Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Disagreement About the Term Structure of Inflation Expectations

Hie Joo Ahn and Leland E. Farmer

2024-084

Please cite this paper as:
Ahn, Hie Joo, and Leland E. Farmer (2024). "Disagreement About the Term Structure of Inflation Expectations," Finance and Economics Discussion Series 2024-084. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2024.084.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Disagreement About the Term Structure of Inflation Expectations*

Hie Joo Ahn†
Federal Reserve Board

Leland E. Farmer ‡
University of Virginia

September 16, 2024

Abstract

We develop a model of the individual term structure of inflation expectations across forecasting horizons. Using the Survey of Professional Forecasters, we decompose disagreement about inflation expectations into individuals' long-term beliefs, private information, and public information. We find that in normal times, long-horizon disagreement is predominantly driven by individuals' long-term beliefs, while short-horizon disagreement stems from private information. During economic downturns, heterogeneous reactions to public information become a key driver of disagreement at all horizons. When forecasters disagree about public information, monetary policy exhibits a delayed response and a price puzzle emerges, underscoring the importance of anchoring inflation 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, 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

"Of course, an extended period of high goods and services inflation resulting from a series of demand and supply shocks associated with the pandemic and the war could lead to a rise in inflation expectations, which would make it much more difficult to bring inflation down. That is why it has been important for monetary policy to take a risk-management posture to defend the expectations anchor. And the evidence from market- and survey-based measures suggests that longer-term inflation expectations are well anchored, while year-ahead measures have recently declined but remain elevated." — **Lael Brainard (January 19, 2023)**

1 Introduction

Households, firms, financial market participants, and policy makers disagree on future economic conditions (Cornand and Hubert, 2022). This disagreement has significant implications for the effectiveness of monetary policy and the anchoring of inflation expectation (Falck et al., 2021; Reis, 2020; Fofana and Reis, 2024). Heterogeneity in agents' expectations has also been central to macroeconomic modeling, serving as a key driver of business cycle dynamics (Lorenzoni, 2009; Angeletos and La'O, 2013; Ilut and Schneider, 2014), inflation dynamics (Michael, 2002; Mackowiak and Wiederholt, 2009), and asset pricing (Scheinkman and Xiong, 2003; Burnside et al., 2016; Barillas and Nimark, 2017). A growing body of literature on expectation formation has focused on disagreement and its time variation to explain departures from rational expectations and to explore the structural mechanisms underlying agents' expectation formation (e.g., Andrade et al., 2016; Maćkowiak et al., 2023; Fofana and Reis, 2024).

Despite the growing interest in disagreement among economic agents, no research has yet fully characterized the cross-sectional distribution of individual inflation expectations across the entire path of forecasting horizons. As Fofana and Reis (2024) note, most academic studies on expectation formation have focused primarily on short-term inflation expectations, typically with a horizon of one year or less. However, central banks closely monitor long-term expectations to assess how well expectations are anchored. Thus, a comprehensive understanding of the term structure of inflation expectations is essential. Unfortunately, expectations data are often aggregated into broad categories of forecasting horizons, leading to an incomplete picture of the term structure.

When evaluating the extent to which monetary policy anchors inflation expectations, it is crucial to identify the information sources driving disagreement over time and across forecasting horizons. If short-term expectations are primarily based on forecasters' private information or long-term beliefs, the effectiveness of monetary policy communication in reducing disagreement about near-term inflation projections may be limited. Conversely, if long-term expectations are mainly influenced by public information, effective monetary policy communication could significantly reduce long-term disagreement and aid in anchoring inflation expectations. In this context, it is important to understand the extent to which each information source shapes forecasters' projections and their disagreement across forecast horizons.

This paper makes three original contributions. First, we develop a new individual-level model of inflation expectations across forecasting horizons, which we call the *individual term-structure of inflation expectations*. This model describes a forecaster's trajectory of inflation forecasts across different horizons using two factors: *level and slope*. This approach is inspired by Nelson and Siegel's term structure of interest rates (Nelson and Siegel, 1987). The *level* reflects long-term inflation projections, which we refer to as *the long end*. The *slope* captures the overall difference between the long end and the current quarter's inflation nowcast.[1] We estimate the model using Bayesian methods, applying it to forecaster-level data from the Survey of Professional Forecasters (SPF). Although the SPF provides only a partial snapshot of a forecaster's term structure of inflation projections, our model is able to recover the cross-sectional distribution of inflation expectations across all forecasting horizons at each point in time.

The estimated short- and long-term inflation expectations exhibit different dynamics. Although the 6-month and 1-year ahead consensus forecasts closely track realized CPI inflation, the 10-year ahead consensus expectation trends downward during the 1990s and stabilizes just below 2.5% from 2000 onward. Even during the COVID-19 pandemic, the 10-year consensus shows limited variation, rising slightly before returning to its pre-pandemic level. Although the estimated consensus forecasts suggest well-anchored long-term inflation expectations, the estimated disagreement across forecasting horizons reveals a more nuanced picture. Forecasters exhibited greater disagreement about long-run inflation during the Great Recession and its recovery than in the early 1990s, when inflation was estimated to have been nonstationary. In particular, during the pandemic, both the variance and skewness of long-term expectations spiked above levels seen in the early 1990s, suggesting that expectations were weakly anchored when measured by higher moments. Our model shows that the consensus forecast and disagreement often yield distinct insights into the anchoring of agents' inflation expectations.

[1]For the sake of parsimony, we consider the curvature only in an appendix (Appendix H), as it plays a limited role in the observed individual-level term structure of inflation expectations.

Second, we develop a novel dynamic factor model where the individual term structure elements are characterized by a common component, an idiosyncratic component, and an individual fixed effect. We provide an economic interpretation to each element by building a noisy information model. In this structural model, a forecaster who has own long-run beliefs observes both public and private signals when updating inflation projections. We demonstrate that this structural model naturally aligns with the dynamic factor characterization, and captures the key empirical results of the statistical model with a simulation exercise. Thus, we interpret the common and idiosyncratic components as reflecting the contributions of public and private information to the forecaster's inflation expectations, while the individual fixed effect captures the forecaster's long-term beliefs.[2] Unlike previous studies (e.g., Herbst and Winkler, 2021), we allow for heterogeneous reactions to public information, which are captured by the individual factor loadings on the common component. To the best of our knowledge, this paper is the first to recover the distribution of sensitivity to public information across forecasters and identify the role of public and private information in disagreement based on a formal statistical model.

We observe distinct roles for the three sources of information in contributing to disagreement across forecasting horizons. Long-term beliefs and private information account for the majority of disagreement in long-run and short-run expectations respectively. We also find that forecasters exhibit heterogeneous reactions to public information, as evidenced by the non-degenerate distribution of individual-level loadings on the common component. The role of public information in disagreement is small on average.[3] However, during economic downturns and periods of high inflation uncertainty, public information becomes a key driver of disagreement. At the peak of the Great Recession, public information accounted for about half of long-term disagreement. Similarly, during the deflation at the onset of the pandemic, public information drove the majority of disagreement, and it contributed to about one-third of disagreement during the subsequent period of high inflation.

Finally, we investigate the effects of each component of disagreement on the effectiveness of monetary policy. First, we show that the Fed's responses to recent data releases (Bauer and Swanson, 2023) reduce the portion of disagreement about 2-year-ahead inflation attributable to public

[2] Throughout the paper we will use the term *long-term belief* to refer to an unconditional bias in an individual's forecast relative to rational expectations. This bias can be interpreted as arising due to some cognitive limitation / behavioral bias, or from the persistent effect of priors about long-run inflation dynamics as in Farmer et al. (2021). We remain agnostic as to the source of this bias throughout our analysis.

[3] This result is in line with previous studies such as Patton and Timmermann (2010), Farmer et al. (2021), and Lahiri and Sheng (2008).

information, but not the components attributable to private information or individual long-run means. Next, we explore how disagreement attributable to public information influences the effectiveness of monetary policy. For this analysis, we construct a new measure, the sensitivity of disagreement to public information, and use it to estimate the nonlinear effects of monetary policy shocks. We find that when public information is the main source of disagreement, the economy's responses to monetary policy shock are delayed and a price puzzle emerges. In contrast, when non-public information is the primary source of disagreement, monetary policy has rapid and statistically significant stabilizing effects on the macroeconomy.

Our empirical findings have important implications for the conduct of monetary policy. First, our approach identifies at each point in time why professional forecasters disagree about future inflation. Second, our results suggest that at important junctures (such as the Covid inflation), public information is a key source of disagreement and, thus, a potential contributor to the de-anchoring of inflation. Clear communication from the monetary authority about the macroeconomic landscape can help anchor the expectations of economic agents by reducing their disagreement about future macroeconomic conditions. Lastly, our findings offer a new perspective on the source of the price puzzle. Our results suggest that the sensitivity of disagreement to public information contributes to delivering a price puzzle, and that this price puzzle can be mitigated through effective expectation management by the monetary authority. We leave further exploration of these implications for future research.

Related Literature This paper makes contributions to several strands of the literature.

The first is the literature on the *aggregate* term structure of inflation expectations. Aruoba (2020) models the term structure of inflation expectations using a Nelson-Siegel yield curve (Nelson and Siegel, 1987). Aruoba treats the level, slope, and curvature factors as latent states and estimates them with a linear state-space model using consensus SPF and Blue Chip forecasts. Clark et al. (2022) construct the term structure of inflation expectations and uncertainty using a state space model in which stochastic volatility and persistent biases in forecasts are allowed. Diebold et al. (2008) employ a Nelson-Siegel model from Diebold and Li (2006) to predict government bond yields in the international context. To model global bond-yield dynamics, Diebold et al. (2008) consider a dynamic factor structure, similar to our modeling scheme, but different in that our main interest is forecaster-specific inflation expectations.

Recent studies focus on modeling individuals' inflation expectations (e.g., Herbst and Winkler, 2021; Crump et al., 2023; Fisher et al., 2022).[4] This paper is closest to Herbst and Winkler (2021) in that inflation forecasts across forecasting horizons are modeled and estimated at the individual level. However, there are some key differences. First, our focus is on the term structure of inflation expectations at the individual level, and hence we model the complete term structure of inflation expectations from the current quarter through ten years out with flexible Laguerre polynomials from the Nelson-Siegel model. Herbst and Winkler only model individual-level inflation expectations over horizons up to one year out, since their goal is to characterize the joint dynamics of various macroeconomic variables.[5] Second, we compute the extent to which public and private information account for changing disagreement about future inflation at each forecasting horizon, with a particular emphasis on forecasters' heterogeneous responses to public information. These heterogeneous responses are not explicitly considered in Herbst and Winkler (2021). Lastly, we provide a unique decomposition that summarizes the role of long-term beliefs, public, and private information in disagreement at each point in time.

Our research directly speaks to a vast literature on disagreement. Studies on learning models have focused on the source of disagreement. 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. Lahiri and Sheng estimate the model parameters using forecaster-level data from seven different countries provided by Consensus Economics. Second, 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. Compared to previous studies, our paper offers three distinct contributions. First, we incorporate long-term beliefs, public information, and private information into a coherent empirical framework, allowing for heterogeneous responses to public information. Second, we estimate the contribution of each information source to individual forecasters'

[4]Crump et al. (2021) estimate a multivariate trend-cycle model using the universe of professional forecasts for the U.S., and show that the multivariate model better fits the data than estimating individual univariate models on each time series. Fisher et al. (2022) model people's expectations of inflation using a trend-cycle decomposition and estimate the term structure of expectations using the full panel structure of the SPF assuming agents receive private and public signals.

[5]They adopt three common factors and an idiosyncratic component, and estimate the model with Bayesian methods.

inflation expectations using a formal statistical model with minimal a priori structure. Third, we explicitly demonstrate how each information source contributes to the disagreement in inflation expectations across various forecasting horizons and at each point in time.

Our paper also relates to recent studies on the interaction between monetary policy and disagreement of economic agents. Andrade et al. (2016) provide empirical evidence for heterogeneous beliefs about forward guidance and analyze the effect of monetary policy in the context of a new Keynesian model. Glas and Hartmann (2016) distinguish individual inflation uncertainty and disagreement between forecasters, and show that overall disagreement increases during periods of contractionary policy. Ehrmann et al. (2019) shows 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) show that a price puzzle arises in states with high disagreement but disappears in states with low disagreement using a state-dependent local projection. Dong et al. (2024) empirically show that inflation disagreement weakens the power of forward guidance and conventional monetary policy and provide a structural model where households have heterogeneous beliefs about the inflation target of central banks. Our research differs from recent studies in that we further identify the source of disagreement and explicitly show that disagreement attributable to public information is the component that delays monetary-policy effects and creates the price puzzle.

Lastly, we contribute to the discussion on how to assess the anchoring of inflation expectations. Clarida (2021) mentions that the assessment of anchored long-run inflation expectations is pivotal to outcome-based forward guidance. Bundick and Smith (2020) analyze the effect of monetary policy on inflation expectations measured from the Michigan Survey of Consumers and TIPS break-even rates, and find that inflation became more anchored after the announcement of the inflation target in 2012. In contrast, Reis (2020) develop a parsimonious structural model that can characterize discrepancies in inflation expectations between financial markets and households, and find that inflation became gradually more unanchored from 2014 onwards, which poses a trade-off in the conduct of monetary policy. Our novel empirical approach contributes to the literature in two key ways: First, it demonstrates that disagreement can serve as an independent metric for assessing anchored explanations; Second, it introduces an empirical method to quantify the extent to which communication about public information, including monetary policy announcements, can reduce inflation expectations and disagreement about future inflation.

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

2 Data: The Survey of Professional Forecasters

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

2.1 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.

2.2 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.[6] That is, we assume

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

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

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 2023Q3. 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.

2.3 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 1-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 1-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 one quarter ahead and a 10 years ahead. Short-run forecasts typically exhibit higher dispersion than long-run forecasts, with a notable exception 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.

Figure 1: SPF FORECAST SUMMARY STATISTICS AND EXAMPLES

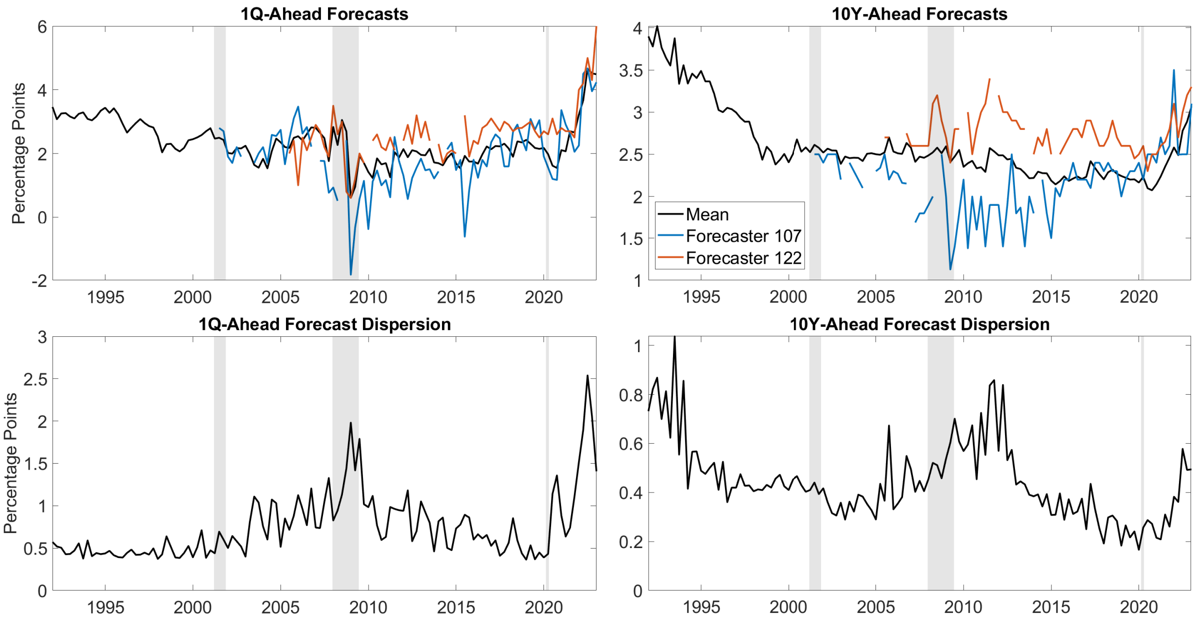

Notes: The black lines in the upper two panels show the average of inflation forecasts 1 quarter (left panel) and 10 years (right panel) ahead. The colored lines are the forecasts of two forecasters whose IDs are 107 and 122 in the survey. The bottom two panels report the cross-sectional standard deviation of the forecasts.

Sources: SPF and authors' calculation.

