal Reserve Break the Phillips Curve? Theory and Evidence of Anchoring Inflation Expectations. Research Working Paper RWP 20-11, Federal Reserve Bank of Kansas City.

Castillo-Martinez, L. and Reis, R. (2024). How do central banks control inflation? A guide for the perplexed. Discussion Papers 2433, Centre for Macroeconomics (CFM).

Chan, J. C. C., Koop, G., and Potter, S. M. (2013). A new model of trend inflation. *Journal of Business & Economic Statistics*, 31(1):94–106.

Clarida, R. (2021). The federal reserve's new framework and outcome-based forward guidance. Speech, Board of Governors of the Federal Reserve System.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159.

Coibion, O. and Gorodnichenko, Y. (2025a). Inflation, expectations and monetary policy: What have we learned and to what end? Working Paper 33858, National Bureau of Economic Research.

Coibion, O. and Gorodnichenko, Y. (2025b). Inflation, expectations and monetary policy: What have we learned and to what end? Working Paper 33858, National Bureau of Economic Research.

Constâncio, V. (2015). Understanding inflation dynamics and monetary policy.

Crump, R. K., Eusepi, S., Moench, E., and Preston, B. (2023). Chapter 17 - the term structure of expectations. In Bachmann, R., Topa, G., and van der Klaauw, W., editors, *Handbook of Economic Expectations*, pages 507–540. Academic Press.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38.

Diebold, F. X., Li, C., and Yue, V. Z. (2008). Global yield curve dynamics and interactions: a dynamic nelson–siegel approach. *Journal of Econometrics*, 146(2):351–363.

Dong, D., Liu, Z., Wang, P., and Wei, M. (2024). Inflation disagreement weakens the power of monetary policy. Working paper, The Federal Reserve Bank of San Francisco.

Ehrmann, M., Gaballo, G., Hoffmann, P., and Strasser, G. (2019). Can more public information raise uncertainty? The international evidence on forward guidance. *Journal of Monetary Economics*, 108(C):93–112.

Falck, E., Hoffmann, M., and Hürtgen, P. (2021). Disagreement about inflation expectations and monetary policy transmission. *Journal of Monetary Economics*, 118(C):15–31.

Farmer, L., Nakamura, E., and Steinsson, J. (2021). Learning About the Long Run. NBER Working Papers 29495, National Bureau of Economic Research, Inc.

Federal Reserve Bank of New York (2023). Multivariate core trend (mct) inflation. https://www.newyorkfed.org/research/policy/mct. Applied Macroeconomics and Econometrics Center.

Fisher, J., Melosi, L., and Rast, S. (2022). Anchoring long-run inflation expectations in a panel of professional forecasters. Working paper.

Fitzgerald, T., Jones, C., Kulish, M., and Nicolini, J. P. (2024). Is there a stable relationship between unemployment and future inflation? *American Economic Journal: Macroeconomics*, 16(4):114–42.

Fofana, S, P. P. and Reis, R. (2024). Household disagreement about expected inflation. CEPR Discussion Paper No. 18956, CEPR Press.

Forbes (2008). Bernanke: Runaway inflation? no way. *Forbes*.

Friedrich, C. (2016). Global inflation dynamics in the post-crisis period: What explains the puzzles? *Economics Letters*, 142(C):31–34.

Fujita, S. and Ramey, G. (2009). The cyclicality of separation and job finding rates*. *International Economic Review*, 50(2):415–430.

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720.

Glas, A. and Hartmann, M. (2016). Inflation uncertainty, disagreement and monetary policy: Evidence from the ecb survey of professional forecasters. *Journal of Empirical Finance*, 39:215–228. SI: Euro Zone in Crisis.

Hazell, J., Herreño, J., Nakamura, E., and Steinsson, J. (2022). The slope of the phillips curve: Evidence from u.s. states*. *The Quarterly Journal of Economics*, 137(3):1299–1344.

Herbst, E. P. and Winkler, F. (2021). The Factor Structure of Disagreement. Finance and Economics Discussion Series 2021-046, Board of Governors of the Federal Reserve System (U.S.).

Husted, L., Rogers, J., and Sun, B. (2020). Monetary policy uncertainty. *Journal of Monetary Economics*, 115:20–36.

Kim, H. (2023). Quantifying the Sources of Forecaster Disagreement. mimeo 4659361, University of Texas at Austin.

Lahiri, K. and Sheng, X. (2008). Evolution of forecast disagreement in a Bayesian learning model. *Journal of Econometrics*, 144(2):325–340.

Maćkowiak, B., Matějka, F., and Wiederholt, M. (2023). Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–73.

McLeay, M. and Tenreyro, S. (2020). Optimal inflation and the identification of the phillips curve. *NBER Macroeconomics Annual*, 34:199–255.

Mertens, E. (2016). Measuring the level and uncertainty of trend inflation. *Review of Economics and Statistics*, 98(5):950–967.

Mertens, T. M. and Williams, J. C. (2019). Monetary policy frameworks and the effective lower bound on interest rates. Staff Reports 877, Federal Reserve Bank of New York.

Nelson, C. R. and Siegel, A. F. (1987). Parsimonious modeling of yield curves. *Journal of business*, pages 473–489.

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

Reis, R. (2020). The People versus the Markets: A Parsimonious Model of Inflation Expectations. Discussion Papers 2033, Centre for Macroeconomics (CFM).

Stock, J. H. and Watson, M. W. (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking*, 39(s1):3–33.

Stock, J. H. and Watson, M. W. (2016). Core Inflation and Trend Inflation. *The Review of Economics and Statistics*, 98(4):770–784.

Stock, J. H. and Watson, M. W. (2018). Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. *Economic Journal*, 128(610):917–948.

Appendix A State-Space Representation of Nelson-Siegel Model

This section provides a full characterization of the state-space representation of the Nelson-Siegel model's equations, along with detailed definitions of all of the coefficient vectors and matrices.

A.1 State Equation

We start with the state equation. Let the $(2(n+1)\times 1)$ state vector $\mathbf{x}_t$ be defined as

$$\mathbf{x}_t \coloneqq \begin{bmatrix} L_t, S_t, \varepsilon^L_{1,t}, \varepsilon^S_{1,t}, \cdots, \varepsilon^L_{n,t}, \varepsilon^S_{n,t} \end{bmatrix}' \tag{A1}$$

Define the transition matrix $\mathbf{F}$ as

$$\mathbf{F} \coloneqq \begin{bmatrix} \rho_L & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & \rho_S & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & \bar{\rho}_L & 0 & 0\dots & 0 & 0 \\ 0 & 0 & 0 & \bar{\rho}_S & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & \bar{\rho}_L & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & \bar{\rho}_S \end{bmatrix}. \tag{A2}$$

Let $\mathbf{u}_t$ be the vector of shocks to the state vector defined as follows:

$$\mathbf{u}_t \coloneqq \begin{bmatrix} u_{L,t}, u_{S,t}, u^L_{1,t}, u^S_{1,t}, \cdots, u^L_{n,t}, u^S_{n,t} \end{bmatrix}'. \tag{A3}$$

The covariance matrix of the shocks is given by:

$$\mathbf{Q} = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & \sigma_L^2 & 0 & 0\dots & 0 & 0 \\ 0 & 0 & 0 & \sigma_S^2 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & \sigma_L^2 & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & \sigma_S^2 \end{bmatrix}. \tag{A4}$$

Finally, we arrive at the state equation:

$$\mathbf{x}_t = \mathbf{F}\mathbf{x}_{t-1} + \mathbf{u}_t, \quad \mathbf{u}_t \sim N(\mathbf{0}, \mathbf{Q}). \tag{A5}$$

A.2 Measurement Equation

The measurement equation is of the form:

$$\mathbf{y}_t = \mu_y + \mathbf{H}\mathbf{x}_t + \mathbf{v}_t, \quad \mathbf{v}_t \sim N(\mathbf{0}, \mathbf{R}), \tag{A6}$$

where μ_y is a vector of forecaster fixed effects, $\mathbf{H}$ is a matrix of factor loadings on the aggregate and idiosyncratic level and slope factors, and $\mathbf{v}_t$ is a vector of measurement errors. The rest of the section details each component in Equation (10).