Figure 2: THREE TERM STRUCTURES OF OBSERVED SPF FORECASTS

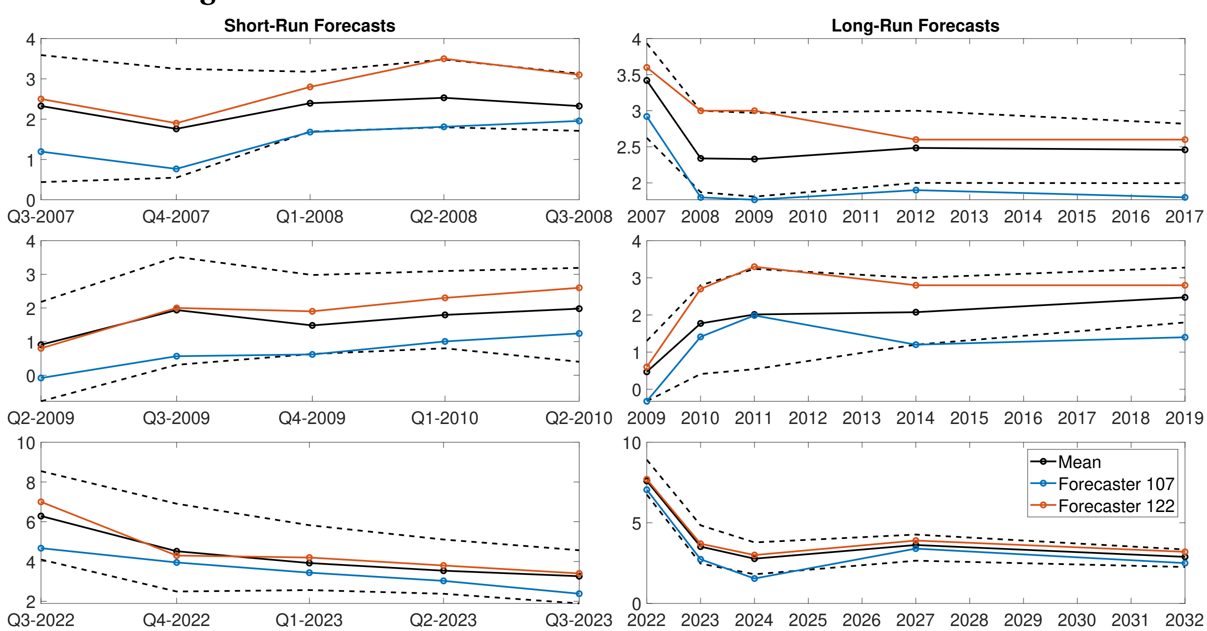

Notes: The left panel displays short-horizon forecasts, while the right panel displays long-horizon forecasts. The top panels present a snapshot of SPF forecasts for 2007Q3, the middle panels show the snapshot for 2009Q2, and the bottom panels display the snapshot for 2022Q3. The black lines capture consensus forecasts. The colored lines capture the forecasts of two forecasters whose IDs are 107 and 122 in the survey. The dashed lines are the forecasts at the 5th and 95th percentiles of the distribution.

Sources: SPF and authors' calculation.

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.

3 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.

3.1 Model

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

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

where $L_{i,t}$ and $S_{i,t}$ 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

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

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

$$L_{i,t} = \alpha_i^L + \beta_i^L L_t + \varepsilon_{i,t}^L \tag{3}$$

$$S_{i,t} = \alpha_i^S + \beta_i^S S_t + \varepsilon_{i,t}^S \tag{4}$$

where L_t and , S_t are level and slope factors which are common to all forecasters, α_i^L, and α_i^S are forecaster-specific constant terms which govern the overall level of the individual factors, β_i^L and β_i^S are forecaster-specific loadings on the common factors, and $\varepsilon_{i,t}^L$ and $\varepsilon_{i,t}^S$ capture the purely idiosyncratic components of the individual factors.[7]

[7] In this decomposition, we assume that the loadings are time-invariant. In the dynamic factor model, what is identified is the common component, which is the product of the loading and the time-varying factor. This common component is sufficient to capture forecaster i's heterogeneous reactions to common shocks. Therefore,

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 find that curvature plays a limited role in the observed individual-level trajectories of inflation expectations.[8] 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:

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

Since the scale of the common factors and the factor loadings are not separately identified, we normalize the shocks to the common factors u_t^L and u_t^S to have unit variance, and we assume the shocks are uncorrelated.[9]

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

In addition, we assume that the idiosyncratic components evolve according to AR(1) processes which are independent across forecasters.[10]

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

We also assume that the covariance matrix is diagonal, so that the factors evolve indepen-

in principle, our model does not preclude time-varying 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 explicitly estimate the time-varying loading and factor. In this paper, we follow the most standard approach—constant loadings and time-varying factors—in the dynamic factor model literature.

[8]We consider the curvature and a more flexible factor dynamics (AR(3)) in Appendix H. Our empirical results remain robust.

[9]This treatment is a standard approach in the literature of dynamic factor model (e.g., Diebold et al., 2008).

[10]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.

dently:

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

For parsimony, we make the further simplifying assumption that $\rho_{i,\varepsilon}^L, \rho_{i,\varepsilon}^S, \sigma_{i,L}^2$, and $\sigma_{i,S}^2$ are the same for all forecasters i and equal to $\rho_\varepsilon^L, \rho_\varepsilon^S, \sigma_L^2$, and σ_S^2 respectively.[11]

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.

3.2 Estimation

Our baseline model has a total of 431 parameters consisting of

- Forecaster-specific means $\{\alpha_i^L, \alpha_i^S\}_{i=1}^n$
- Forecaster-specific factor loadings $\{\beta_i^L, \beta_i^S\}_{i=1}^n$
- Factor autocorrelation parameters $\rho_L, \rho_S, \rho_\varepsilon^L$, and ρ_ε^S
- Idiosyncratic factor conditional variances σ_L^2 and σ_S^2
- Shape parameter λ

[11] See Appendix G.1 for more discussions on the modeling assumptions and their implication. The 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 over 200 additional parameters, likely increasing the uncertainty of the estimates. Therefore, our current approach strikes a good balance between flexibility and parsimony.

- Measurement error variances $\sigma^2_{v,1}, \ldots, \sigma^2_{v,20}$.[12]

The parameter vector is denoted as

$$\boldsymbol{\theta} = \left[\alpha^L_1, \ldots, \alpha^S_n, \beta^L_1, \ldots, \beta^S_n, \rho_L, \rho_S, \rho^L_\varepsilon, \rho^S_\varepsilon, \sigma^2_L, \sigma^2_S, \lambda, \sigma^2_{v,1}, \ldots, \sigma^2_{v,20}\right]'.$$

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. The model is estimated with a Gibbs sampler, detailed in Appendix B.[13]

4 Estimation Results

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

4.1 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 β^S_i, we report the median, 5th, and 95th percentiles across posterior draws of the average value across forecasters. The average value of α^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 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 α^S_i. Both the common and idiosyncratic factors are estimated to be highly persistent, with autocorrelation coefficients of between 0.73 and 0.95 at the quarterly frequency. Finally, measurement error standard deviations are estimated to be between 10 basis points (for three year forward expectations) and 67 basis points (for one quarter ahead expectations).

[12] See Appendix **??** 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 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.

[13] Alternatively, we also considered a Frequentist approach for robustness checks, which is outlined in Appendix G.2. The overall results, available upon request, remain qualitatively robust.

Table 1: POSTERIOR PARAMETER DISTRIBUTION STATISTICS

Parameter	**Median**	**95% CI**	**Parameter**	**Median**	**95% CI**
α_i^L	2.482	[2.422,2.548]	$\sigma_{\nu,6}$	0.261	[0.247,0.276]
α_i^S	0.192	[0.071,0.319]	$\sigma_{\nu,7}$	0.216	[0.202,0.232]
β_i^L	0.011	[-0.027,0.051]	$\sigma_{\nu,8}$	0.095	[0.079,0.112]
β_i^S	0.375	[0.315,0.447]	$\sigma_{\nu,9}$	0.304	[0.285,0.326]
ρ_L	0.904	[0.786,0.996]	$\sigma_{\nu,10}$	0.323	[0.303,0.346]
ρ_S	0.835	[0.685,0.979]	$\sigma_{\nu,11}$	0.297	[0.279,0.318]
ρ_ε^L	0.947	[0.928,0.964]	$\sigma_{\nu,12}$	0.331	[0.311,0.352]
ρ_ε^S	0.732	[0.696,0.767]	$\sigma_{\nu,13}$	0.185	[0.172,0.199]
σ_L	0.125	[0.117,0.134]	$\sigma_{\nu,14}$	0.197	[0.183,0.211]
σ_S	0.440	[0.421,0.458]	$\sigma_{\nu,15}$	0.237	[0.222,0.254]
λ	0.174	[0.166,0.182]	$\sigma_{\nu,16}$	0.258	[0.242,0.275]
$\sigma_{\nu,1}$	0.672	[0.656,0.688]	$\sigma_{\nu,17}$	0.193	[0.181,0.206]
$\sigma_{\nu,2}$	0.469	[0.458,0.480]	$\sigma_{\nu,18}$	0.188	[0.176,0.201]
$\sigma_{\nu,3}$	0.444	[0.434,0.454]	$\sigma_{\nu,19}$	0.181	[0.169,0.194]
$\sigma_{\nu,4}$	0.443	[0.433,0.454]	$\sigma_{\nu,20}$	0.221	[0.208,0.235]
$\sigma_{\nu,5}$	0.270	[0.256,0.285]			

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 α_i^L, α_i^S, β_i^L, and β_i^S, the table reports statistics for the average value across forecasters.

Sources: Authors' calculation

4.2 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 sharp downward trend in the 1990s. In the later part of the sample, the estimate drops sharply after the Great Recession and stays depressed for several years afterwards. During the COVID-19 pandemic, the estimate dips in the early phase of the pandemic, quickly recovers, and does not exhibit any notable changes thereafter.

The bottom panel plots the smoothed common slope factor. The estimate tracks expected changes in inflation between the current quarter and the long run at each point in time. The slope is typically positive, indicating that the term structure of inflation expectations is upward sloping on average. The steepest positive slopes occur in the years following the dotcom bubble and in the Great Recession. The slope systematically declined over the 2010s, reflecting the decade

Figure 3: SMOOTHED COMMON FACTORS

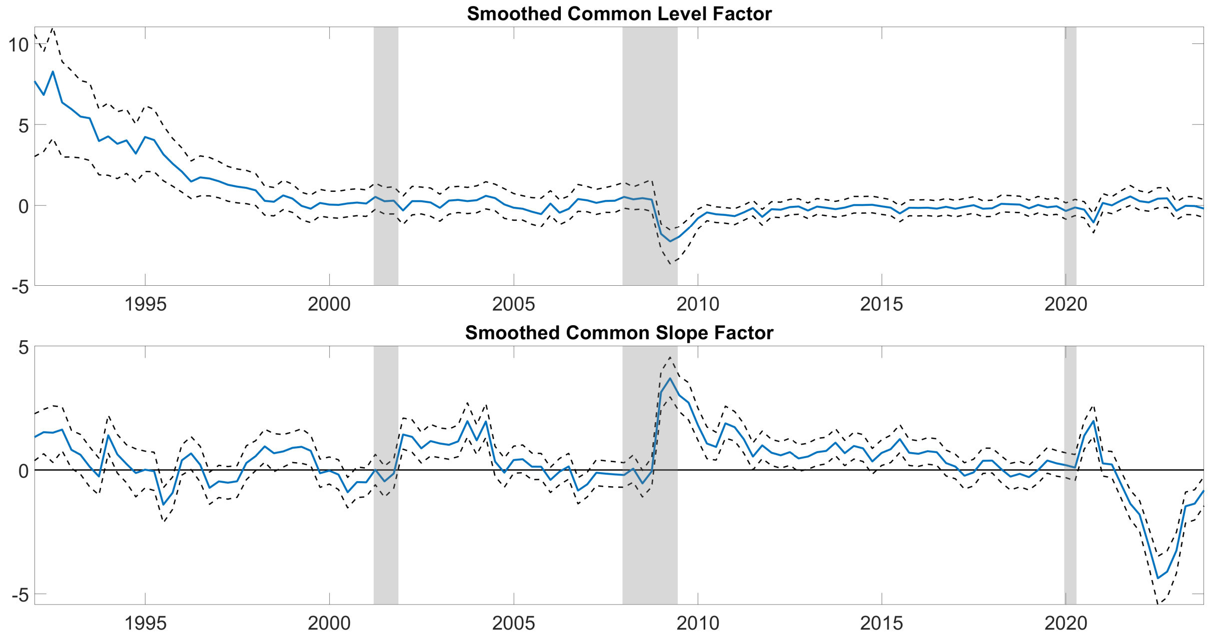

Notes: The upper panel plots the posterior median of the smoothed common level factor (blue line) along with 95 percent credible intervals (black dashed line). The bottom panel plots the posterior median of the smoothed common slope factor (blue line) along with the 95 percent credible intervals (black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

Sources: Authors' calculation

of low and stable inflation after the Great Recession. The slope increases with the onset of the COVID-19 pandemic. This reflects that forecasters expect inflation to increase relative to inflation in the current quarter, given that inflation plunged in the early phase of the pandemic and long-run expectations did not change much. From 2021 onward, when inflation picked up rapidly, the slope estimate sharply declined and stayed negative, reflecting forecasters' expectation that inflation would eventually decline.

Figure 4 plots the median and 95-percent credible interval for the mean of the forecasting distribution (across forecasters) at four different forecast horizons. The top panels plot mean 6-month and 1-year ahead inflation expectations, which track the consensus inflation nowcast and thus realized inflation to a large extent. The bottom panels display mean 5-year and 10-year ahead inflation expectations, which are significantly less variable than the short-term expectations as they incorporate less of the variation in realized inflation. The 10-year-ahead inflation consensus forecast largely tracks changes in the common level factor.

Figure 4: SMOOTHED CONSENSUS FORECASTS

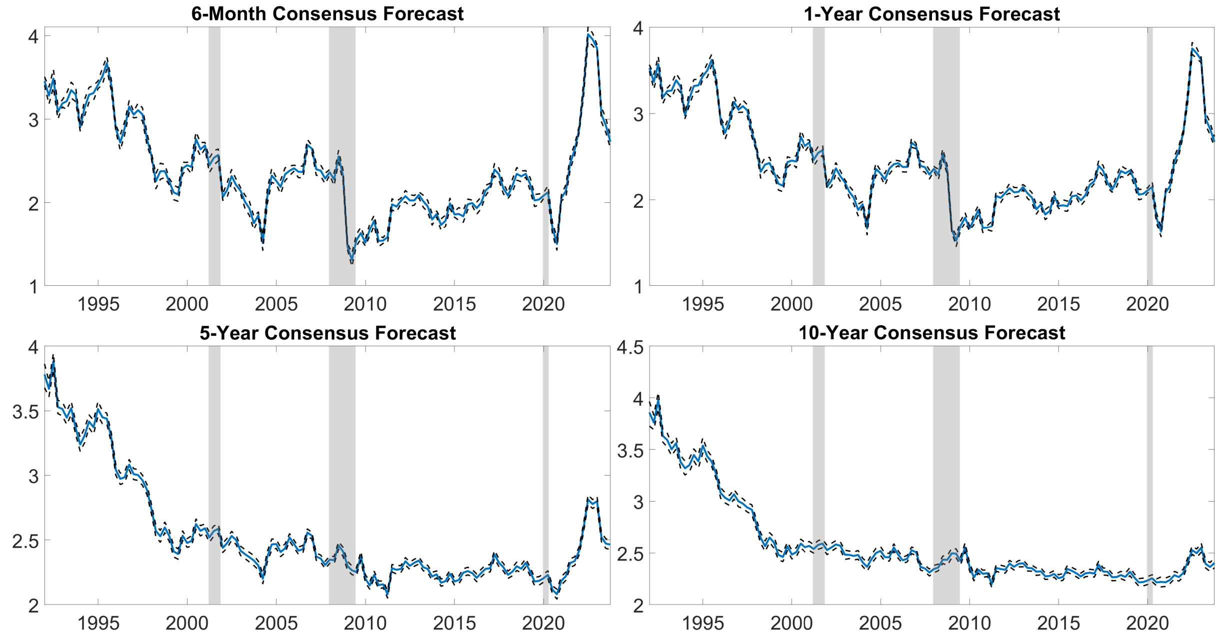

Notes: This figure plots the posterior median of the average inflation forecast posterior distribution (blue line) recovered from the individual-level term structure model along with 95-percent credible intervals (black dashed line). The four panels correspond to different forecasting horizons, 6-month, 1-year, 5-year, and 10-year respectively. The shaded areas denote NBER recessions.

Sources: Authors' calculation

4.3 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.