A.2.1 The Vector of Observations: $\mathbf{y}_t$

For estimation, we use one-quarter to four-quarter ahead fixed-horizon forecasts and two-year forward, three-year forward, five-year average, and ten-year average fixed event forecasts. For the five-year and ten-year average forecasts, we use observed nowcasts and one-quarter backcasts

when available, and realized inflation two quarters and three quarters prior from the most recent CPI vintage at the time the survey was conducted to capture realized inflation.

The observation vector for any period $\mathbf{y}_t$ is given by

$$\mathbf{y}_t = \begin{bmatrix} \pi_{1,t\to t+1|t}, & \pi_{1,t+1\to t+2|t}, & \pi_{1,t+2\to t+3|t}, & \pi_{1,t+3\to t+4|t}, & \cdots \\ \pi_{1,t+3\to t+7|t}, & \pi_{1,t+2\to t+6|t}, & \pi_{1,t+1\to t+5|t}, & \pi_{1,t\to t+4|t}, & \cdots \\ \pi_{1,t+7\to t+11|t}, & \pi_{1,t+6\to t+10|t}, & \pi_{1,t+5\to t+9|t}, & \pi_{1,t+4\to t+8|t}, & \cdots \\ \pi_{1,t\to t+19|t}, & \pi_{1,t\to t+18|t}, & \pi_{1,t\to t+17|t}, & \pi_{1,t\to t+16|t}, & \cdots \\ \pi_{1,t\to t+39|t}, & \pi_{1,t\to t+38|t}, & \pi_{1,t\to t+37|t}, & \pi_{1,t\to t+36|t}, & \cdots \\ & & \cdots\cdots & & \\ & & \cdots \quad \pi_{n,t\to t+37|t}, & \pi_{n,t\to t+36|t} & \end{bmatrix}'. \tag{A7}$$

The first four elements of $\mathbf{y}_t$ correspond to fixed horizon forecasts of one to four quarters ahead and are typically observed every period. Only four of the final sixteen elements of $\mathbf{y}_t$ are observed in any given quarter. These final sixteen elements correspond to fixed event forecasts, where each group of four correspond to the fixed event correctly mapped to the quarter in which the survey was conducted.

For the final eight elements, which correspond to forecasts of average inflation over five and ten year periods including the current calendar year, we must adjust them to account for the fact that they include realized inflation over previous quarters. Specifically,

- In Q1, we define

$$\pi_{it\to t+19|t} = \frac{4}{19}\left(5\pi_{it-1\to t+19|t} - \frac{1}{4}\pi_{it-1\to t|t}\right)$$
$$\pi_{it\to t+39|t} = \frac{4}{39}\left(10\pi_{it-1\to t+19|t} - \frac{1}{4}\pi_{it-1\to t|t}\right)$$

- In Q2, we define

$$\pi_{it\to t+18|t} = \frac{4}{18}\left(5\pi_{it-1\to t+19|t} - \frac{1}{4}\pi_{it-1\to t|t} - \frac{1}{4}\pi_{it-2\to t-1|t}\right)$$

$$\pi_{it\to t+38|t} = \frac{4}{38}\left(10\pi_{it-1\to t+19|t} - \frac{1}{4}\pi_{it-1\to t|t} - \frac{1}{4}\pi_{it-2\to t-1|t}\right)$$

- In Q3, we define

$$\pi_{it\to t+17|t} = \frac{4}{17}\left(5\pi_{it-1\to t+19|t} - \frac{1}{4}\pi_{it-1\to t|t} - \frac{1}{4}\pi_{it-2\to t-1|t} - \frac{1}{4}\pi_{it-3\to t-2|t}\right)$$
$$\pi_{it\to t+37|t} = \frac{4}{37}\left(10\pi_{it-1\to t+19|t} - \frac{1}{4}\pi_{it-1\to t|t} - \frac{1}{4}\pi_{it-2\to t-1|t} - \frac{1}{4}\pi_{it-3\to t-2|t}\right)$$

- In Q4, we define

$$\pi_{it\to t+16|t} = \frac{4}{16}\left(5\pi_{it-1\to t+19|t} - \frac{1}{4}\pi_{it-1\to t|t} - \frac{1}{4}\pi_{it-2\to t-1|t} - \frac{1}{4}\pi_{it-3\to t-2|t} - \frac{1}{4}\pi_{it-4\to t-3|t}\right)$$
$$\pi_{it\to t+36|t} = \frac{4}{36}\left(10\pi_{it-1\to t+19|t} - \frac{1}{4}\pi_{it-1\to t|t} - \frac{1}{4}\pi_{it-2\to t-1|t} - \frac{1}{4}\pi_{it-3\to t-2|t} - \frac{1}{4}\pi_{it-4\to t-3|t}\right)$$

For the nowcasts $\pi_{it-1\to t|t}$ and backcasts $\pi_{it-2\to t-1|t}$, we use the reported values from the SPF. For the two and three period backcasts $\pi_{it-3\to t-2|t}$ and $\pi_{it-4\to t-3|t}$, we use the most recently available vintage of the CPI at the time that the forecast was made.

A.2.2 The Vector of Forecaster Fixed Effects: μ_y

We define the loading function on the slope factor for forecasts of inflation between horizons at two dates $t+h_1$ and $t+h_2$ as

$$f_S(h_1,h_2) = \frac{e^{-\lambda h_1} - e^{-\lambda h_2}}{\lambda(h_2-h_1)} \tag{A8}$$

This expression is used in μ_y and $\mathbf{H}$. Define the constant vector in the measurement equation, μ_y, as

$$\mu_y \coloneqq \begin{bmatrix} \alpha_L - \alpha_S f_S(0,1), & \alpha_L - \alpha_S f_S(1,2), & \alpha_L - \alpha_S f_S(2,3), & \alpha_L - \alpha_S f_S(3,4), \\ \alpha_L - \alpha_S f_S(3,7), & \alpha_L - \alpha_S f_S(2,6), & \alpha_L - \alpha_S f_S(1,5), & \alpha_L - \alpha_S f_S(0,4), \\ \alpha_L - \alpha_S f_S(7,11), & \alpha_L - \alpha_S f_S(6,10), & \alpha_L - \alpha_S f_S(5,9), & \alpha_L - \alpha_S f_S(4,8) \\ \alpha_L - \alpha_S f_S(0,19), & \alpha_L - \alpha_S f_S(0,18), & \alpha_L - \alpha_S f_S(0,17), & \alpha_L - \alpha_S f_S(0,16), \\ \alpha_L - \alpha_S f_S(0,39), & \alpha_L - \alpha_S f_S(0,38), & \alpha_L - \alpha_S f_S(0,37), & \alpha_L - \alpha_S f_S(0,36), \\ & \cdots\cdots & & \\ & \cdots\cdots \alpha_L - \alpha_S f_S(0,37), & \alpha_L - \alpha_S f_S(0,36) & \end{bmatrix}'. \tag{A9}$$

Define the error term in the measurement equation, $\mathbf{v}_t$, as

$$\mathbf{v}_t \coloneqq \begin{bmatrix} v_{1,1,t}, & v_{1,2,t}, & v_{1,3,t}, & v_{1,4,t}, \\ v_{1,5,t}, & v_{1,6,t}, & v_{1,7,t}, & v_{1,8,t}, \\ v_{1,9,t}, & v_{1,10,t}, & v_{1,11,t}, & v_{1,12,t}, \\ v_{1,13,t}, & v_{1,14,t}, & v_{1,15,t}, & v_{1,16,t}, \\ v_{1,17,t}, & v_{1,18,t}, & v_{1,19,t}, & v_{1,20,t}, \\ & \cdots\cdots & & \\ & \cdots\cdots v_{n,19,t}, & v_{n,20,t} & \end{bmatrix}'. \tag{A10}$$

The covariance matrix of the measurement error vector ($\mathbf{v}_t$), $\mathbf{R}$, is given by the following

diagonal matrix.

$$\mathbf{R} := \text{diag}\left(\left[\begin{array}{cccc} \sigma^2_{v,1}, & \sigma^2_{v,2}, & \sigma^2_{v,3}, & \sigma^2_{v,4}, \\ \sigma^2_{v,5}, & \sigma^2_{v,6}, & \sigma^2_{v,7}, & \sigma^2_{v,8}, \\ \sigma^2_{v,9}, & \sigma^2_{v,10}, & \sigma^2_{v,11}, & \sigma^2_{v,12}, \\ \sigma^2_{v,13}, & \sigma^2_{v,14}, & \sigma^2_{v,15}, & \sigma^2_{v,16} \\ \sigma^2_{v,17}, & \sigma^2_{v,18}, & \sigma^2_{v,19}, & \sigma^2_{v,20} \\ & & \cdots\cdots & \\ & & \cdots\cdots, \sigma^2_{v,19}, & \sigma^2_{v,20} \end{array}\right]'\right) \tag{A11}$$

Note that the argument in the square bracket is a vector.

Finally, we define the measurement equation mapping matrix $\mathbf{H}$ as

$$\mathbf{H} \coloneqq \begin{bmatrix}
\beta_1^L & -\beta_1^S f_S(0,1) & 1 & f_S(0,1) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(1,2) & 1 & f_S(1,2) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(2,3) & 1 & f_S(2,3) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(3,4) & 1 & f_S(3,4) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(3,7) & 1 & f_S(3,7) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(2,6) & 1 & f_S(2,6) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(1,5) & 1 & f_S(1,5) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(0,4) & 1 & f_S(0,4) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(7,11) & 1 & f_S(7,11) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(6,10) & 1 & f_S(6,10) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(5,9) & 1 & f_S(5,9) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(4,8) & 1 & f_S(4,8) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(0,19) & 1 & f_S(0,19) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(0,18) & 1 & f_S(0,18) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(0,17) & 1 & f_S(0,17) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(0,16) & 1 & f_S(0,16) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(0,39) & 1 & f_S(0,39) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(0,38) & 1 & f_S(0,38) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(0,37) & 1 & f_S(0,37) & \dots & 0 & 0 \\
\beta_1^L & -\beta_1^S f_S(0,36) & 1 & f_S(0,36) & \dots & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
\beta_n^L & -\beta_n^S f_S(0,37) & 0 & 0 & \dots & 1 & f_S(0,37) \\
\beta_n^L & -\beta_n^S f_S(0,36) & 0 & 0 & \dots & 1 & f_S(0,36)
\end{bmatrix}. \tag{A12}$$

A.2.3 Remark

In the measurement equation, each series in $\mathbf{y}_t$ is assumed to be observed with measurement error. The fixed event forecasts are treated separately in each quarter throughout the calendar year to reflect the fact that the forecasting horizon shrinks as the calendar year progresses. This leaves us with a total of 20 observables for each forecaster in each quarter, 12 of which are missing by construction.

Appendix B Gibbs Sampler

1. Sample $\boldsymbol{\theta}_1 = \{\alpha_{L,i}, \alpha_{S,i}, \beta_{L,i}, \beta_{S,i}\}_{i=1}^{N}$ conditional on remaining parameters.

   Since all of the shocks are assumed to be independent, this can be treated as a separate regression model for each forecaster i. We assume independent, multivariate normal priors for each group of four parameters $[\alpha_{L,i}, \alpha_{S,i}, \beta_{L,i}, \beta_{S,i}]'$ across each foracaster i, where

$$\begin{bmatrix} \alpha_{L,i} \\ \alpha_{S,i} \\ \beta_{L,i} \\ \beta_{S,i} \end{bmatrix} \sim N(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$$

2. Sample $\boldsymbol{\theta}_2 = \text{vec}(\boldsymbol{A})$ conditional on remaining parameters. Since L_t and S_t are observed, this is a standard multivariate regression model. We impose stationarity through the prior.

3. Sample $\boldsymbol{\theta}_3 = \text{vec}(\boldsymbol{B})$ conditional on remaining parameters. Since $\varepsilon_{L,it}$ and $\varepsilon_{S,it}$ are observed for every forecaster $i = 1, \ldots, n$, this is a standard multivariate regression model with known covariance matrix where we pool the data across forecasters. We impose stationarity through the prior.

4. Sample $\boldsymbol{\theta}_4 = [\sigma_L^2, \sigma_S^2]'$ conditional on remaining parameters. Since $\varepsilon_{L,it}$ and $\varepsilon_{S,it}$ are observed for every forecaster $i = 1, \ldots, n$, this is a standard variance estimation problem with known regression coefficients where we pool the data across forecasters.

5. Sample $\boldsymbol{\theta}_5 = \lambda$ with a Metropolis Hastings step conditional on remaining parameters. It boils down to a nonlinear regression problem.

6. Sample $\boldsymbol{\theta}_6 = \left[\sigma_{v,1}^2, \ldots, \sigma_{v,20}^2\right]'$ conditional on remaining parameters. Given other parameters, $\mathbf{v}_t$ is directly observed.

7. Sample $\boldsymbol{\theta}_7 = \{\boldsymbol{x}_t\}_{t=1}^{T}$ conditional on remaining parameters using a simulation smoother.

Appendix C Noisy Information Model Proofs

Proof of Proposition 1. We seek to derive an approximate factor structure that aligns with our statistical model.

We start from the evolution of the nowcast

$$\mathbb{F}_{it}[\boldsymbol{x}_t] = \boldsymbol{b}_i + \sum_{k=0}^{\infty} \boldsymbol{\lambda}_{i,k}\boldsymbol{\varepsilon}_{t-k} + \boldsymbol{u}_{it} + \boldsymbol{v}_{it},$$

and focus on one equation at a time. Beginning with the trend, we have that the weights on fundamental shocks are given by

$$\begin{aligned}
\lambda_{ik}^{\tau} &= \sum_{j=0}^{k} (1-g_{i,\tau})^j g_{i,\tau}\sigma_\tau \\
&= \sigma_\tau g_{i,\tau} \frac{1-(1-g_{i,\tau})^{k+1}}{g_{i,\tau}} \\
&= \sigma_\tau \left[1-(1-g_{i,\tau})^{k+1}\right]
\end{aligned}$$

The above expression allows us to rewrite the beliefs about the trend as

$$\begin{aligned}
\mathbb{F}_{it}[\tau_t] &= b_i^{\tau} + \sum_{k=0}^{\infty} \lambda_{ik}^{\tau}\varepsilon_{t-k}^{\tau} + u_{it}^{\tau} + v_{it}^{\tau} \\
&= \underbrace{b_i^{\tau}}_{\text{long-term belief}} + \underbrace{\sum_{k=0}^{\infty}\left[1-(1-g_{i,\tau})^{k+1}\right]\sigma_\tau\varepsilon_{t-k}^{\tau} + \sum_{k=0}^{\infty} g_{i,\tau}(1-g_{i,\tau})^k\theta_\tau z_{t-k}^{\tau}}_{\text{common component}} \\
&\quad + \underbrace{\sum_{k=0}^{\infty} g_{i,\tau}(1-g_{i,\tau})^k\omega_{i,\tau}w_{it-k}^{\tau}}_{\text{idiosyncratic component}} \qquad \text{(C13)} \\
&= b_i^{\tau} + \mathbb{F}_{it,common}[\tau_t] + \mathbb{F}_{it,idiosyncratic}[\tau_t] \qquad \text{(C14)}
\end{aligned}$$