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 forecasters. The factors are in units of annualized percentage points. The top panel displays the distribution of individual smoothed level factors. Although consensus long-run inflation expectations remain low and stable after 2005, the dispersion of estimates changes substantially over time. In particular, the dispersion seen 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. During the COVID-19 pandemic, the median level factor edged up slightly, but the distribution dramatically skewed to the right, reflecting the perceived upside risk in long-run inflation. The bottom panel displays the distribution of individual smoothed slope factors. Similar to the level estimates, the slope dispersion increases substantially over the course

Figure 5: SMOOTHED FACTOR DISTRIBUTIONS

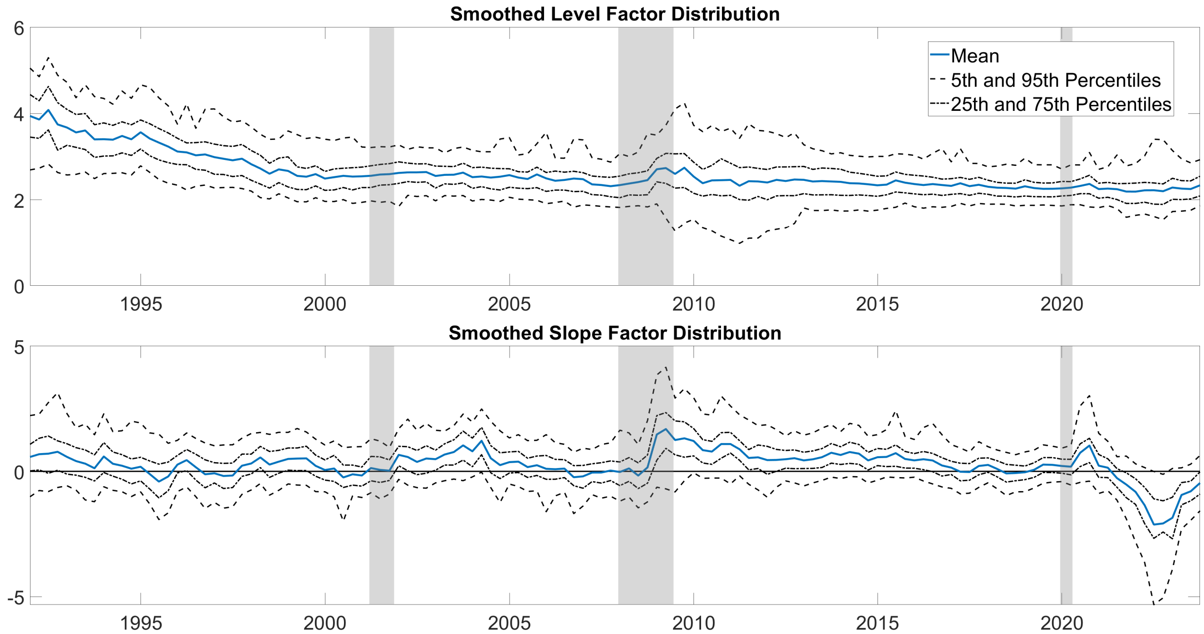

Notes: The figure shows the cross-sectional distributions of the individual level factors (upper panel) and individual slope factors (bottom panel). The solid blue line is the posterior median of the median factor across forecasters. The dashed-dotted lines depict the posterior medians of the 25th and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

Sources: Authors' calculation

of the Great Recession and the COVID-19 pandemic. In particular, the distribution becomes skewed to the right at the onset of the pandemic, but it becomes skewed dramatically to the left following a rapid rise in inflation in 2021.

Figure 6 plots the cross-sectional distribution of the constants (α_i^L and α_i^S) and the factor loadings (β_i^L and β_i^S). We plot the distribubtion across forecasters of the posterior median of each parameter. The upper panels show that forecasters exhibit significant 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 fraction anticipates it to be above 2.5% (the upper right graph). For the long-run 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. The bottom panels display the distributions of factor loadings. For both the level factor and the slope factor, the loadings show significant dispersion. Specifically, the loading for the level factor ranges from -1 to 2, while the loading for the slope factor spans from -1.5 to 1.5. The distributions of long-run means show that forecasters exhibit considerable heterogeneity in their reactions to common shocks when forming expectations about future inflation.

Figure 6: Distributions of Posterior Medians: Constants and Factor Loadings

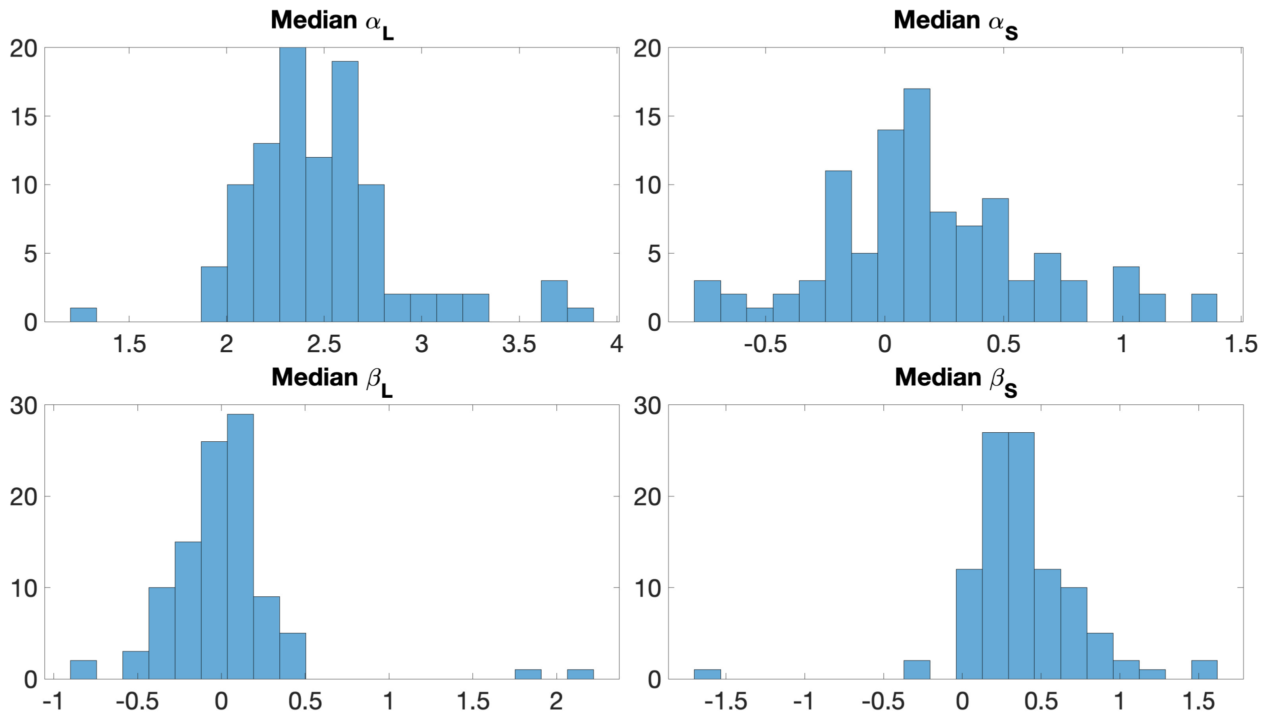

Notes: The figure shows the cross-sectional distribution of posterior median constant terms and factor loadings for both the level and slope factors.
Sources: Authors' calculation

Figure 7 displays the distribution of inflation forecasts at four different horizons: 6 months, 1 year, 5 years, and 10 years from top left to bottom right. This figure reveals stark changes in the forecast distributions that are obscured in the consensus expectations, offering insights into the anchoring of inflation expectations that differ significantly from what the consensus suggests. For example, when the mean inflation projections are low and stable, from 2005 onward, the dispersions become larger than those observed in the 1990s (when the level of inflation was higher and more volatile). It is also notable that right skewness increased significantly during the COVID-19 pandemic, indicating that forecasters had differing opinions about the upside risk of inflation, even as consensus expectations appeared relatively stable.

To ensure that our conclusions from Figure 7 are not purely driven by outliers, we examine the standard deviation (disagreement) and skewness of the forecasting distribution, displayed in Figure 8. Consistent with our observation in Figure 7, 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

Figure 7: DISTRIBUTION OF FORECASTS

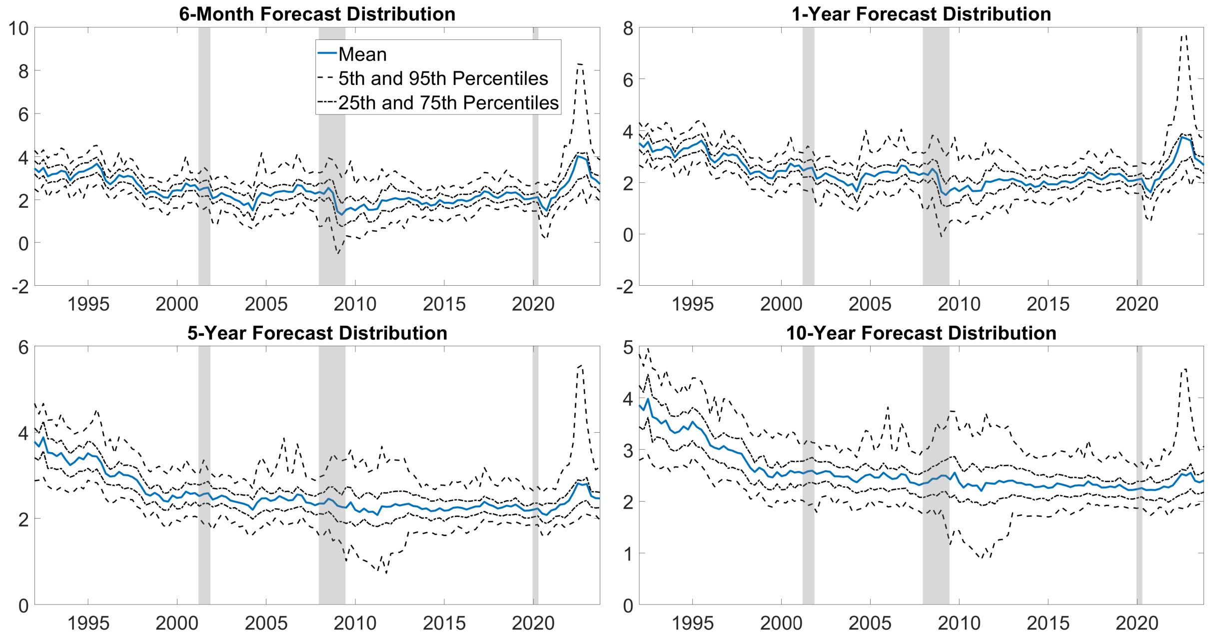

Notes: The figure shows the cross-sectional distribution of individual inflation forecasts at four different forecast horizons. The solid blue line is the posterior median of the mean forecast across forecasters. The dotted lines depict the posterior medians of the 25th and 75th percentiles. The dashed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

Sources: Authors' calculation

levels in the sample period. Both disagreement and skewness declined after 2022 but remained elevated relative to the pre-pandemic levels at the end of 2023. This observation suggests that disagreement carries information about the degree of anchoring in inflation expectations quite different from what is captured by the consensus.

4.4 Individual Dynamics

Next, we examine the individual-level term structure of inflation expectations. The purpose of this section is to show that the individual term structures of inflation expectations differ significantly from the consensus term structure, revealing that forecasters have substantively different views about the overall path of inflation in the future. This disagreement regarding the term structure of inflation expectations cannot be captured by a model based solely on the consensus term structure. Meanwhile, the individual term structures not only comove but also exhibit distinct variations across individuals, which our dynamic factor model successfully captures. It is important to note that our parsimonious individual-level model is flexible enough to accommodate the various patterns of disagreement across different forecasting horizons.

Figure 8: Disagreement About and Skewness of Forecasts

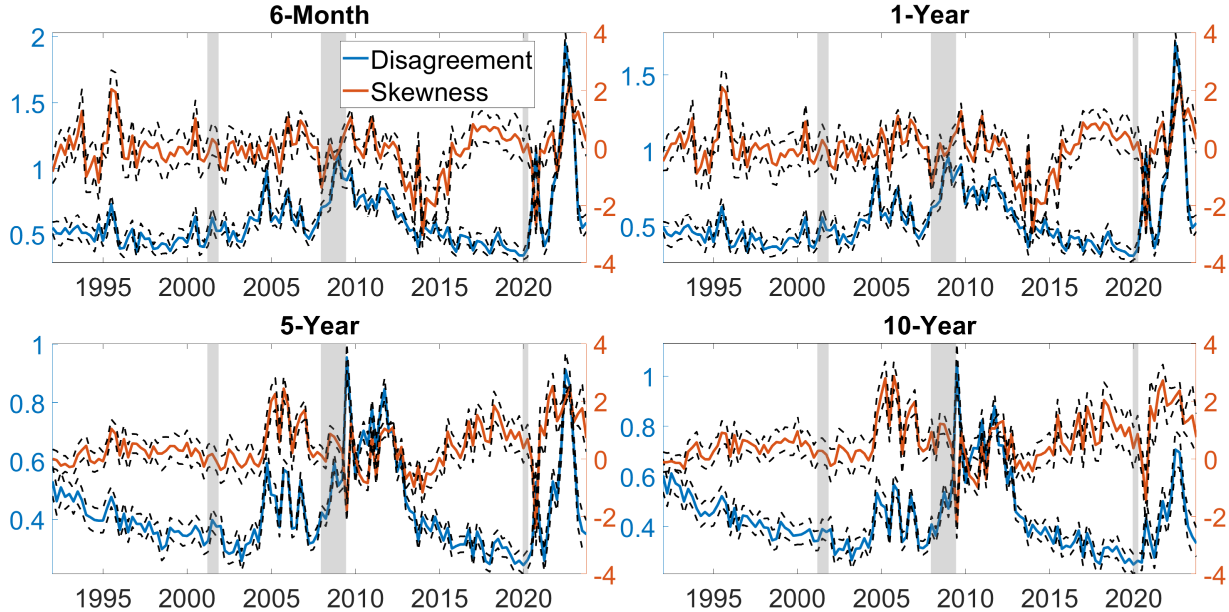

Notes: The figure shows the standard deviation and skewness of individual inflation forecasts at four different forecast horizons. The solid blue line is the posterior median of the disagreement across forecasters. The solid red line is the posterior median of the skewness across forecasters. The dashed lines depict the posterior 5th and 95th percentiles of the disagreement and skewness. The shaded areas denote NBER recessions.

Sources: Authors' calculation

Figure 9 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 downward sloping term structure of forecaster 535 (red line) and the upward sloping term structure of forecaster 518 (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 structures are upward-sloping for both individual forecasters and the consensus. Although the two forecasters significantly disagree about near-term inflation, they are still expecting a rise in inflation over the forecast horizon, in line with the consensus. 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 1.05% (middle right panel). This non-linearity signals higher

Figure 9: TERM STRUCTURE OF INFLATION EXPECTATIONS AT THREE DATES

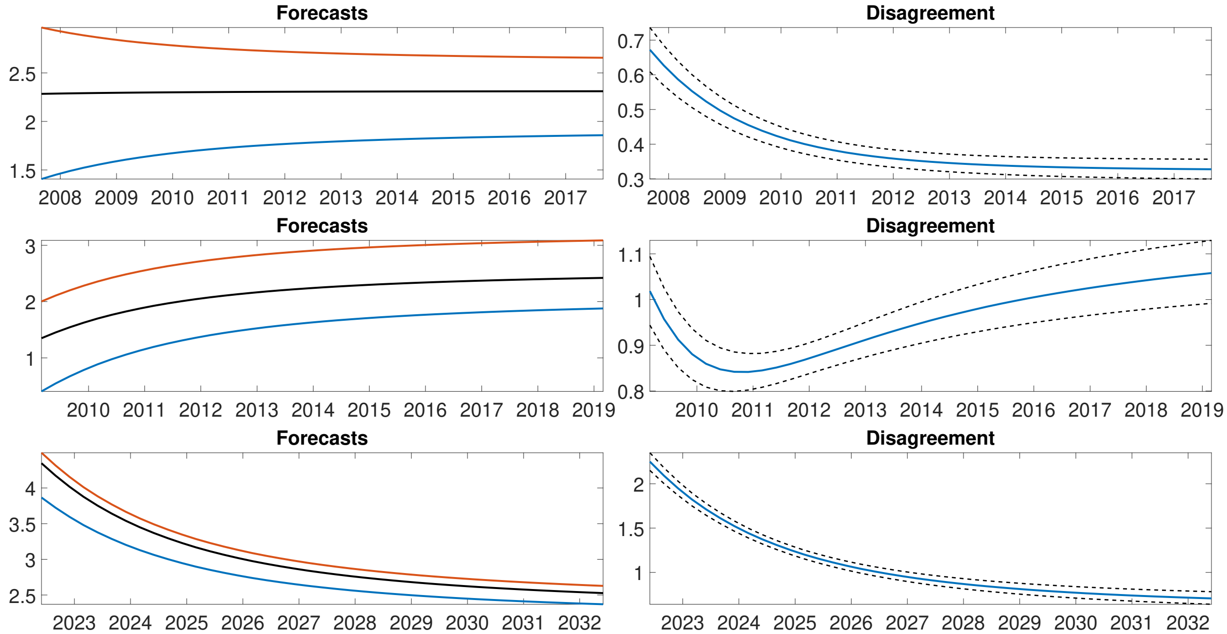

Notes: The left panels display estimated posterior medians of the term structure of inflation expectations for the consensus (black line) and two individual forecasters (forecaster 518 in blue and forecaster 535 in red). The right panels display posterior medians of the term structure of disagreement (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

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 monotonically decreasing for both forecasters and in the consensus. 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 prior to the Great Recession, the magnitude of disagreement is three times as large as that seen in 2007 in the short run.

5 A Noisy Information Model of Disagreement

This section introduces a noisy information model that provides economic interpretations of the three components of the individual-level and slope factors: the individual-specific constants, common components, and idiosyncratic components. The structural model also helps to motivate the disagreement decomposition we will present in Section 6. In the noisy information model, a forecaster with individual long-term beliefs updates inflation projections based on both public and private signals. We demonstrate that this structural model naturally aligns with the dynamic factor characterization presented in Section 3. We show that the common and idiosyncratic components of the factor model reflect the contributions of public and private information to a forecaster's inflation expectations, while the individual fixed effect captures a forecaster's long-run belief.[14]

Following Stock and Watson (2016), we describe inflation at time t (π_t) as the sum of its low-frequency component (τ_t) and transitory component (c_t). Consider the following trend-cycle state-space model for inflation:

$$\begin{aligned}
\pi_t &= \tau_t + c_t \\
\tau_t &= \tau_{t-1} + \varepsilon_t^\tau, \quad \varepsilon_t^\tau \sim N\left(0, \left(\sigma_\varepsilon^\tau\right)^2\right) \\
c_t &= \rho_c c_{t-1} + \varepsilon_t^c, \quad \varepsilon_t^c \sim N\left(0, \left(\sigma_\varepsilon^c\right)^2\right).
\end{aligned}$$

As in Beveridge and Nelson (1981), the trend component of a variable captures its long-run forecast ($\pi_{t\to t+\infty|t}$), whereas the forecast for a shorter horizon represents transitory deviations from this long-run forecast. Deviations of current inflation from its trend component are captured by the transitory component. This notion of Beveridge and Nelson decomposition naturally maps the trend to the level element of the term-structure model and the transitory component to the slope element of the model.[15]

We assume that forecasters receive signals about both the trend and cycle components separately[16], which are subject to private noise shocks $v_{i,t}^\tau$ and $v_{i,t}^c$, and public noise shocks u_t^τ and u_t^c:

$$\begin{aligned}
y_t^\tau &= \tau_t + u_t^\tau, \quad u_t^\tau \sim N\left(0, \left(\sigma_u^\tau\right)^2\right) \\
z_{i,t}^\tau &= \tau_t + v_{i,t}^\tau, \quad v_{i,t}^\tau \sim N\left(0, \left(\sigma_{i,v}^\tau\right)^2\right)
\end{aligned}$$

[14] Our structural characterization combines elements from both noisy information models (e.g., Coibion and Gorodnichenko, 2012a) and learning models (e.g., Lahiri and Sheng, 2008). The noisy information model includes both public and private information but lacks a long-run belief component. In contrast, the learning model incorporates public information and a long-run belief component but does not account for private information.

[15] The deviation of h-period-ahead inflation (for $h < 10$ years) formed at time t (denoted $\pi_{t\to t+h|t}$) from its trend ($\pi_{t\to t+\infty|t}$) is also considered transitory, with the magnitude of the deviation being captured by the slope element at time t within the term-structure model.

[16] This assumption is not necessary but serves to simplify the exposition. If only signals of inflation are observed, then the dynamics of the slope and level factors in our statistical model become a vector autoregression instead of independent univariate autoregressive processes.

$$y_t^c = c_t + u_t^c, \quad u_t^c \sim N\left(0, \left(\sigma_u^c\right)^2\right)$$
$$z_{i,t}^c = c_t + v_{i,t}^c, \quad v_{i,t}^c \sim N\left(0, \left(\sigma_{i,v}^c\right)^2\right)$$

y_t^τ and y_t^c are public signals of the trend and cycle components respectively, while $z_{i,t}^\tau$ and $z_{i,t}^c$ are private signals of the trend and cycle components respectively. We allow for the possibility that the variance of the private signal, $\sigma_{i,v}$, is different across forecasters for both the trend and cycle components.

Forecasters maintain time-invariant long-run beliefs about the trend and cycle components of inflation. For simplicity, we assume that these long-run beliefs are zero for all forecasters. This assumption can be easily relaxed to allow for individual-specific non-zero constants. Although the transitory component is modeled as a mean-zero stationary process, empirically, forecaster i's realized mean prediction of the transitory component can be non-zero. In fact, while the median of the estimated $\alpha_{i,s}$ is close to zero, the distribution of $\alpha_{i,s}$ is non-degenerate.[17] Furthermore, we characterize the trend component as a non-stationary process without a well-defined first moment. Nevertheless, we can still estimate the empirical average of each forecaster's prediction for the trend.[18] These empirical estimates of individual forecasters' mean trend and transitory components reflect their time-invariant long-run beliefs. [19]

Define the gains that forecaster i places on errors made in forecasting the public and private signals of the trend component as $g_{i,y}^\tau$ and $g_{i,z}^\tau$ respectively. We define $g_{i,y}^c$ and $g_{i,z}^c$ analogously for the cyclical component. If all agents are rational, the gains placed on forecasts of the public signals will be the same for all agents. Similarly, if the variance of private signals is the same across all agents, then the gains placed on forecasts of private signals will be the same across all agents. In our model, agents can have different forecasts either because they are rational but have heterogeneous private signal precisions or because they have different gains due to behavioral reasons or cognitive limitations.

Let $F_{i,t}$ be the expectation of agent i formed with time t information. In steady-state, agent i

[17]See the upper right panel of Figure 6.

[18]See the upper left panel of Figure 6.

[19]These long-run beliefs can be thought of as coming from some type of behavioral bias or cognitive limitation, or as capturing deviations from the truth which result from learning about the long run starting from an informative but biased prior belief (Farmer et al., 2021).

updates their beliefs according to the following equations

$$F_{i,t}\tau_t = F_{i,t-1}\tau_t + g^{\tau}_{i,y}\left(y^{\tau}_t - F_{i,t-1}\tau_t\right) + g^{\tau}_{i,z}\left(z^{\tau}_{i,t} - F_{i,t-1}\tau_t\right)$$
$$F_{i,t}c_t = F_{i,t-1}c_t + g^{c}_{i,y}\left(y^{c}_t - F_{i,t-1}c_t\right) + g^{c}_{i,z}\left(z^{c}_{i,t} - F_{i,t-1}c_t\right).$$

After some simplification the above equations can be rewritten as

$$F_{i,t}\tau_t = \left(1-g^{\tau}_i\right)F_{i,t-1}\tau_{t-1} + g^{\tau}_i\tau_t + g^{\tau}_{i,y}u^{\tau}_t + g^{\tau}_{i,z}v^{\tau}_{i,t}$$
$$F_{i,t}c_t = \left(1-g^{c}_i\right)\rho_c F_{i,t-1}c_{t-1} + g^{c}_i c_t + g^{c}_{i,y}u^{c}_t + g^{c}_{i,z}v^{c}_{i,t},$$

where $g^{\tau}_i := g^{\tau}_{i,y} + g^{\tau}_{i,z}$ and $g^{c}_i := g^{c}_{i,y} + g^{c}_{i,z}$.

We are now ready to state our main proposition:

Proposition 1. *Given the dynamic factor model characterized by equations* (3) – (8)*, there is an exact equivalence with the noisy information model for the parameters*

$$\rho_S = \rho_c \qquad\qquad \rho_L = 1$$
$$\rho^S_{i,\varepsilon} = \left(1-g^c_i\right)\rho_c \qquad\qquad \rho^L_{i,\varepsilon} = 1-g^{\tau}_i$$
$$\alpha^S_i = 0 \qquad\qquad \alpha^L_i = 0$$
$$\beta^S_i = \left(\frac{\left(g^c_i\right)^2(\sigma^c_\varepsilon)^2 + \left(g^c_{i,y}\right)^2(\sigma^c_u)^2\left(1-\rho^2_c\right)}{1+\left(\rho_c\bar{g}^c\right)^2-\rho^2_c}\right)^{1/2} \qquad\qquad \beta^L_i = \left[\left(g^{\tau}_i\right)^2(\sigma^{\tau}_\varepsilon)^2 + \left(g^{\tau}_{i,y}\right)^2(\sigma^{\tau}_u)^2\right]^{1/2}$$
$$\sigma^S_{i,v} = \left[\left(g^c_{i,z}\sigma^c_{i,v}\right)^2 - \frac{\left(\beta^S_i\right)^2\left[\left(g^c_i-\bar{g}^c\right)\rho_c\right]^2}{1-\rho^2_c}\right]^{1/2} \qquad\qquad \sigma^L_{i,v} = g^{\tau}_{i,z}\sigma^{\tau}_{i,v}$$

where $\bar{g}^c$ is the cross-sectional population average of g^c_i.

See Appendix C for the proof. Our individual-level dynamic factor model aligns with the economic interpretations of the noisy information model. Specifically, the common components L_t and S_t capture the influence of public information, while the idiosyncratic components $\epsilon^L_{i,t}$ and $\epsilon^S_{i,t}$ capture the influence of private information. Note that reactions to public information—β^S_i and β^L_i—differ across forecasters. Finally, we set the long-run mean parameters—α^S_i and α^L_i—to zero for simplicity, to clearly demonstrate the mapping between our statistical model and the standard noisy information model. As mentioned earlier, these parameters can take non-zero values without compromising the structural characterization or the mapping to the statistical model. In other words, if we assume that agents believe the permanent component has a drift and the cyclical component has a non-zero mean, these beliefs will be captured by the empirical fixed effects α_i^L and α_i^S, respectively.[20] This will introduce an additional source of disagreement, which we characterize as the 'long-term beliefs' component of disagreement.

6 Empirical Decomposition of Disagreement

In this section, we show how our statistical model decomposes individual inflation forecasts at each point in time into three distinct components: (1) individual long-term beliefs (or individual fixed effects); (2) heterogeneous responses to public information (the common component); and (3) private information (the idiosyncratic component). Section 6.1 illustrates the decomposition of individual forecasts. Section 6.2 provides a new measure for the sensitivity of disagreement to each source of information. Section 6.3 discusses the results of the decomposition. Section 6.4 provides a simulation exercise based on the noisy information model to further justify our economic interpretations of the statistical results. The simulation confirms that the empirical findings align with the theoretical predictions, even in the absence of individual long-term beliefs.

6.1 Decomposing 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.

$$L_{i,t} = \underbrace{\alpha_i^L}_{\text{long-term belief}} + \underbrace{\beta_i^L L_t}_{\text{public info.}} + \underbrace{\varepsilon_{i,t}^L}_{\text{private info.}} \tag{11}$$

The level factor of forecaster i, $L_{i,t}$, is decomposed into portions representing *long-term*

[20]In this case, the time-varying elements in the noisy information model will be demeaned values, analogous to those in the dynamic factor model.

beliefs (α_i^L), *public information* ($\beta_i^L L_t$), and *private information* ($\varepsilon_{i,t}^L$).[21] Note that α_i^L is estimated with individual fixed effects, and $\beta_i^L L_t$ and $\varepsilon_{i,t}^L$ 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 $L_{i,t}$ to be $L_i^{ltb} = \alpha_i^L$, the common component of $L_{i,t}$ to be $L_{i,t}^{pub} = \beta_i^L L_t$, and the idiosyncratic component of $L_{i,t}$ to be $L_{i,t}^{priv} = \varepsilon_{i,t}^L$. We then rewrite $L_{i,t}$ as

$$L_{i,t} = L_i^{ltb} + L_{i,t}^{pub} + L_{i,t}^{priv}. \tag{12}$$

Likewise, we give similar economic interpretations to each component of $S_{i,t}$ in Equation (4).[22] Thus, $S_{i,t}$ is rewritten as:

$$S_{i,t} = S_i^{ltb} + S_{i,t}^{pub} + S_{i,t}^{priv}. \tag{13}$$

Using the decompositions in Equations (12) and (13), forecaster i's h-quarter-ahead inflation forecast at time t, $\pi_{i,t\to t+h|t}$, is expressed as three components representing long-run beliefs ($\pi_{i,t\to t+h|t}^{ltb}$), public information ($\pi_{i,t\to t+h|t}^{pub}$), and private information ($\pi_{i,t\to t+h|t}^{priv}$):

$$\begin{aligned}
\pi_{i,t\to t+h|t} &= L_{i,t} - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right) S_{i,t} \\
&= \alpha_i^L + \beta_i^L L_t + \varepsilon_{i,t}^L - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right)\left(\alpha_i^S + \beta_i^S S_t + \varepsilon_{i,t}^S\right) \\
&= \alpha_i^L - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right)\alpha_i^S + \beta_i^L L_t - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right)\beta_i^S S_t + \varepsilon_{i,t}^L - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right)\varepsilon_{i,t}^S \\
&= \underbrace{L_i^{ltb} - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right) S_i^{ltb}}_{\text{long-term belief}} + \underbrace{L_{i,t}^{pub} - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right) S_{i,t}^{pub}}_{\text{public information}} + \underbrace{L_{i,t}^{priv} - \left(\frac{1-e^{-\lambda h}}{\lambda h}\right) S_{i,t}^{priv}}_{\text{private information}} \\
&= \pi_{i,t\to t+h|t}^{ltb} + \pi_{i,t\to t+h|t}^{pub} + \pi_{i,t\to t+h|t}^{priv}
\end{aligned}$$

The factor structure leads us to decompose the cross-sectional dispersion of $\pi_{i,t\to t+h|t}$ into the dispersion components driven by the three information sources.

$$\mathrm{Var}_i(\pi_{i,t\to t+h|t}) \approx \mathrm{Var}_i(\pi_{i,t\to t+h|t}^{ltb}) + \mathrm{Var}_i(\pi_{i,t\to t+h|t}^{pub}) + \mathrm{Var}_i(\pi_{i,t\to t+h|t}^{priv}) \tag{14}$$

[21] We use public (private) information and common (idiosyncratic) information interchangeably, but will use public (private) information in this section given the economic interpretation from the noisy information model.

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

Note that the variance of individual long-term beliefs (the first term in (14)) will not create changes in the dispersion, if the pool of forecasters does not change over time. Since each forecaster has different average forecasts, the compositional change in forecasters over time creates variation in long-term belief dispersion. In this sense, dispersion due to long-term beliefs is interpreted as the natural level of disagreement or fundamental disagreement (Andrade et al., 2016).

This variance decomposition is possible because innovations to the common and idiosyncratic components are assumed to be independent of each other in the dynamic factor model. However, realized shocks to the common and idiosyncratic components can exhibit finite sample comovement. This comovement creates a wedge between the two sides of (14). Nonetheless, the finite-sample comovement will be close to zero in large samples.

6.2 Disagreement Shares

This section provides an alternative decomposition of disagreement about inflation h periods ahead into the components of three information sources—what we call *the disagreement shares*. Our goal is to assess the extent to which each information source influences the overall cross-sectional variance in forecasts of inflation h periods ahead at each point in time. This new measure avoids the caveat of (14) that obscures the contribution of each information source in the disagreement.

Conceptually, our disagreement shares resemble the idea of *beta* in finance, which gauges the sensitivity of stocks to a common factor. As an alternative to simply looking at the cross-sectional variance of each component, we construct the information share β for disagreement using an approach proposed in Fujita and Ramey (2009).[23] This new measure overcomes the issue that finite-sample comovements among components cause the individual variance shares not add up to one, which is a well-known issue in the literature on dynamic factor models (e.g., Ahn and Luciani, 2024). Fujita and Ramey (2009) show when a variable is expressed as the sum of different sub-components, the ratio of the covariance between the sum and each component divided by the variance of the sum — the beta — add up to one. Note that the h-period ahead inflation forecast of individual i at time t is composed of the portion accounted for by long-term

[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 unemployment rate, a similar problem in the literature of unemployment dynamics.

beliefs, public information, and private information. Therefore, we can apply Fujita and Ramey's decomposition to our case.

Restating our decomposition of inflation projections we have that:

$$\pi_{i,t\to t+h|t} = \pi^{ltb}_{i,t\to t+h|t} + \pi^{pub}_{i,t\to t+h|t} + \pi^{priv}_{i,t\to t+h|t}. \tag{15}$$

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

$$1 = \beta^{ltb}_{h,t} + \beta^{pub}_{h,t} + \beta^{priv}_{h,t}, \tag{16}$$

where

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

Note that $\beta^{ltb}_{h,t}$, $\beta^{pub}_{h,t}$, and $\beta^{priv}_{h,t}$ can technically be negative or go above 1.