Define the following weights:

$$a_{ik}^{\tau} := \left[1-(1-g_{i,\tau})^{k+1}\right]\sigma_\tau$$

$$
\begin{aligned}
d_{ik}^{\tau} &:= g_{i,\tau}(1-g_{i,\tau})^{k}\theta_{\tau} \\
\tilde{a}_{k}^{\tau} &:= \left[1-(1-g_{\tau})^{k+1}\right]\sigma_{\tau} \\
\tilde{d}_{k}^{\tau} &:= g_{\tau}(1-g_{\tau})^{k}\theta_{\tau}
\end{aligned}
$$

We seek to approximate the common component in (C14) as an i-specific constant times an aggregate component, so that:

$$
\mathbb{F}_{it,common}[\tau_t] \approx \tilde{\beta}_i^L \tilde{L}_t.
$$

and propose the following decomposition:

$$
\tilde{\beta}_i^L = \frac{g_{i,\tau}}{g_{\tau}} \tag{C15}
$$

$$
\tilde{L}_t = \sum_{k=0}^{\infty} \tilde{a}_k^{\tau} \varepsilon_{t-k}^{\tau} + \sum_{k=0}^{\infty} \tilde{d}_k^{\tau} z_{t-k}^{\tau}. \tag{C16}
$$

Working first with $\tilde{a}_{ik}^{\tau}$ we have

$$
\begin{aligned}
a_{ik}^{\tau} &= \left[1-(1-g_{i,\tau})^{k+1}\right]\sigma_{\tau} \\
&= g_{i,\tau}\frac{1-(1-g_{i,\tau})^{k+1}}{g_{i,\tau}}\sigma_{\tau} \\
&\approx g_{i,\tau}\frac{1-(1-g_{\tau})^{k+1}}{g_{\tau}}\sigma_{\tau} \\
&= \frac{g_{i,\tau}}{g_{\tau}}\left[1-(1-g_{\tau})^{k+1}\right]\sigma_{\tau} \\
&= \tilde{\beta}_i^L \tilde{a}_k^{\tau}
\end{aligned}
$$

Note that the third approximate equality holds because

$$
1-(1-g)^{k+1} = \sigma_{\tau}\left[(k+1) - \frac{(k+1)k}{2}g + \mathcal{O}\left(g^2\right)\right]
$$

Hence, for small g, the function $g \mapsto 1-(1-g)^{k+1}$ is approximately flat in g.

Turning to d_{ik}^{τ}, we have

$$\begin{aligned} d_{ik}^{\tau} &= g_{i,\tau}(1-g_{i,\tau})^k\theta_\tau \\ &= \frac{g_{i,\tau}}{g_\tau} g_\tau (1-g_{i,\tau})^k \theta_\tau \\ &\approx \frac{g_{i,\tau}}{g_\tau} g_\tau (1-g_\tau)^k \theta_\tau \\ &= \tilde{\beta}_i^L \tilde{d}_k^{\tau} \end{aligned}$$

Note that the third approximate equality holds because

$$(1-g_{i,\tau})^k = (1-g_\tau)^k + \mathcal{O}\left(k(1-g_\tau)^{k-1}(g_{i,\tau}-g_\tau)\right).$$

Combining these two results for the weights, we obtain the representation

$$\begin{aligned} \mathbb{F}_{it,common}[\tau_t] &= \sum_{k=0}^{\infty} a_{ik}^{\tau}\varepsilon_{t-k}^{\tau} + \sum_{k=0}^{\infty} d_{ik}^{\tau} z_{t-k}^{\tau} \\ &\approx \frac{g_{i,\tau}}{g_\tau}\left(\sum_{k=0}^{\infty}\tilde{a}_k^{\tau}\varepsilon_{t-k}^{\tau} + \sum_{k=0}^{\infty}\tilde{d}_k^{\tau} z_{t-k}^{\tau}\right) \\ &= \tilde{\beta}_i^L \tilde{L}_t \end{aligned}$$

In order to pin down the relative magnitudes of β_i^L and L_t, we use our identification assumption that the conditional variance of L_t is 1. The conditional variance of common component of beliefs is

$$\mathrm{Var}_{t-1}\left(\beta_i^L L_t\right) = \mathrm{Var}_{t-1}\left(\mathbb{F}_{it,common}[\tau_t]\right) \approx \frac{g_{i,\tau}^2}{g_\tau^2} g_\tau^2\left(\sigma_\tau^2+\theta_\tau^2\right)$$

Since we know the conditional variance of L_t must be 1, this gives us the expression for the loading β_i^L:

$$\beta_i^L = \tilde{\beta}_i^L g_\tau\left(\sigma_\tau^2+\theta_\tau^2\right)^{1/2} = \left(\frac{g_{i,\tau}}{g_\tau}\right) g_\tau\left(\sigma_\tau^2+\theta_\tau^2\right)^{1/2} \tag{C17}$$

Bai and Ng (2013; JOE) show that for the identification of a common factor and the loadings, one needs an identifying restriction on the variance of the common component and a restriction on the loadings. Note that the unit conditional variance of L_t satisfies one condition. Below, we show that the second condition, the restriction on the loading, is also satisfied. Averaging across all N forecasters, we have:

$$\frac{1}{N}\sum_{i=1}^{N}\beta_i^L \approx g_\tau\left(\sigma_\tau^2+\theta_\tau^2\right)^{1/2}\underbrace{\left\{\frac{1}{N}\sum_{i=1}^{N}\left(\frac{g_{i,\tau}}{g_\tau}\right)\right\}}_{\approx 1} \tag{C18}$$

where the term in brackets is approximately equal to 1 for large N because $g_\tau = \mathbb{E}_i[g_{i,\tau}]$. $\tilde{\rho}_L \approx 1$ because $\bar{\mathbb{F}}_{t,common}[\tau_t]$ is the sum of a term that is approximately a random walk, $\sum_{k=0}^{\infty}\left[1-(1-g_\tau)^{k+1}\right]\sigma_\tau\varepsilon_{t-k}^{\tau}$, and a stationary AR(1) process, $\sum_{k=0}^{\infty}g_\tau(1-g_\tau)^k\theta_\tau z_{t-k}^{\tau}$.

Turning to the idiosyncratic component v_{it}^{τ}, we see that it follows an AR(1) process with persistence $1-g_{i,\tau}$ and conditional standard deviation $\omega_{i,\tau}$, which immediately gives us $\rho_{L,i} = 1-g_{i,\tau}$ and $\sigma_{L,i} = \omega_{i,\tau}$. Finally, if there is a non-zero long-term belief in the forecasts, b_i^{τ}, this will be picked up by the fixed effect of the forecaster, α_i^L.

We repeat a similar derivation for the cyclical component. We have that the weights on fundamental shocks are given by

$$\begin{aligned}
\lambda_{i,k}^c &= \sum_{j=0}^{k} g_{i,c}\sigma_c(1-g_{i,c})^j\rho_i^j\rho^{k-j} \\
&= \sigma_c g_{i,c}\rho^k\sum_{j=0}^{k}\left[\frac{\rho_i}{\rho}(1-g_{i,c})\right]^j \\
&= \sigma_c g_{i,c}\rho^k\frac{1-\left[\frac{\rho_i}{\rho}(1-g_{i,c})\right]^{k+1}}{1-\left[\frac{\rho_i}{\rho}(1-g_{i,c})\right]},
\end{aligned}$$

where ρ is the cross-sectional average of ρ_i.