6.3 Decomposition Results

Figure 10 presents the results of estimating Equation (16) for each time period and forecasting horizon. For clarity, we will focus on discussing the results for the 1-year and 10-year ahead forecast horizons.[24] The upper panels show the level of disagreement about 1-year and 10-year ahead inflation: total (black line), the portion attributable to long-term beliefs (blue line), the portion attributable to public information (orange line), and the portion attributable to private information (yellow line). The bottom panels display the disagreement shares of public information for 1-year and 10-year ahead inflation projections.

We make a few key observations. First, private information is the primary source of short-run (1-year ahead) disagreement, explaining approximately 60 percent of the dispersion. Second,

[24] Results for additional forecasting horizons can be found in Section Appendix D of the appendix.

individual long-term beliefs are the main source of disagreement for long-run inflation forecasts, accounting for about 55 percent of the dispersion. Finally, public information contributes the smallest share of disagreement across all forecasting horizons, making up around 10 percent for both short-run and long-run forecasts. However, public information plays a significantly larger role in explaining disagreement during three major periods: the early and mid-1990s, the Great Recession, and the COVID-19 pandemic. These periods correspond to economic recessions or episodes of heightened inflation uncertainty.

Figure 10: FORECAST VARIANCE DECOMPOSITION

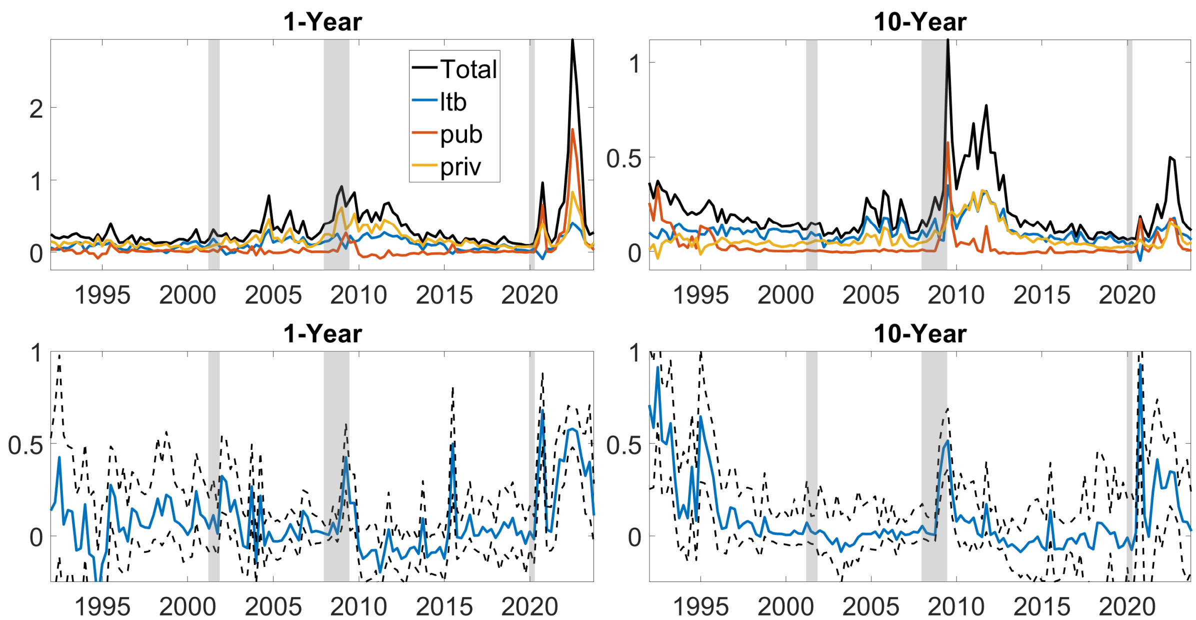

Notes: The top two panels show the decomposition of the cross-sectional variance of inflation forecasts (black line) into the components driven by individual long-term beliefs (denoted by ltb, blue line), heterogeneous responses to public information (denoted by pub, red line), and private information (denoted by priv, yellow line). Each line corresponds to the posterior median. The bottom panels show the variance share of public information, $\beta_{h,t}^{pub}$. The solid blue line corresponds to the posterior median and the dotted black lines correspond to pointwise 95% credible intervals. The left and right columns correspond to 1- and 10-year forecasting horizons respectively. The shaded areas denote NBER recessions. The bottom panels show the variance share of public information.

Sources: Authors' calculation

To better understand this time-varying role of public information as a driver of disagreement, we examine $\beta_{h,t}^{pub}$ in the bottom panels of Figure 10 more closely.[25] The contribution of public information to disagreement changes dramatically over time. 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. For

[25]Section Appendix D of the appendix provides the disagreement shares of public information for additional forecasting horizons.

Figure 11: TERM STRUCTURE OF VARIANCE SHARES AT THREE DATES

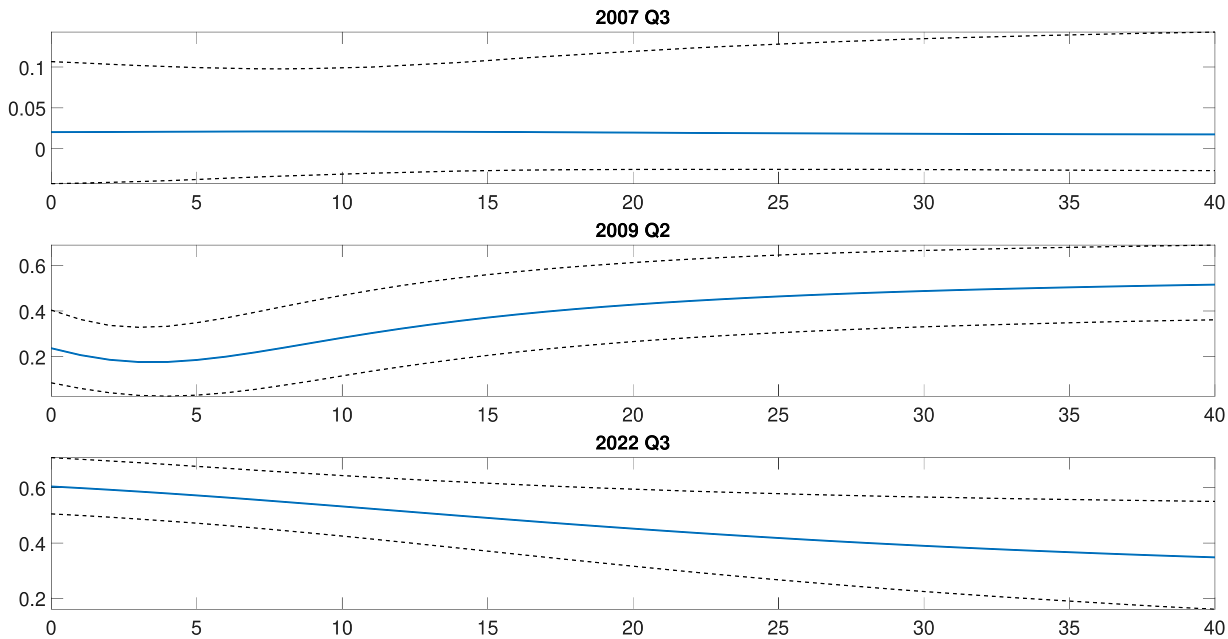

Notes: The figure shows the fraction of overall disagreement about inflation expectations driven by heterogeneous responses to common information over a ten-year forecast horizon at three different points in time as measured by the β measure proposed in Fujita and Ramey (2009). The top panel reports 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

example, it spikes to approximately 65 percent for short-horizon forecasts and 90 percent for long-horizon forecasts during the pandemic. This observation suggests that forecasters pay more attention to public information during economic downturns and periods of higher inflation uncertainty, but translate the public information into their forecasts in different ways.

Furthermore, the term structure of common disagreement varies dramatically over time as well. Figure 11 shows the share of public information in disagreement across forecasting horizons for three particular time periods. In 2007Q3, just before the Great Recession, the role of common 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 50% compared to 20%. The COVID-19 pandemic is unique in that public information was the primary driver of increased disagreement across all forecasting horizons (bottom panel). The share is greater than 30 percent across all forecasting horizons, declining from about 60% in the short run to just below 35% in the long run. In the presence of shocks that are unprecedented in nature and magnitude and the consequent policy response, forecasters paid attention to public information during the pandemic. However, due to high economic uncertainty, forecasters had very different interpretations about the same public information and produced quite different inflation expectations across forecasting horizons.

The importance of public information in long-run disagreement in times of high economic uncertainty suggests that monetary policy may be able to anchor long-horizon expectations effectively with clearer communication.

6.4 Relation to Noisy Information Model

Next, we conduct a simulation exercise based on the noisy information model described in Section 5, and show that our empirical disagreement decomposition results are consistent with our theoretical model's predictions.

Our model allows forecasters to respond heterogeneously to the same public news. These heterogeneous responses drive the significant increase in disagreement attributable to public information during periods of large shocks, even though we do not account for individual-level uncertainty. Moreover, by allowing for varying sensitivity to public information across different forecasting horizons for each individual, our model captures the distinct roles of public information in driving disagreement in the short run versus the long run. In contrast, previous literature on the noisy information model has typically assumed uniform reactions to public information across individuals and has not accounted for differing reactions across forecasting horizons.[26]

Using simulated data, we demonstrate how the noisy information model presented in Section 5 can account for the stylized facts on time-varying disagreement and time-varying shares of disagreement. We assume that $g_i^c = g$ and $\sigma_{i,v}^c = \sigma_v$ for all agents i, and we assume that the trend component τ_t is known to be constant over time and equal to 0. Thus, all time-variation in inflation comes from the cyclical component c_t. These simplifications are for illustrative purposes only, to help isolate the time variation in the relative importance public vs private information. By allowing for a non-zero trend component, we can generate variations in the source of disagreement across forecasting horizons over time. Similarly, we abstract from long-term beliefs, because they do not contribute to time-variation in disagreement unless we also

[26] A few previous studies have considered a noisy information model where agents update their information at different frequencies, resulting in limited dispersion of forecasts in response to public information. A detailed literature review on this topic is provided in Section Appendix F of the appendix.

model forecaster turnover, as is observed in the SPF.

Given the above simplifications, beliefs of agent i about inflation follow the stochastic process

$$F_{i,t}\pi_t = \left(1-g\right)\rho_c F_{i,t-1}\pi_{t-1} + g\pi_t + g^c_{i,y}u^c_t + g^c_{i,z}v^c_{i,t}$$

Let g^c_y denote the cross-sectional population average of $g^c_{i,y}$. Furthermore, let $\sigma^2_{g,y}$ and $\sigma^2_{g,z}$ denote the cross-sectional population variances of $g^c_{i,y}$ and $g^c_{i,z}$ respectively. The cross-sectional mean of beliefs follows

$$\bar{F}_{i,t}\pi_t = \left(1-g\right)\rho_c \bar{F}_{i,t-1}\pi_{t-1} + g\pi_t + \bar{g}^c_y u^c_t \tag{17}$$

and the cross-sectional variance follows

$$\mathrm{Var}_i\left(F_{i,t}\pi_t\right) = \left[\left(1-g\right)\rho_c\right]^2 \mathrm{Var}_i\left(F_{i,t-1}\pi_{t-1}\right) + \sigma^2_{g,y}\left(u^c_t\right)^2 + \sigma^2_{g,z}\sigma^2_v \tag{18}$$

Equation (18) implies that time-variation in disagreement is driven entirely by public noise shocks u^c_t. The magnitude of the impact on cross-sectional disagreement is driven by the magnitude of the shock and the variance of the gains $g^c_{i,y}$, $\sigma^2_{g,y}$, across agents. Larger shocks (in magnitude) and a larger dispersion of gains on the public signal lead to larger and more persistent swings in disagreement.

To illustrate the connection between the noisy information model and our statistical model, we consider a simulation experiment. We parameterize the process for beliefs using $\rho_c = 0.95$, $\sigma^c_\varepsilon = 0.83$, $g = 0.4$, $\sigma_u = 2.1$, $\sigma_v = 2.1$, $\bar{g}^c_y = 0.2$, and $\sigma_{g,y} = 0.1$. The volatility of shocks to the cyclical component, σ^c_ε, is chosen so that the unconditional variance of c_t is the same as the unconditional variance of inflation over our sample period. The variance of the public and private noise shocks are both chosen to be equal to twice the variance of the cyclical component.[27] The overall gain and average public signal gain g and $\bar{g}^c_y$ are set to be the gains that would be chosen by a rational agent facing the specified signal extraction problem. The $g^c_{i,y}$ are randomly drawn from a Normal distribution with mean $\bar{g}^c_y$ and standard deviation $\sigma_{g,y}$. The $g^c_{i,z}$ are set equal to $g - g^c_{i,y}$ for each i.

We simulate a balanced panel of forty forecasters for three-hundred time periods. We then construct common and idiosyncratic components by decomposing the simulated nowcasts using our statistical dynamic factor model in Section 3 assuming no measurement error. The

[27]This implies a signal to noise ratio of one third for both the public and private signals. This is chosen to approximately match the observed cross-sectional dispersion in inflation nowcasts from the SPF.

Figure 12: NOISY INFORMATION MODEL SIMULATIONS

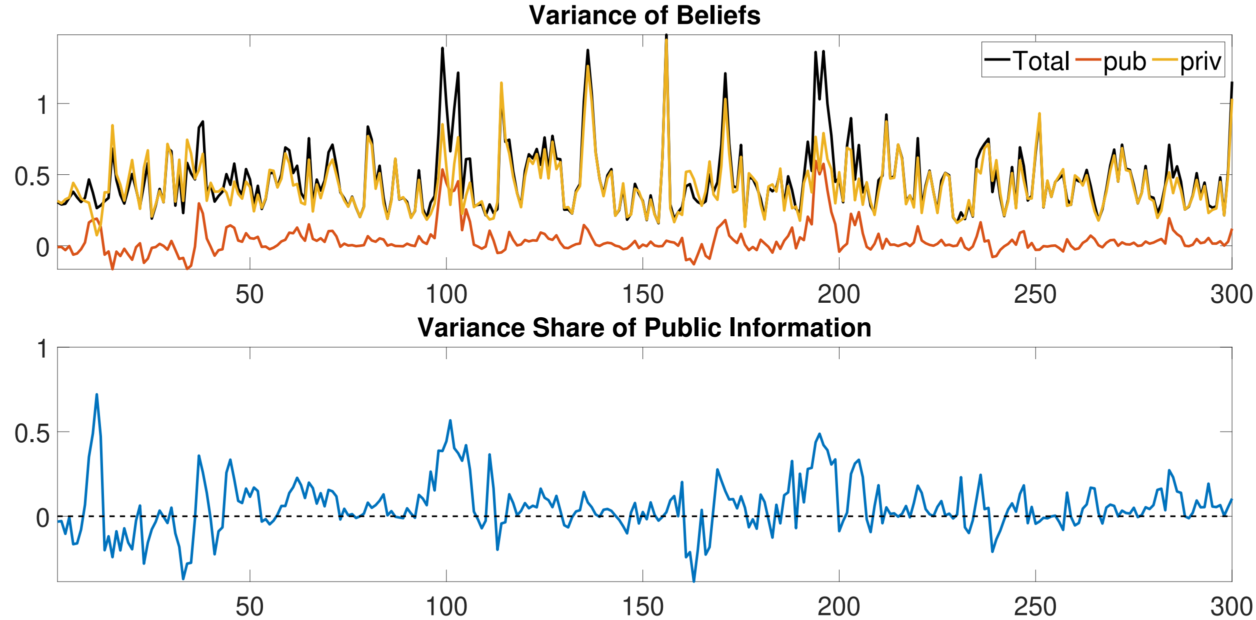

Notes: The figure shows the dynamics of disagreement and disagreement shares in a simulation experiment from our noisy information model. The top panel plots the overall cross-sectional variance of beliefs about inflation in black, the variance attributable to public information in red, and the variance attributable to private information in yellow. The bottom panel plots the share of the total variance attributable to public information.

Sources: Authors' calculation

disagreement and variance share measures are constructed exactly as in Section 6.

Figure 12 presents the results for disagreement. The top panel plots the total variance of beliefs in black along with the portions of variance driven by public and private information in red and yellow respectively. The bottom panel plots the share of the variance driven by public information. These panels broadly match the patterns observed in Figure 7. Disagreement exhibits conditional heteroskedasticity corresponding to periods of large public noise shocks. In periods of large public noise shocks, the share of disagreement explained by public information rises. There are also spikes in disagreement which arise from a large realized dispersion of private signals, and in these periods there is no corresponding rise in the variance share of public information, as predicted by our theory.

The structural model effectively captures our key empirical finding—the increased importance of public information in driving disagreement during periods of large inflationary shocks. It is important to note that this result was achieved even without allowing for changes in individual-level uncertainty. This observation further validates the economic interpretations of the three elements of the dynamic factor model as long-term beliefs, heterogeneous reactions to public information, and private information.

7 Implications for Monetary Policy

This section explores the link between the source of disagreement and the effectiveness of monetary policy. Section 7.1 examines whether the news component of monetary policy surprises reduces the disagreement attributable to public information and discusses implications for anchoring inflation expectations. Section 7.2 investigates the effect of disagreement about public information on the transmission of monetary policy shocks.

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

So far, we have implicitly assumed that disagreement about public information can be mitigated through monetary policy communication. 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 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. If forecasters pay attention to the Fed's reactions to data releases, this news component should reduce the disagreement about public information among forecasters.

Let y^p_{t+h} denote the disagreement about 8-quarter-ahead inflation attributable to public information. Similarly, let y^o_{t+h} denote the disagreement about 8-quarter-ahead inflation attributable to non-public information, which includes both private information and long-term beliefs. We measure disagreement using the standard deviation of the portion of individual-level forecasts attributable to each information source. We focus on an 8-quarter horizon to account for policy lags.

The local projection model is specified as follows:

$$y^j_{t+h} = \alpha^j_h + \beta^j_h z_t + \Gamma^j_h \mathbf{X}_{t-1} + e^j_{t+h} \qquad \text{for} \quad j \in [p, o] \quad h = 0, 1, \cdots, H, \tag{19}$$

where α^j_h is a constant, z_t is the news-component shock, β^j_h captures the magnitude of pass-trough of the shock h quarters after impact, and e^j_{t+h} is an error term. The notation $\mathbf{X}_{t-1}$ denotes a set of macroeconomic controls, 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 also include two lags of z_t in the controls.[28] Note that the parameter draws from the Gibbs sampler are used to construct y^p_{t+h} and y^o_{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.[29]

Figure 13 displays the responses of two disagreement measures to the news component of a monetary policy shock. We report the effects of a one-standard-deviation innovation in the news component. As shown in Panel A, the Fed's positive reaction to economic news immediately and significantly reduces disagreement about 8-quarter-ahead inflation attributable to public information. However, the same shock has a smaller and statistically insignificant effect on the corresponding disagreement attributable to non-public information. This result confirms that disagreement about public information is the portion that can be reduced by monetary policy communication and is therefore relevant for anchoring inflation expectations.

We further investigate whether the effects of the Fed's reactions to news change when forecasters have recently experienced substantial disagreement about public information. To explore this possibility, 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 6, with a focus on disagreement regarding 8-quarter-ahead inflation. Consider the following nonlinear local projection model.

$$y^j_{t+h} = \alpha^j_h + \beta^j_{1,h}(1 - s_{t-1})z_t + \beta^j_{2,h} s_{t-1} z_t + \Gamma^j_h \mathbf{X}_{t-1} + e^j_{t+h} \quad \text{for} \quad j \in [p, o], \quad h = 0, 1, \cdots, H \tag{20}$$

In this model, s_{t-1} is the indicator of Regime 2 and is measured using $\beta^c_{8,t-1}$ from Equation (16). The indicator of Regime 1 is thus $(1 - s_{t-1})$, measured with $(1 - \beta^c_{8,t-1})$. The parameter $\beta^j_{1,h}$ captures the magnitude of pass-through when disagreement is driven by non-public information (Regime 1), while $\beta^j_{2,h}$ reflects the magnitude when public information is the source of disagreement (Regime 2). We use a one-quarter lag of s_t to avoid contemporaneous feedback from policy actions (Auerbach and Gorodnichenko, 2013). The parameter draws from the Gibbs sampler are

[28] We use the controls of macroeconomic variables as considered by Bauer and Swanson (2022).

[29] 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 13: Propagation of Fed's reactions to news

Panel A. Average Responses

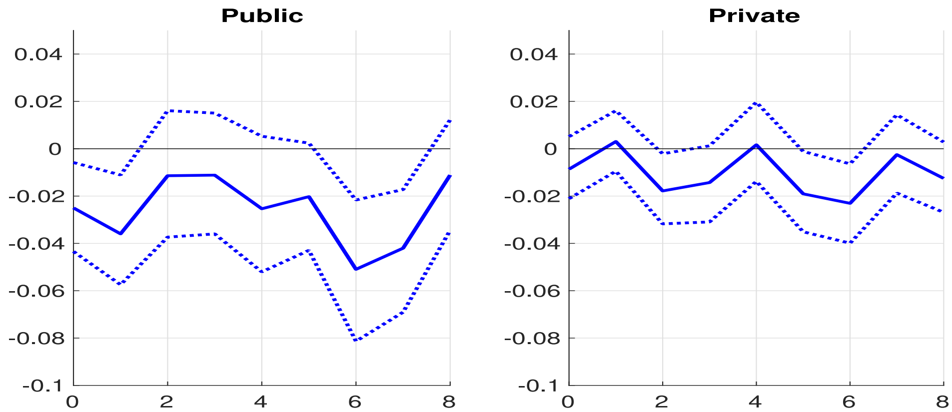

Panel B. Nonlinear Responses

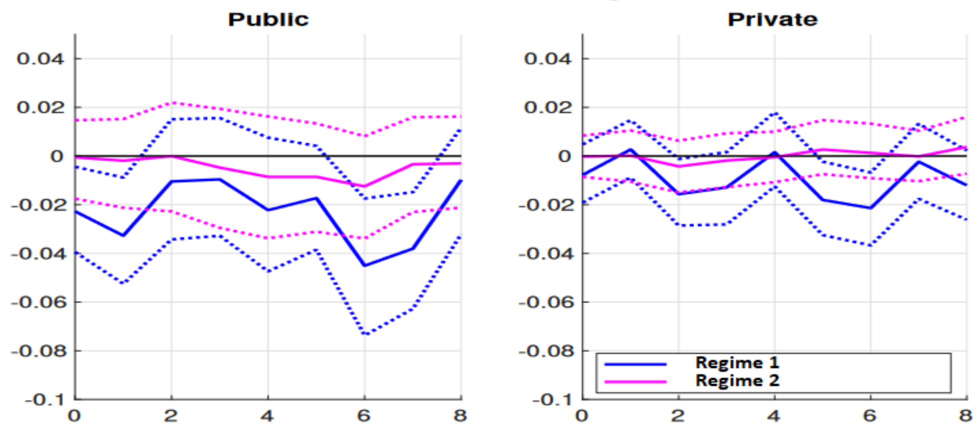

Notes: The figure shows the responses of disagreement about 8-quarter-ahead inflation expectations attributable to public information and non-public information following the Fed's response to news from Bauer and Swanson (2023). HAC standard errors are reported. Panel A presents the results from the linear model in Equation. (19), while Panel B shows the results from the nonlinear model in Equation (20). In Panel B, the blue lines represent the responses in Regime 1, while the magenta lines represent those in Regime 2. Non-public information is the primary source of disagreement in Regime 1, whereas public information is the main source in Regime 2. The average disagreement shares of non-public and public information are 0.9 and 0.1, respectively. In Panel B, the estimated responses for Regimes 1 and 2 are scaled by 0.9 and 0.1, respectively, so that their sum closely approximates the average impulse responses in Panel A. The dashed lines indicate the 95% posterior intervals.

Source: Authors' calculation

used to estimate $\beta^c_{8,t-1}$ and hence s_{t-1}, which are then used to compute impulse responses.

Panel B of Figure 13 shows the impulse responses for the two regimes. We report the effects of a one-standard-deviation innovation in the news component. The responses of regimes 1 and 2 are scaled by 0.9 and 0.1, corresponding to the average disagreement shares of non-public information and public information, respectively. This approach ensures that the sum of the two impulse responses closely approximates the average responses reported in Panel A. Notable differences emerge across the two regimes. When public information is the source of disagreement (magenta lines), the news component does not statistically significantly reduce disagreement. In contrast, when non-public information is the source, the reduction of disagreement is statistically significant on impact. This finding suggests that increased disagreement about public information may reflect situations where forecasters disagree about monetary policy communication, given that the Fed's communication is also public information. Thus, heightened disagreement attributable to public information may suggest a weakened effectiveness of monetary policy communication in anchoring inflation expectations.

This result has a few important implications. First, our empirical model effectively identifies the sources of information contributing to disagreement, highlighting that the portion attributable to public information is the amount that is reducible by monetary policy communication and is relevant in anchoring inflation expectations. Second, the Fed's interpretation of economic news plays a crucial role in shaping public information essential for inflation forecasting. In this sense, clear communication of the monetary authority can help anchor economic agents' expectations by reducing their disagreement about future macroeconomic conditions. Finally, the extent of disagreement driven by public information can serve as an auxiliary indicator of how well-anchored inflation expectations are.

In appendix E.1, we provide extensive robustness checks for the analysis presented in this section. First, we further consider disagreement about 10-year-ahead inflation instead of about 8-quarter ahead inflation projections. Second, we consider expanding the macroeconomic controls to include a measure of individual uncertainty from Binder (2017) and the consensus inflation expectations about next year.[30] Third, we consider not including lags for the news component of the monetary policy shock. Fourth, we also consider alternative methodologies such as the two-stage local projection with external instrumental variable (LP-IV). The results are robust.

7.2 Disagreement and Monetary Policy Effectiveness

This section explores the relationship between the role of public information in disagreement and the stabilizing effects of monetary policy shocks. Specifically, we assess the extent to which the disagreement share of public information affects the effectiveness of monetary policy.

[30] For this robustness check, we also include an external consensus measure, rather than using the model estimates, to ensure that our results are robust to the inclusion of external data.

Figure 14: PROPAGATION OF MONETARY POLICY SHOCKS OF THE TWO REGIMES (8-QUARTER AHEAD)

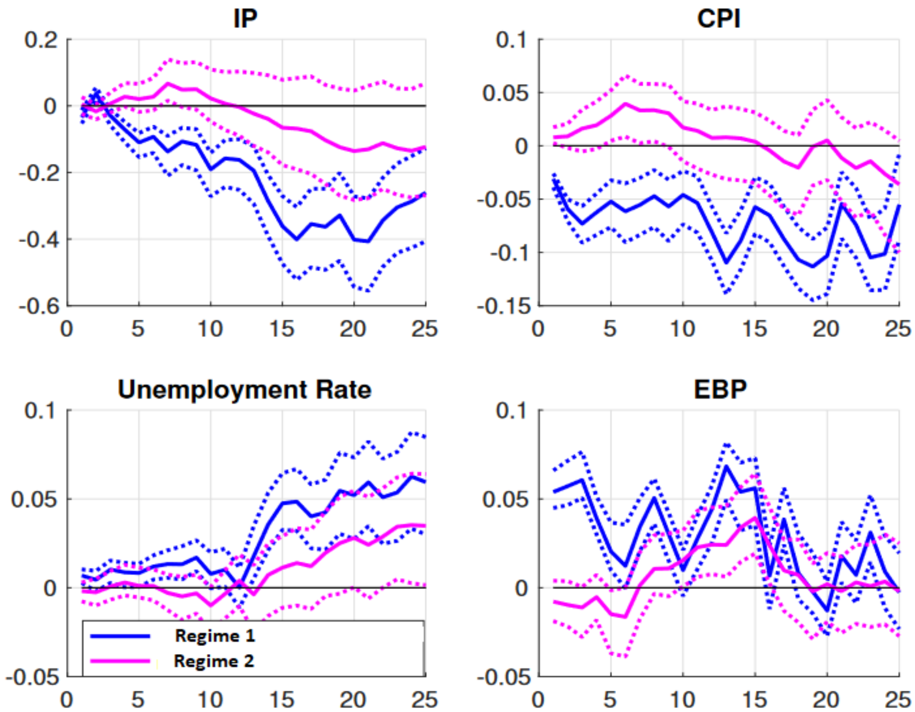

Notes: The figure reports the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2023). The figure shows the impulse responses of Regimes 1 and 2 scaled by 0.9 and 0.1, respectively. The blue lines show the responses when non-public information is the source of disagreement (Regime 1), while the magenta lines represent 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 dashed lines represent the 90% posterior intervals.

Sources: Authors' calculation

For this analysis, we utilize a nonlinear local projection model similar to the one in equation (20), with some modifications. First, the dependent variables are macroeconomic variables, namely 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. 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.[31]

[31]As a robustness check, we include the uncertainty measure from Binder (2017) and the consensus inflation expectations from the SPF (long-run and next year) as macro controls. For the 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. All robustness checks are reported in the appendix.

The nonlinear model has two regimes which are the same as those in equation (20). Non-public information is the source of the disagreement in Regime 1, and public information in Regime 2. Each regime is also distinguished by the disagreement share of public information 8 quarters ahead. Again, to avoid contemporaneous feedback from policy actions, we consider a one-period lag for the regime indicator. Since the disagreement share is quarterly, we assign the same quarterly value to the three months of the corresponding quarter and use the previous quarter's share as the regime indicator for month t.[32] The sample period is from October 1991 to December 2019.

Figure 14 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, again reflecting the average disagreement shares of non-public information and public information.[33] For the responses of IP growth and CPI inflation, we report the cumulative effects of the monetary policy shock, 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, the sensitivity of disagreement to public information—captured by the disagreement share of public information—is an important determinant of monetary policy effectiveness. When forecasters disagree about public information, including the stance of monetary policy, the stabilizing effects of monetary policy weaken significantly.[34] Second, our empirical result also offers a new perspective on the source of the price puzzle. Our result suggests that the

[32] As a robustness check, we also consider a one-month lag for s_t. Since the SPF is conducted during the first week in the second month of each quarter, the submitted forecasts largely reflect the information set through the previous quarter. For this reason, the one-month lag for s_t is less likely to create a feedback effect in the current period. The estimation results remain robust regardless.

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

[34] 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. However, beyond the difference data source, we further identify that the key component of disagreement that weakens the effectiveness of monetary policy is the portion of that disagreement which is attributable to public information.

contribution of public information to disagreement is an important channel driving the price puzzle. This implies that the price puzzle may be related to the effectiveness of monetary policy communication or the credibility of the central bank.[35] Overall, our empirical results suggest that clear communication of monetary policy can enhance policy effectiveness by reducing disagreement attributable to public information.

Appendix E.2 provides extensive robustness checks for the analysis presented in this section. First, we consider the disagreement share for 10-year-ahead inflation instead of 8-quarter-ahead inflation projections. Second, we explore the use of regime indicators based on both a one-month lag and the current quarter to assess the sensitivity of the results to the timing of the regime indicator. Third, we expand the macroeconomic controls to include the uncertainty measure from Binder (2017), the monetary policy uncertainty from Husted et al. (2020), and the consensus inflation expectations for the next year. Note that we control for individual-level uncertainty and uncertainty regarding monetary policy to isolate the effects of disagreement about public information on the effectiveness of monetary policy. Fourth, we present the impulse responses with the average disagreement shares seen during economic recessions. The empirical results remain robust across these alternative specifications.[36]

8 Conclusion

This paper makes three key contributions. First, we develop a parametric model which we call the 'individual term-structure of inflation expectations,' which uses two factors—level and slope—to describe forecasters' inflation predictions across different time horizons. Second, we extend this model to a dynamic factor framework, decomposing individual-level elements into

[35] Falck et al. (2021) also observe that the price puzzle becomes more pronounced when professional forecasters disagree, based on SPF data. Our finding differs from Falck et al. (2021) in that we identify public information as the key driver. The sensitivity of disagreement to public information increases when overall disagreement rises, explaining the observed consistency between the two findings.

[36] Additionally, we examine alternative regimes based on high- and low- level disagreement without further decomposing by information source. See Appendix E.2 for more details. In this analysis, the regimes are determined by whether total 4-quarter-ahead disagreement—the forecast horizon frequently often considered by previous studies. Similar to the earlier findings, we find larger and more statistically significant contractionary effects of monetary policy in the regime of low disagreement, while the effects are muted and less statistically significant in the regime of high disagreement. However, the price puzzle is observed in the low disagreement regime not in the high disagreement regime, which is somewhat different from the earlier studies. The difference may stem from differences in data and the sample period. All told, the difference in policy effectiveness between the two regimes is less pronounced when distinguishing the regimes based on the level of disagreement compared to our baseline. This observation reinforces the importance of the information source in evaluating the effectiveness of monetary policy.

common and idiosyncratic components and an individual constant term. We build a noisy information model where forecasters exhibit heterogeneous reactions to public information. We show that our structural model maps directly into the estimates from our dynamic factor model, which allows to interpret the common and idiosyncratic components as responses to public and private information, respectively, and the individual fixed effects as long-term beliefs. This unique decomposition allows us to separate disagreement about inflation projections into these three sources at each point in time and across forecasting horizons. Third, we investigate how the sensitivity of disagreement to public information impacts the effectiveness of monetary policy and the anchoring of inflation expectations. All of these contributions are entirely novel to the literature.

Our research highlights the importance of considering disagreement when evaluating the anchoring of inflation expectations. Although consensus forecasts suggest well-anchored long-term expectations, inflation forecasts across forecasting horizons show greater disagreement and increased skewness, particularly during the Great Recession and the COVID-19 pandemic. This observation indicates that expectations were likely less well-anchored than previously thought. Our model shows that the consensus forecast and disagreement often yield distinct insights into the anchoring of agents' inflation expectations.