This allows us to rewrite the beliefs as

$$\mathbb{F}_{it}[c_t] = b_i^c + \sum_{k=0}^{\infty}\lambda_{i,k}^c\varepsilon_{t-k}^c + u_{it}^c + v_{it}^c$$

$$= \underbrace{b_i^c}_{\text{bias}} + \underbrace{\sum_{k=0}^{\infty} \sigma_c g_{i,c} \rho^k \frac{1 - \left[\frac{\rho_i}{\rho}(1-g_{i,c})\right]^{k+1}}{1 - \left[\frac{\rho_i}{\rho}(1-g_{i,c})\right]} \varepsilon_{t-k}^c + \sum_{k=0}^{\infty} g_{i,c}(1-g_{i,c})^k \rho_i^k \theta_c z_{t-k}^c}_{\text{common component}}$$

$$+ \underbrace{\sum_{k=0}^{\infty} g_{i,c}(1-g_{i,c})^k \rho_i^k \omega_{i,c} w_{it-k}^c}_{\text{idiosyncratic component}}$$

$$= b_i^c + \mathbb{F}_{it,common}[c_t] + \mathbb{F}_{it,idiosyncratic}[c_t] \tag{C19}$$

Define the following quantities:

$$\begin{aligned}
g_c &\coloneqq \mathbb{E}_i[g_{i,c}] \\
\rho_c &\coloneqq \mathbb{E}_i[\rho_i] \\
\kappa_i &\coloneqq \frac{\rho_i}{\rho}(1-g_{i,c}) \\
\kappa &\coloneqq \frac{\rho_c}{\rho}(1-g_c) \\
a_{ik}^c &\coloneqq g_{i,c}\sigma_c \rho^k \frac{1-(\kappa_i)^{k+1}}{1-\kappa_i} \\
d_{ik}^c &\coloneqq g_{i,c}(1-g_{i,c})^k \rho_i^k \theta_c \\
\tilde{a}_k^c &\coloneqq g_c \sigma_c \rho^k \frac{1-\kappa^{k+1}}{1-\kappa} \\
\tilde{d}_k^c &\coloneqq g_c(1-g_c)^k \rho_c^k \theta_c
\end{aligned}$$

We seek to approximate the common component in (C19) as an i-specific constant times an aggregate component, so that:

$$\mathbb{F}_{it,common}[c_t] \approx \tilde{\beta}_i^S \tilde{S}_t.$$

and propose the following decomposition:

$$\tilde{\beta}_i^S = \frac{g_{i,c}}{g_c} \tag{C20}$$

$$\tilde{S}_t = \sum_{k=0}^{\infty} \tilde{a}_k^c \varepsilon_{t-k}^c + \sum_{k=0}^{\infty} \tilde{d}_k^c z_{t-k}^c. \tag{C21}$$

Working first with $\tilde{a}^c_{ik}$ we have

$$\begin{aligned} a^{\tau}_{ik} &= g_{i,c}\sigma_c\rho^k\frac{1-(\kappa_i)^{k+1}}{1-\kappa_i} \\ &= \frac{g_{i,c}}{g_c}g_c\sigma_c\rho^k\frac{1-(\kappa_i)^{k+1}}{1-\kappa_i} \\ &\approx \frac{g_{i,c}}{g_c}g_c\sigma_c\rho^k\frac{1-(\kappa)^{k+1}}{1-\kappa} \\ &= \tilde{\beta}^S_i\tilde{a}^c_k \end{aligned}$$

Turning to d^c_{ik}, we have

$$\begin{aligned} d^c_{ik} &= g_{i,c}(1-g_{i,c})^k\rho^k_i\theta_c \\ &= \frac{g_{i,c}}{g_c}g_c(1-g_{i,c})^k\rho^k_i\theta_c \\ &\approx \frac{g_{i,c}}{g_c}g_c(1-g_c)^k\rho^k_c\theta_c \\ &= \tilde{\beta}^S_i\tilde{d}^c_k \end{aligned}$$

Note that both approximations hold to first order with an approximation term which is $\mathcal{O}\left((g_{i,c}-g_c)^2+(\rho_i-\rho_c)^2\right)$, as can be shown through a first-order multivariate Taylor expansion of both functions around the point (g_c,ρ_c).

Combining these two results for the weights, we obtain the representation

$$\begin{aligned} \mathbb{F}_{it,common}[c_t] &= \sum_{k=0}^{\infty}a^c_{ik}\varepsilon^c_{t-k}+\sum_{k=0}^{\infty}d^c_{ik}z^c_{t-k} \\ &\approx \frac{g_{i,c}}{g_c}\left(\sum_{k=0}^{\infty}\tilde{a}^c_k\varepsilon^c_{t-k}+\sum_{k=0}^{\infty}\tilde{d}^c_kz^c_{t-k}\right) \\ &= \tilde{\beta}^S_i\tilde{S}_t \end{aligned}$$

Our identification assumption is that the conditional variance of S_t is 1. We have

$$\text{Var}_{t-1}\left(\beta_i^S S_t\right) = \text{Var}_{t-1}\left(\mathbb{F}_{it,common}[c_t]\right) \approx \frac{g_{i,c}^2}{g_c^2} g_c^2\left(\sigma_c^2 + \theta_c^2\right)$$

This gives us the approximate expression for the loading β_i^S

$$\beta_i^S \approx \frac{g_{i,c}}{g_c} g_c\left(\sigma_c^2 + \theta_c^2\right)^{1/2}, \tag{C22}$$

and the restriction on the loading.

$$\frac{1}{N}\sum_{i=1}^{N} \beta_i^S \approx g_c\left(\sigma_c^2 + \theta_c^2\right)^{1/2} \underbrace{\{\frac{1}{N}\sum_{i=1}^{N}\left(\frac{g_{i,c}}{g_c}\right)\}}_{\approx 1}. \tag{C23}$$

ρ_S will capture the first order autocorrelation of the process $\bar{\mathbb{F}}_{t,common}[c_t]$, which is approximately ARMA(2,1).[32] S_t is approximately the sum of two independent AR(1) processes, one with approximate autocorrelation parameter $\rho - g_c(\rho - \rho_c)$ and conditional variance $\frac{\sigma_c^2}{\sigma_c^2+\theta_c^2}$, and the other with autocorrelation parameter $(1-g_c)\rho_c$ and conditional variance $\frac{\theta_c^2}{\sigma_c^2+\theta_c^2}$. This tells us that

$$\begin{aligned}
\bar{\rho}_S &\approx \frac{\tilde{a}_1^c \tilde{a}_0^c + \tilde{d}_1^c \tilde{d}_0^c}{\left(\tilde{a}_0^c\right)^2 + \left(\tilde{d}_0^c\right)^2} \\
&\approx \frac{(\sigma_c g_c)^2\left[\rho - g_c(\rho - \rho_c)\right] + (\theta_c g_c)^2(1-g_c)\rho_c}{(\sigma_c g_c)^2 + (\theta_c g_c)^2} \\
&= \frac{\sigma_c^2\left[\rho - g_c(\rho - \rho_c)\right] + \theta_c^2(1-g_c)\rho_c}{\sigma_c^2 + \theta_c^2}
\end{aligned}$$

Analogous to the idiosyncratic component of the trend, the idiosyncratic component of the transitory component, v_{it}^c, follows an AR(1) process with persistence $(1-g_{i,c})\rho_i$ and conditional standard deviation $\omega_{i,c}$. Thus, in our statistical dynamic factor model, the persistence of the idiosyncratic component $\rho_{S,i} = (1-g_{i,c})\rho_i$ and the conditional standard deviation of the idiosyncratic component $\sigma_{S,i} = \omega_{i,c}$. If there are non-zero long-term beliefs in the forecasts of the

[32] The sum of two independent AR(1) processes is ARMA(2,1).

transitory component, b_i^c, this will be picked up by the fixed effect of the forecaster, α_i^S.

Finally, λ_i captures the decay in the slope factor in the Nelson-Siegel model when constructing the forecasts. This is directly tied to the persistence of the transitory component, so we have $\lambda_i = -\log(\rho_i)$.

In summary, the approximate mapping from the noisy information model to our statistical model is given by

$$
\begin{aligned}
\alpha_i^L &= b_i^\tau & \alpha_i^S &= b_i^c \\
\beta_i^L &= g_{i,\tau}\left(\sigma_\tau^2 + \theta_\tau^2\right)^{1/2} & \beta_i^S &= g_{i,c}\left(\sigma_c^2 + \theta_c^2\right)^{1/2} \\
\bar{\rho}_L &= 1 & \bar{\rho}_S &= \frac{\sigma_c^2\left[\rho - g_c(\rho - \rho_c)\right] + \theta_c^2(1-g_c)\rho_c}{\sigma_c^2 + \theta_c^2} \\
\rho_{L,i} &= 1 - g_{i,\tau} & \rho_{S,i} &= (1 - g_{i,c})\rho \\
\sigma_{L,i}^2 &= \omega_{i,\tau}^2 & \sigma_{S,i}^2 &= \omega_{i,c}^2 \\
\lambda_i &= -\log\rho_i
\end{aligned}
$$

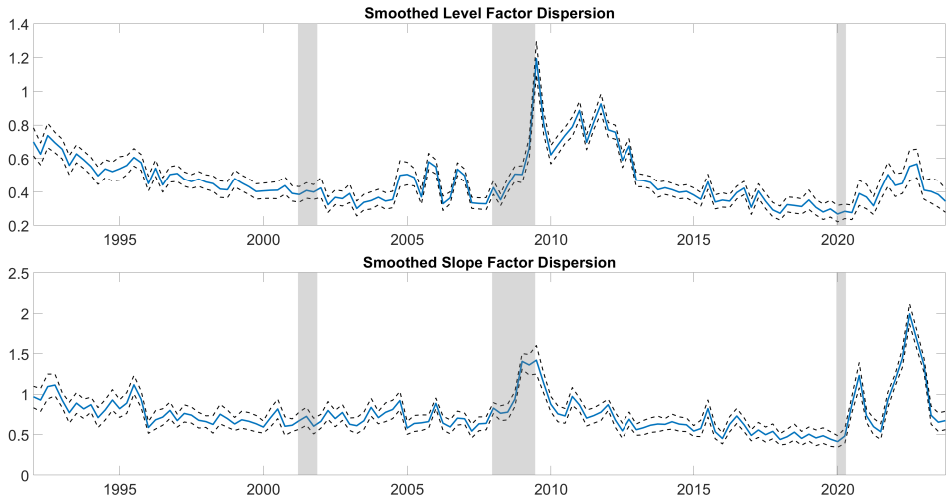
In this setting α_i^L can be interpreted as forecaster i's beliefs about long-run inflation, i.e. their perception of the inflation target, while α_i^S denotes any long-run deviation they perceive from that target.

□

Appendix D Additional Distributional Results

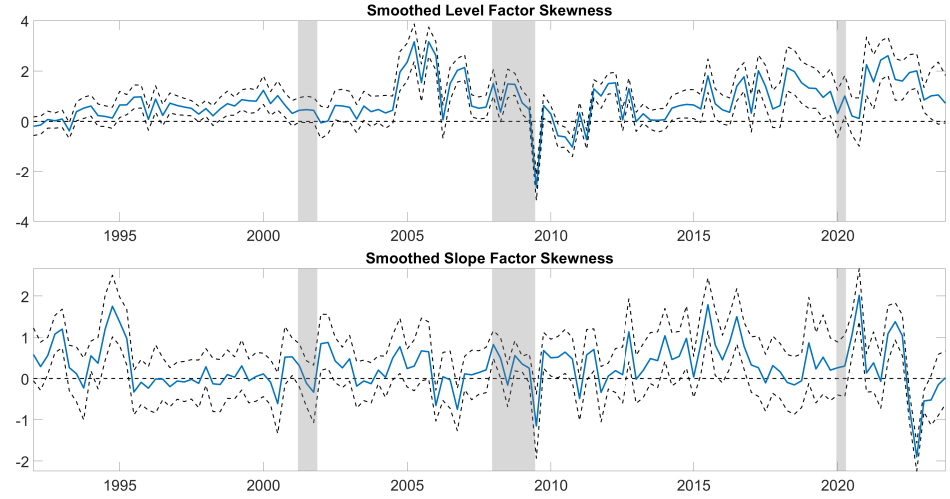
This section shows figures on moments of the factor distributions, the forecast variance decomposition, and the forecast variance share of public information, that are not included in the main text.

Figure D1: SMOOTHED IDIOSYNCRATIC FACTOR DISPERSION



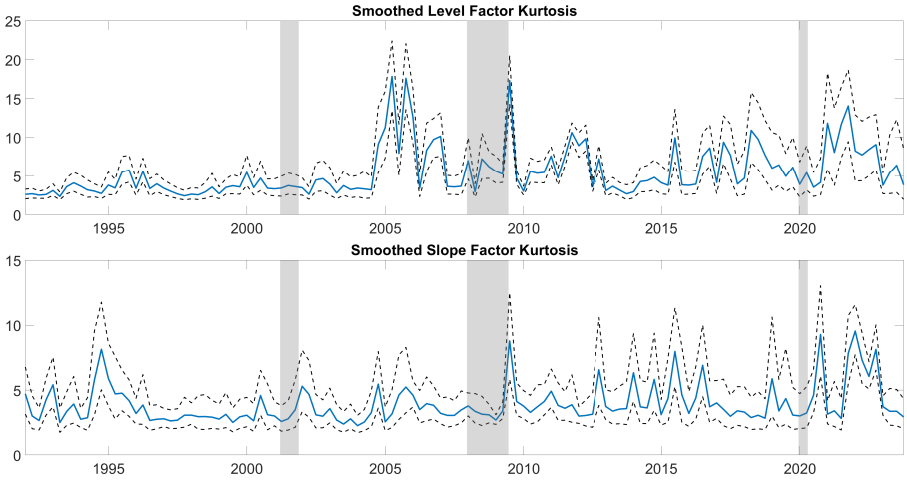
Notes: The shaded areas denote the NBER recessions.
Sources: Authors' calculation

Figure D2: SMOOTHED IDIOSYNCRATIC FACTOR SKEWNESS



Notes: The shaded areas denote the NBER recessions.
Sources: Authors' calculation