We find distinct roles for the three sources of disagreement across forecasting horizons. Long-term beliefs and private information account for the majority of disagreement in long-run and short-run expectations respectively. The role of public information in disagreement is small on average. However, during economic downturns and periods of high inflation uncertainty, public information becomes a key driver of disagreement. A noisy information model with heterogeneous reactions to public information predicts that this varying importance of public information in the disagreement arises in response to large public news shocks.

Finally, we find that when public information is the main source of disagreement, the economy's responses to monetary policy shock are delayed significantly and a price puzzle emerges. When public information is not important in disagreement, monetary policy has rapid and statistically significant stabilizing effects. These results suggests that disagreement about public information is an important determinant of the effectiveness of monetary policy, underscoring the importance of anchoring inflation expectations.

Central bank communication about the macroeconomic outlook plays a crucial role in managing inflation expectations, especially during times when economic agents are highly attentive to monetary policy and macroeconomic news. Clear communication by policymakers can reduce disagreement and provide a stronger anchor for inflation expectations during periods of heightened uncertainty. Our findings offer a new perspective on the source of the price puzzle and its relationship with central banks' expectations management. We leave this topic to future research.

References

Ahn, H. J. and Luciani, M. (2024). Common and idiosyncratic information. Finance and Economics Discussion Series 2020-024, Board of Governors of the Federal Reserve System (U.S.).

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

Angeletos, G. and La'O, J. (2013). Sentiments. *Econometrica*, 81(2):739–779.

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.

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.

Barillas, F. and Nimark, K. P. (2017). Speculation and the Term Structure of Interest Rates. *The Review of Financial Studies*, 30(11):4003–4037.

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.

Burnside, C., Eichenbaum, M., and Rebelo, S. (2016). Understanding booms and busts in housing markets. *Journal of Political Economy*, 124(4):1088–1147.

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

Clark, T. E., Ganics, G., and Mertens, E. (2022). Constructing the term structure of uncertainty from the ragged edge of spf forecasts. Working paper.

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

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

Cornand, C. and Hubert, P. (2022). Information frictions across various types of inflation expectations. *European Economic Review*, 146:104175.

Crump, R. K., Eusepi, S., Moench, E., and Preston, B. (2021). The Term Structure of Expectations. Staff Reports 992, Federal Reserve Bank of New York.

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.

Diebold, F. X. and Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2):337–364.

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.

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

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

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.

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.

Ilut, C. L. and Schneider, M. (2014). Ambiguous Business Cycles. *American Economic Review*, 104(8):2368–2399.

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

Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review*, 99(5):2050–84.

Mackowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.

Mankiw, N. G. and Reis, R. (2002). Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve*. *The Quarterly Journal of Economics*, 117(4):1295–1328.

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

Michael, W. (2002). Inflation Stabilization and Welfare. *The B.E. Journal of Macroeconomics*, 2(1):1–53.

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).

Scheinkman, J. A. and Xiong, W. (2003). Overconfidence and Speculative Bubbles. *Journal of Political Economy*, 111(6):1183–1219.

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.

Woodford, M. (2001). Imperfect Common Knowledge and the Effects of Monetary Policy. NBER Working Papers 8673, National Bureau of Economic Research, Inc.

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

$$\mathbf{x}_t := \left[\begin{array}{c} L_t, S_t, \varepsilon^L_{1,t}, \varepsilon^S_{1,t}, \cdots, \varepsilon^L_{n,t}, \varepsilon^S_{n,t} \end{array}\right]' \tag{A1}$$

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

$$\mathbf{F} := \begin{bmatrix} \rho_L & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & \rho_S & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & \rho^L_\varepsilon & 0 & 0\dots & 0 & 0 \\ 0 & 0 & 0 & \rho^S_\varepsilon & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & \rho^L_\varepsilon & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & \rho^S_\varepsilon \end{bmatrix}. \tag{A2}$$

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

$$\mathbf{u}_t := \left[\begin{array}{c} u^L_t, u^S_t, u^L_{1,t}, u^S_{1,t}, \cdots, u^L_{n,t}, u^S_{n,t} \end{array}\right]'. \tag{A3}$$

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

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

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

$$\begin{aligned} \pi_{i,t\to t+19|t} &= \frac{4}{19}\left(5\pi_{i,t-1\to t+19|t} - \frac{1}{4}\pi_{i,t-1\to t|t}\right) \\ \pi_{i,t\to t+39|t} &= \frac{4}{39}\left(10\pi_{i,t-1\to t+19|t} - \frac{1}{4}\pi_{i,t-1\to t|t}\right) \end{aligned}$$

- In Q2, we define

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

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

- In Q3, we define

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

- In Q4, we define

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

For the nowcasts $\pi_{i,t-1\to t|t}$ and backcasts $\pi_{i,t-2\to t-1|t}$, we use the reported values from the SPF. For the two and three period backcasts $\pi_{i,t-3\to t-2|t}$ and $\pi_{i,t-4\to t-3|t}$, 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

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

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

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

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

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

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

Finally, we define the measurement equation mapping matrix **H** as

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

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 $\boldsymbol{\theta}_1 = \left\{\alpha_i^L, \alpha_i^S, \beta_i^L, \beta_i^S\right\}_{i=1}^N$ 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 $\left[\alpha_i^L, \alpha_i^S, \beta_i^L, \beta_i^S\right]'$ across each foracaster i, where

$$\begin{bmatrix} \alpha_i^L \\ \alpha_i^S \\ \beta_i^L \\ \beta_i^S \end{bmatrix} \sim N\left(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i\right)$$

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.

3. Sample $\boldsymbol{\theta}_3 = \text{vec}(\boldsymbol{B})$ conditional on remaining parameters. Since $\varepsilon_{i,t}^L$ and $\varepsilon_{i,t}^S$ 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.

4. Sample $\boldsymbol{\theta}_4 = \left[\sigma_L^2, \sigma_S^2\right]'$ conditional on remaining parameters. Since $\varepsilon_{i,t}^L$ and $\varepsilon_{i,t}^S$ 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. After some algebraic manipulation one can show that

$$L_{i,t} = \rho^L_{i,\varepsilon} L_{i,t-1} + \beta^L_i \left[\left(1 - \frac{\rho^L_{i,\varepsilon}}{\rho_L} \right) L_t + \frac{\rho^L_{i,\varepsilon}}{\rho_L} u^L_t \right] + v^L_{i,t}$$
$$S_{i,t} = \rho^S_{i,\varepsilon} S_{i,t-1} + \beta^S_i \left[\left(1 - \frac{\rho^S_{i,\varepsilon}}{\rho_S} \right) S_t + \frac{\rho^S_{i,\varepsilon}}{\rho_S} u^S_t \right] + v^S_{i,t}$$

Taking the cross-sectional mean across forecasters of the common term (involving β_i) and adding and subtracting we get

$$L_{i,t} = \rho^L_{i,\varepsilon} L_{i,t-1} + \beta^L_i \left[\left(1 - \frac{\bar{\rho}^L_{\varepsilon}}{\rho_L} \right) L_t + \frac{\bar{\rho}^L_{\varepsilon}}{\rho_L} u^L_t \right] + v^L_{i,t} - \beta^L_i \frac{\rho^L_{i,\varepsilon} - \bar{\rho}^L_{\varepsilon}}{\rho_L} \left(L_t - u^L_t \right)$$
$$S_{i,t} = \rho^S_{i,\varepsilon} S_{i,t-1} + \beta^S_i \left[\left(1 - \frac{\bar{\rho}^S_{\varepsilon}}{\rho_S} \right) S_t + \frac{\bar{\rho}^S_{\varepsilon}}{\rho_S} u^S_t \right] + v^S_{i,t} - \beta^S_i \frac{\rho^S_{i,\varepsilon} - \bar{\rho}^S_{\varepsilon}}{\rho_S} \left(S_t - u^S_t \right)$$

Matching terms to those in the equations which govern the evolution of beliefs in the noisy information model, we can solve for the statistical model parameters as functions of the noisy information model parameters. Starting with the slope dynamics we have

$$\rho^S_{i,\varepsilon} S_{i,t-1} = \left(1 - g^c_i \right) \rho_c F_{i,t-1} c_{t-1} \tag{C13}$$
$$\beta^S_i \left[\left(1 - \frac{\bar{\rho}^S_{\varepsilon}}{\rho_S} \right) S_t + \frac{\bar{\rho}^S_{\varepsilon}}{\rho_S} u^S_t \right] = g^c_i c_t + g^c_{i,y} u^c_t \tag{C14}$$
$$v^S_{i,t} - \beta^S_i \frac{\rho^S_{i,\varepsilon} - \bar{\rho}^S_{\varepsilon}}{\rho_S} \left(S_t - u^S_t \right) = g^c_i v^c_{i,t} \tag{C15}$$

We start by recognizing that $S_{i,t} = F_{i,t} c_t$, and thus from the first equation we immediately obtain

$$\rho^S_{i,\varepsilon} = \left(1 - g^c_i \right) \rho_c \tag{C16}$$

We now assume that $g^c_i \sim i.i.d. \left(\bar{g}^c, (\sigma^c_g)^2 \right)$ across forecasters (this can be justified by agents having different variances of private signals).

Working with equation (C14) and assuming that $\rho_S = \rho_c$,

$$
\begin{aligned}
&\beta_i^S \left[\left(1 - \frac{\bar{\rho}_\varepsilon^S}{\rho_S}\right) S_t + \frac{\bar{\rho}_\varepsilon^S}{\rho_S} u_t^S \right] \\
&= \beta_i^S \left[\left(1 - \frac{\left(1-\bar{g}^c\right)\rho_c}{\rho_S}\right) S_t + \frac{\left(1-\bar{g}^c\right)\rho_c}{\rho_S} u_t^S \right] \\
&= \beta_i^S \left[\bar{g}^c S_t + \left(1-\bar{g}^c\right) u_t^S \right] \\
&= \beta_i^S \left[\bar{g}^c \left(\rho_S S_{t-1} + u_t^S\right) + \left(1-\bar{g}^c\right) u_t^S \right] \\
&= \beta_i^S \left[\rho_c \bar{g}^c S_{t-1} + u_t^S \right] \\
&= g_i^c c_t + g_{i,y}^c u_t^c
\end{aligned}
$$

We know that the last two expressions must have the same time-series variance, thus

$$
\frac{\left(\rho_c \beta_i^S \bar{g}^c\right)^2}{1-\rho_c^2} + \left(\beta_i^S\right)^2 = \left(g_i^c\right)^2 \frac{(\sigma_\varepsilon^c)^2}{1-\rho_c^2} + \left(g_{i,y}^c\right)^2 (\sigma_u^c)^2
$$

Solving for β_i^S gives

$$
\beta_i^S = \left(\frac{\left(g_i^c\right)^2 (\sigma_\varepsilon^c)^2 + \left(g_{i,y}^c\right)^2 (\sigma_u^c)^2 \left(1-\rho_c^2\right)}{1 + \left(\rho_c \bar{g}^c\right)^2 - \rho_c^2} \right)^{1/2} \tag{C17}
$$

Next we work with equation (C15),

$$
\begin{aligned}
&v_{i,t}^S - \beta_i^S \frac{\rho_{i,\varepsilon}^S - \bar{\rho}_\varepsilon^S}{\rho_S} \left(S_t - u_t^S\right) \\
&= v_{i,t}^S + \beta_i^S \left(g_i^c - \bar{g}^c\right) \rho_c S_{t-1} \\
&= g_{i,z}^c v_{i,t}^c
\end{aligned}
$$

As before, we know that the last two expressions must have the same time-series variance, thus

$$
(\sigma_{i,v}^S)^2 + \frac{\left(\beta_i^S\right)^2 \left[\left(g_i^c - \bar{g}^c\right)\rho_c\right]^2}{1-\rho_c^2} = \left(g_{i,z}^c \sigma_{i,v}^c\right)^2
$$

Solving for $(\sigma^S_{i,v})^2$ gives

$$(\sigma^S_{i,v})^2 = \left(g^c_{i,z}\sigma^c_{i,v}\right)^2 - \frac{\left(\beta^S_i\right)^2\left[\left(g^c_i - \bar{g}^c\right)\rho_c\right]^2}{1-\rho_c^2} \tag{C18}$$

Using similar arguments for the level factor / permanent component (but matching the conditional variance instead of the unconditional variance), we obtain

$$\rho^L_{i,\varepsilon} = 1 - g^\tau_i \tag{C19}$$

$$\beta^L_i = \left[\left(g^\tau_i\right)^2(\sigma^\tau_\varepsilon)^2 + \left(g^\tau_{i,y}\right)^2(\sigma^\tau_u)^2\right]^{1/2} \tag{C20}$$

$$\sigma^L_{i,v} = g^\tau_{i,z}\sigma^\tau_{i,v} \tag{C21}$$

□

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

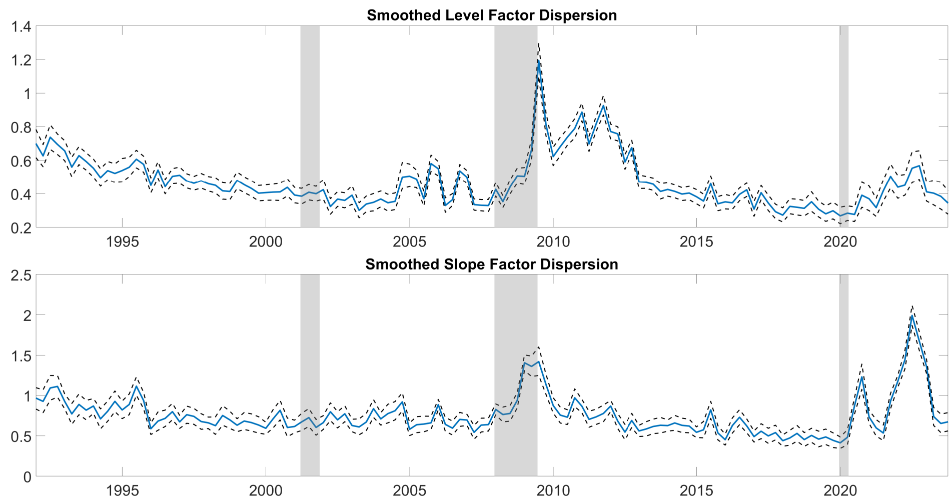

Notes: The shaded areas denote the NBER recessions.
Sources: Authors' calculation

Figure D2: SMOOTHED IDIOSYNCRATIC FACTOR SKEWNESS

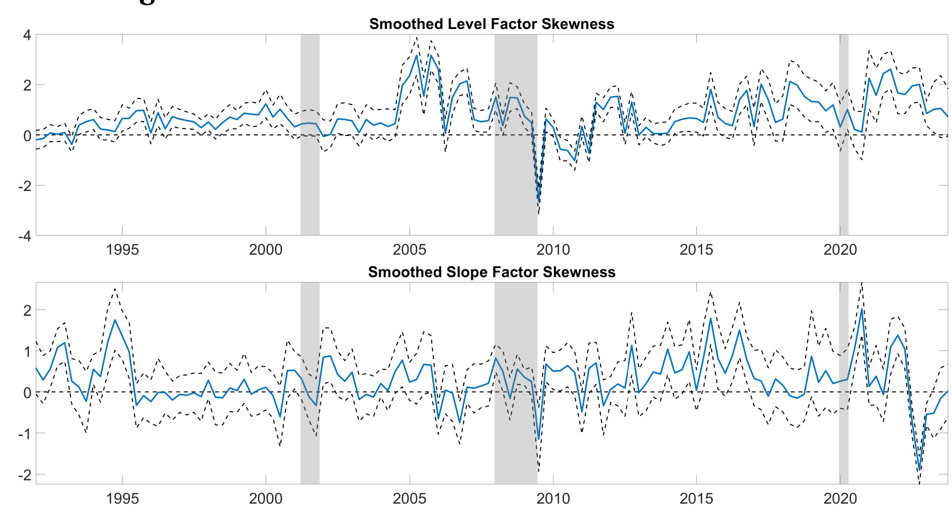

Notes: The shaded areas denote the NBER recessions.
Sources: Authors' calculation

Figure D3: Smoothed Idiosyncratic Factor Kurtosis

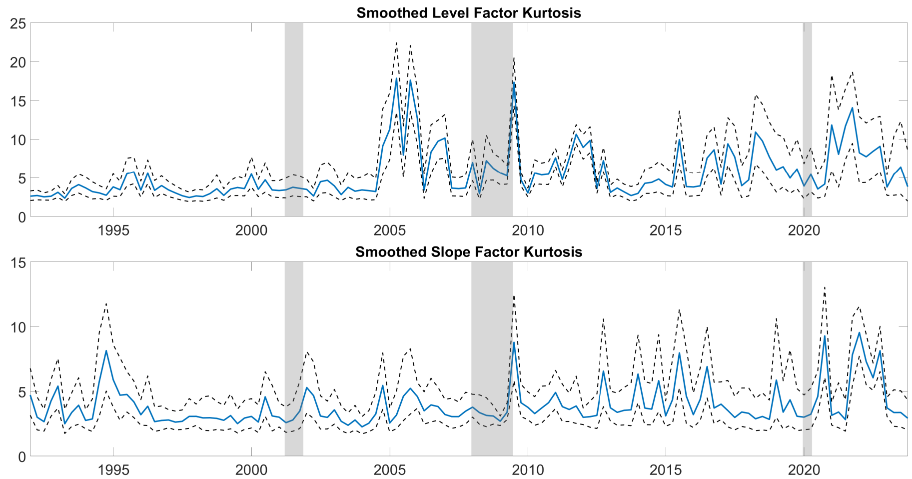

5

Notes: The shaded areas denote the NBER recessions.