Figure D3: SMOOTHED IDIOSYNCRATIC FACTOR KURTOSIS



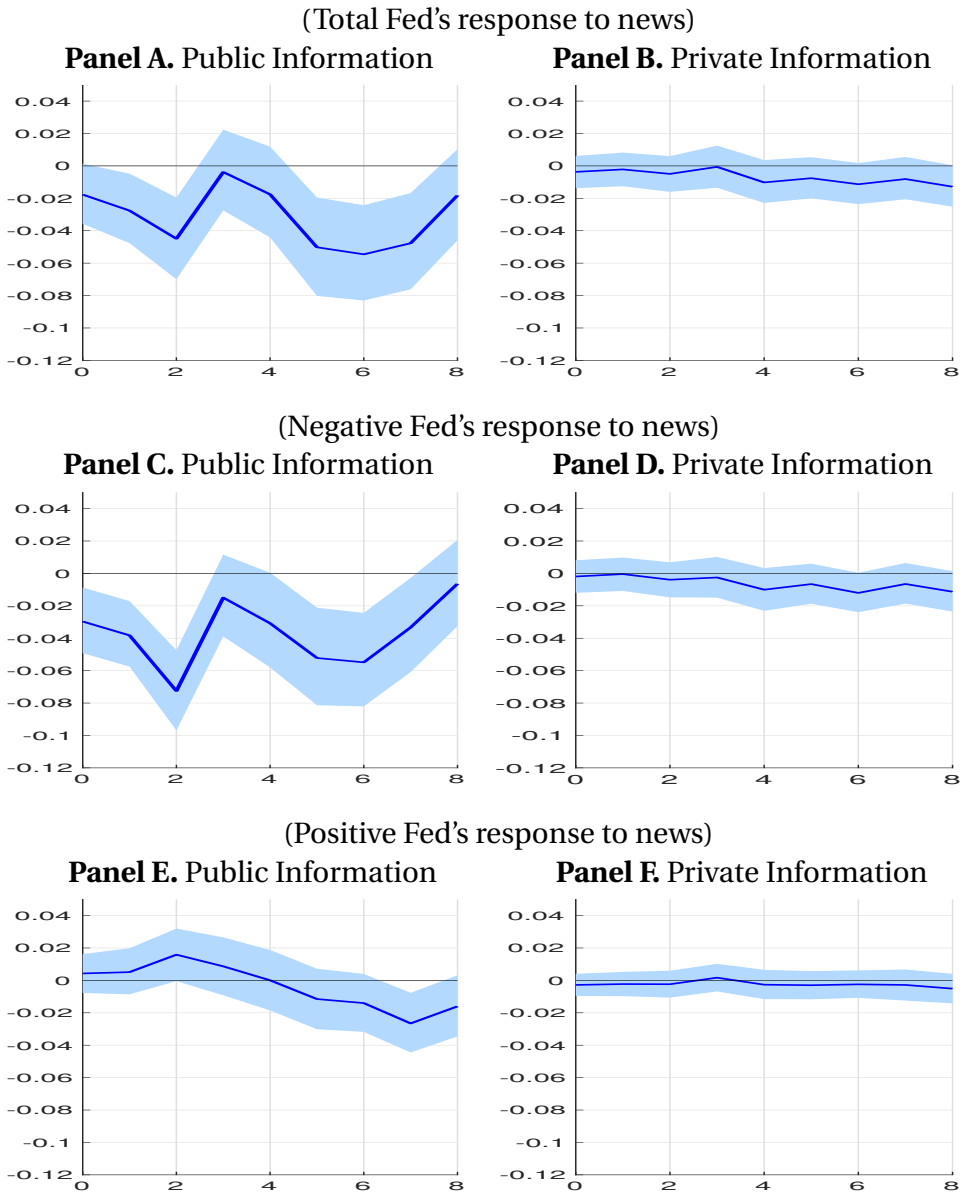
Notes: The shaded areas denote the NBER recessions.

Sources: Authors' calculation

Appendix E Additional Results on Monetary Policy

This section reports the additional estimation results about the transmission of monetary policy shocks. Figure E4 only consider private information as non-public information excluding the individual long-term beliefs. Effects of Fed's response to news are essentially zero and statistically insignificant. Figure E5 displays the transmission of Fed's reaction to news where the local projection model further includes individual-level uncertainty from Binder (2017) and the consensus expectations for the next year from the survey of professional forecasters. Figure E6 displays the transmission of orthogonalized monetary policy shocks from Bauer and Swanson (2023) where the local projection model further includes individual-level uncertainty from Binder (2017), the measure of monetary policy uncertainty from Husted et al. (2020), and the consensus expectations for the next year from the survey of professional forecasters are additionally included.

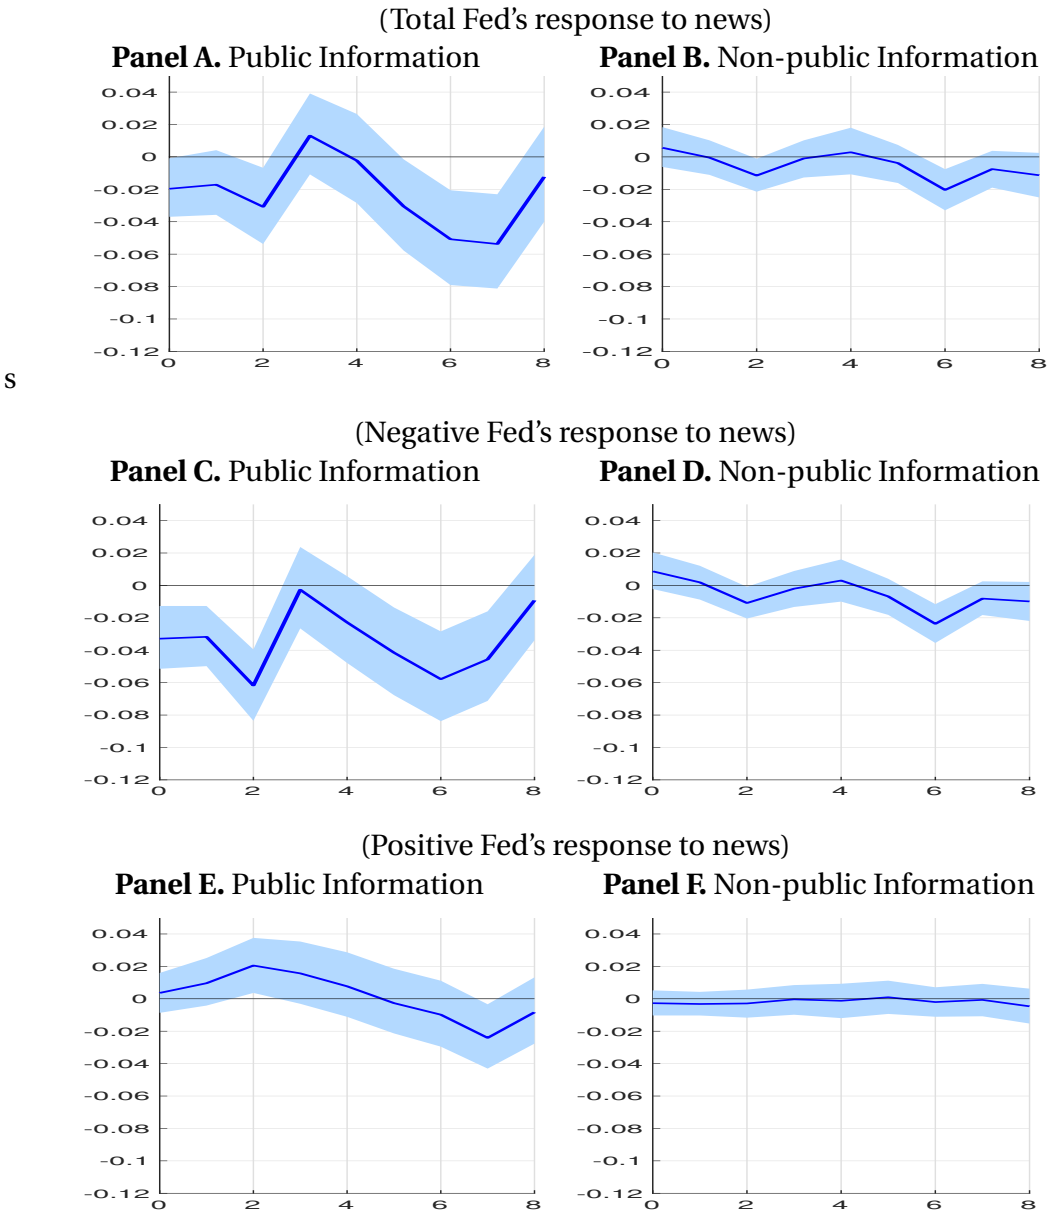
Figure E4: PROPAGATION OF FED'S RESPONSE TO NEWS (PRIVATE INFORMATION)



Notes: The figure shows the responses of disagreement about five-year five-year forward inflation expectations attributable to public information and non-public information following the Fed's response to news from Bauer and Swanson (2023), estimated from the linear model in Equation (31). HAC standard errors are reported. The shaded area indicate the 95% posterior intervals.

Source: Authors' calculation

Figure E5: PROPAGATION OF FED'S REACTIONS TO NEWS (WITH INDIVIDUAL UNCERTAINTY)

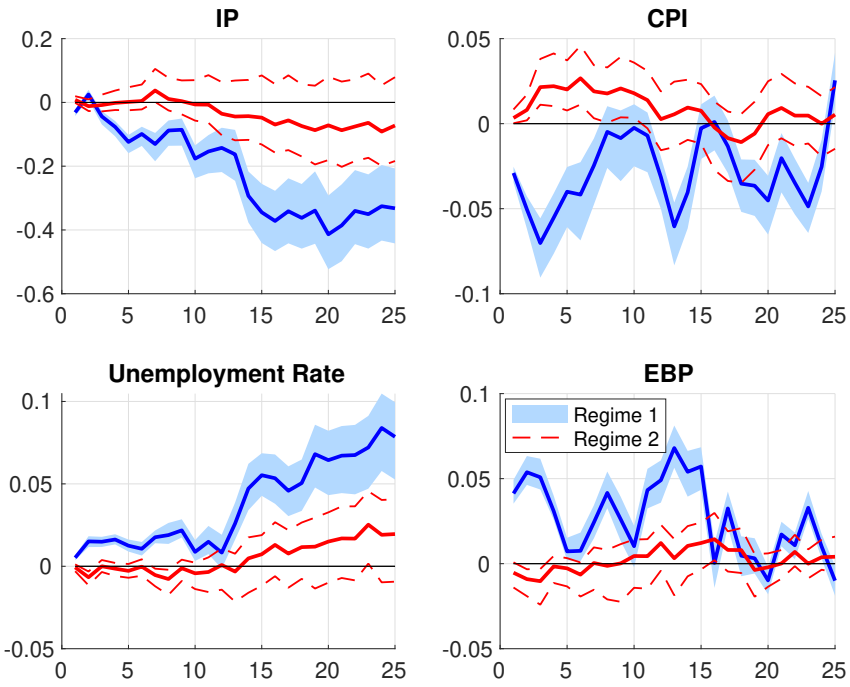


s

Notes: The figure shows the responses of disagreement about five-year five-year forward inflation expectations attributable to public information and non-public information following the Fed's response to news from Bauer and Swanson (2023), estimated from the linear model in Equation (31). The individual-level uncertainty measure from Binder (2017) is additionally included as a macro control. HAC standard errors are reported. The shaded area indicate the 95% posterior intervals.

Source: Authors' calculation

Figure E6: PROPAGATION OF MONETARY POLICY SHOCKS OF THE TWO REGIMES (WITH INDIVIDUAL UNCERTAINTY AND MONETARY POLICY UNCERTAINTY)



Notes: The figure reports the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2023). As the controls, the uncertainty measure from Binder (2017), the measure of monetary policy uncertainty from Husted et al. (2020), and the consensus expectations for the next year are additionally included. The figure shows the impulse responses of Regimes 1 and 2 scaled by 0.9 and 0.1, respectively. The blue line shows the responses when non-public information is the source of disagreement (Regime 1), while the red line represents the responses when public information is the source of disagreement (Regime 2). The upper left figure shows the cumulative response of percent changes in industrial production; the upper right figure shows the cumulative responses of percent changes in the CPI; the bottom left figure shows the responses of the unemployment rate; and the bottom right figure displays the response of the excess bond premium (EBP). The shaded areas and dashed lines represent the 95% posterior intervals.

Sources: Authors' calculation

Appendix F Discussion of Modeling Choices

This section discusses an alternative model and methodology. Section F.1 considers a time-varying parameter model. Section F.2 considers a non-parametric approach as an alternative of our parametric model.

F.1 Dynamic Factor Model

A. Fixed Factor Loadings

In our model, the factor loadings of each forecaster are fixed. A potential concern is that our model is limited in capturing the changing responsiveness of forecasters to common information. However, an individual forecaster stays in the sample for only 27 quarters on average, which is too short to allow for regime changes in the factor loadings for each individual.

Note that our model allows each forecaster to have unique loadings, although the loadings are constant for a forecaster. Therefore, in our model, two similar forecasters observed at two different points in time have different loadings, reflecting increased or decreased attention to potentially similar policy changes, for instance.

Our goal is to parse out the portion of cross-sectional variance attributed to common information. In other words, as long as the common component—the product of common factor and factor loading—is distinguished from the idiosyncrasy, the distinction of factor loading from common factor is not necessary. For instance, if the common component does not change in spite of an increase in the factor loading, the portion of disagreement driven by common information does not change and hence the increase in the factor loading does not matter.

B. Stochastic Volatility

We do not allow for time-varying variances in the dynamics of the factors. However, by allowing forecaster-specific loadings on the common factor and accounting for forecasters moving in and out of the sample, the model can indirectly capture the stochastic volatility of

aggregate inflation projections. As the composition of forecasters changes, these forecaster-specific loadings can effectively represent slow-moving stochastic volatility in the aggregate.

The resilience of the model estimates to the COVID-19 shock is a practical concern. If the pandemic observations dramatically alter the parameter estimates, the pre-pandemic inference may dramatically with the inclusion of a handful of pandemic observations. In this case, including stochastic volatility may robustify the inference, as it discounts the pandemic observations and largely prevents the model from carrying backward the COVID shock for pre-pandemic inference. To check how reliable the estimates are to the COVID shock, we compare the model estimates through 2019:Q4 with those through 2023:Q3. In particular, the estimates and the main conclusion are robust to the inclusion of the pandemic observations for the period prior to the COVID era. This observation suggests that our result is robust even in the absence of stochastic volatility. The results are available upon request.

F.2 Evidence from a Non-Parametric Model

Our baseline model is a highly parameterized model with a large number of parameters. Since the model is estimated with MCMC sampling, the estimation is costly and time-consuming. One may argue that our conclusions are sensitive to the particular parametric assumptions that we impose and that the parameter estimates may be unstable because of the size of the model.

Alternatively, we can consider a non-parametric two-step model that is estimable least squares and MLE. The first step is to construct a non-parametric model that describes the individual term structure of inflation expectations. In this model, we use Legendre polynomials to fit the short-term (less than 1 year ahead) inflation forecasts and a log function to fit the long-term (more than 1 year ahead) inflation forecasts. We fit each forecaster's observed inflation forecasts in each quarter with least squares. This results in individual level and slope factor estimates. In the second step, we estimate a dynamic factor model for each factor to parse out common and idiosyncratic components using the algorithm of Banbura and Modugno (2014). Finally, we recover the fractions of the term structure attributable to long-term beliefs, common

information, and idiosyncratic information.

Relative to the baseline model, the alternative model is less costly to estimate. In addition, we can allow for more than one common factor for the level or slope without much additional effort. However, this convenience comes at a cost. The log function is not flexible enough to capture observed forecasts beyond one year out and hence produces unrealistic long-end estimates—a noticeable decline in the long end— during the COVID-19 pandemic. This problem is not observed in our baseline model. That said, the overall conclusions about the drivers of disagreement are robust. Further details on the empirical approach and the results are available upon request.