Sources: Authors' calculation

Figure D4: Forecast Variance Decomposition

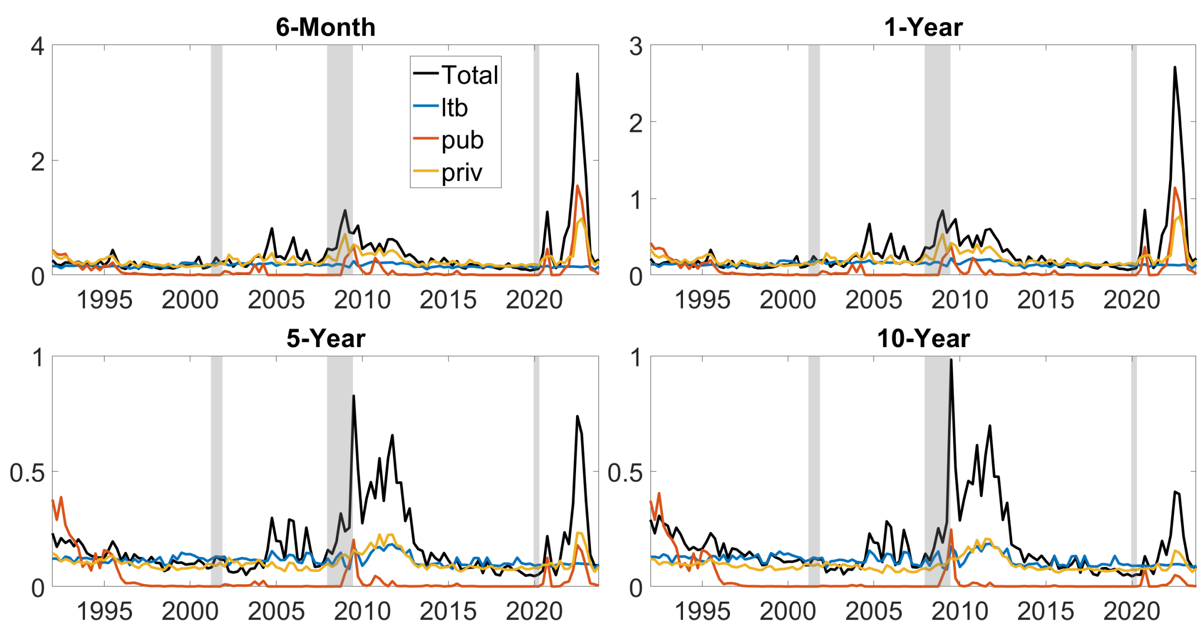

Notes: The figure shows the decomposition of the cross-sectional variance of inflation forecasts (black line) into the components driven by individual long-term beliefs (denoted by ltb, blue line), heterogeneous responses to public information (denoted by pub, red line), and private information (denoted by priv, yellow line). Each line corresponds to the posterior median. The shaded areas denote NBER recessions.

Sources: Authors' calculation

Figure D5: Forecast Variance Share of Common Component

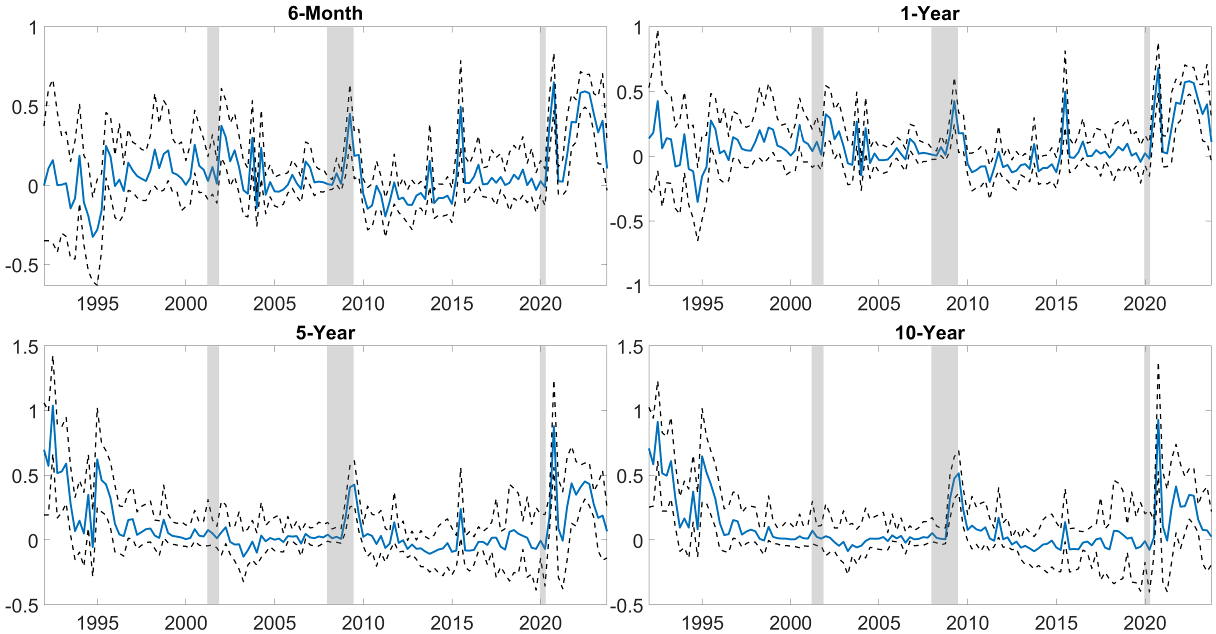

Notes: The figure shows the fraction of overall disagreement about four sets of inflation expectations driven by heterogeneous responses to common information over time as measured by the β measure proposed in Fujita and Ramey (2009). The shaded areas denote NBER recessions.

Sources: Authors' calculation

Appendix E Additional Results on Monetary Policy Effectiveness

E.1 Effects of News Component on Disagreement

Figure E6: (ROBUSTNESS CHECK 1) PROPAGATION OF FED'S REACTIONS TO NEWS: 10-YEAR AHEAD

Panel A. Average Responses

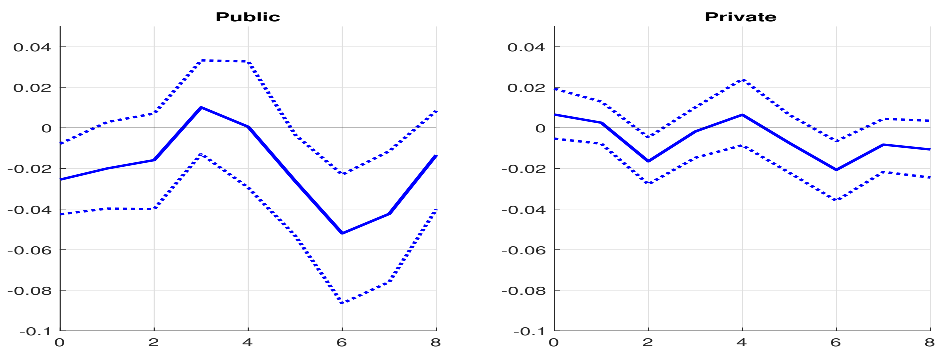

Panel B. Nonlinear Responses

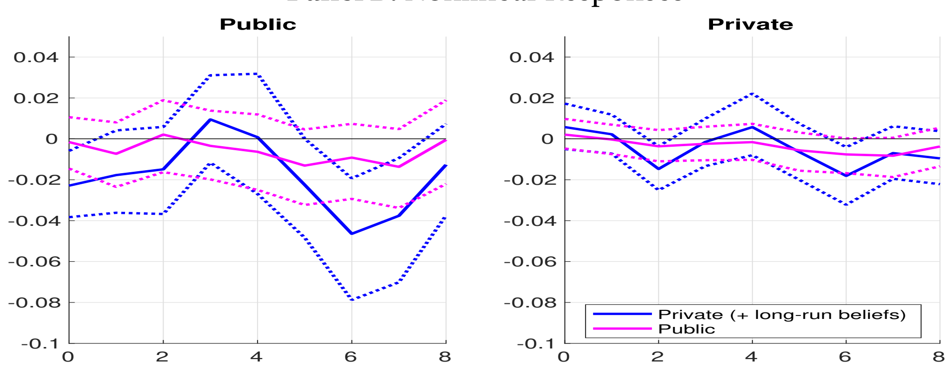

Notes: The figure shows the responses of disagreement about 40-quarter-ahead inflation expectations attributable to public information and non-public information following the Fed's response to news from Bauer and Swanson (2022). The alternative model also includes the individual-level uncertainty from Binder (2017) and the consensus inflation expectation for the next year as the additional controls. In addition, the regimes are determined by the disagreement shares for 40-quarter ahead. HAC standard errors are reported. Panel A presents the results from the linear model in Equation. (19), while Panel B shows the results from the nonlinear model in Equation (20). In Panel B, the magenta lines represent the responses when the previous period's disagreement is attributable to public information, and the blue lines represent the responses when the previous period's disagreement is attributable to non-public information. Estimates for regimes 1 and 2 are scaled by 0.9 and 0.1, respectively, corresponding to the average values of the disagreement shares for public and non-public information for 10-year ahead. The dashed lines indicate the 95% posterior intervals.

Source: Authors' calculation

Figure E7: (Robustness check 2) Propagation of Fed's reactions to news: Alternative controls

Panel A. Average Responses

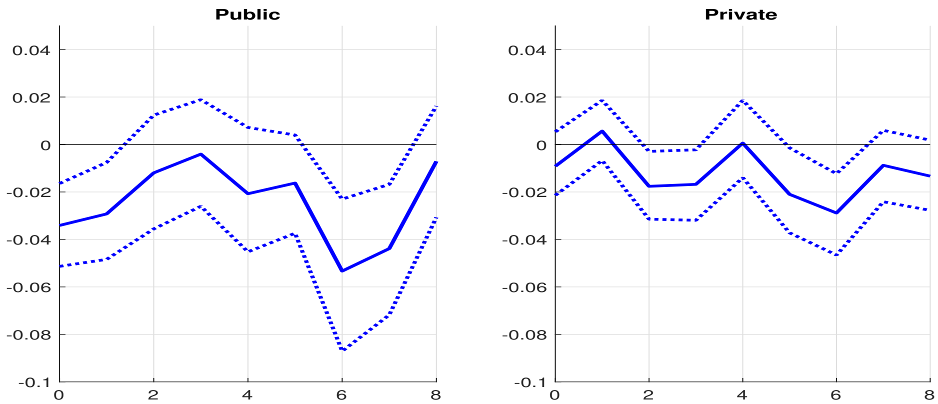

Panel B. Nonlinear Responses

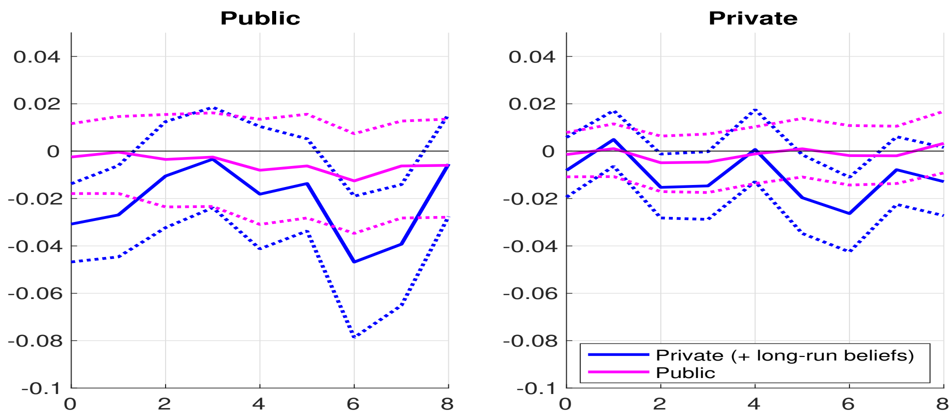

Notes: The figure shows the responses of disagreement about 8-quarter-ahead inflation expectations attributable to public information and non-public information following the Fed's response to news from Bauer and Swanson (2022). The alternative model with the individual-level uncertainty measure from Binder (2017) and the consensus inflation expectations for the next year are additionally considered as the controls. HAC standard errors are reported. Panel A presents the results from the linear model in Equation. (19), while Panel B shows the results from the nonlinear model in Equation (20). In Panel B, the magenta lines represent the responses when the previous period's disagreement is attributable to public information, and the blue lines represent the responses when the previous period's disagreement is attributable to non-public information. Estimates for regimes 1 and 2 are scaled by 0.9 and 0.1, respectively, corresponding to the average values of the disagreement shares for public and non-public information for 8-quarter ahead. The dashed lines indicate the 95% posterior intervals.

Source: Authors' calculation

Figure E8: (ROBUSTNESS CHECK 3) PROPAGATION OF FED'S REACTIONS TO NEWS: NO LAGS OF z_t

Panel A. Average Responses

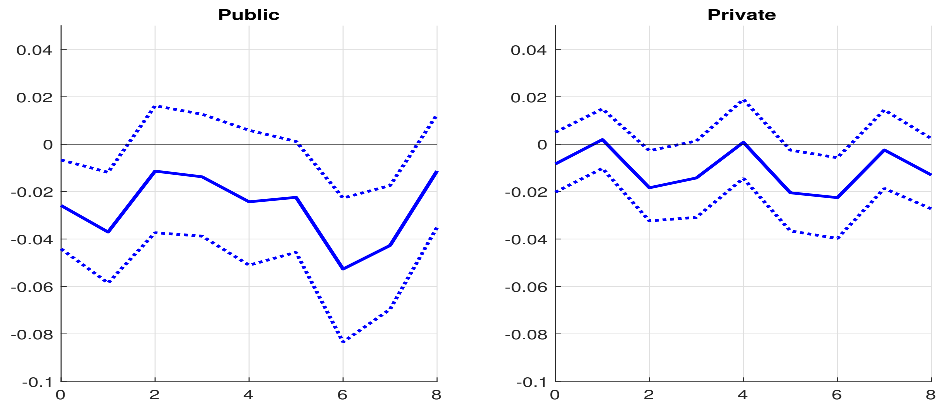

Panel B. Nonlinear Responses

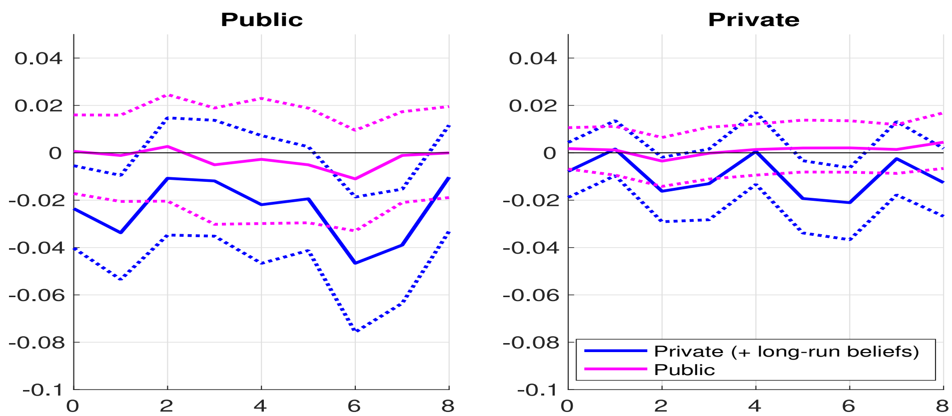

Notes: The figure shows the responses of disagreement about 8-quarter-ahead inflation expectations attributable to public information and non-public information following the Fed's response to news from Bauer and Swanson (2022). The alternative model with zero lags for z_t is considered. HAC standard errors are reported. Panel A presents the results from the linear model in Equation. (19), while Panel B shows the results from the nonlinear model in Equation (20). In Panel B, the magenta lines represent the responses when the previous period's disagreement is attributable to public information, and the blue lines represent the responses when the previous period's disagreement is attributable to non-public information. Estimates for regimes 1 and 2 are scaled by 0.9 and 0.1, respectively, corresponding to the average values of the disagreement shares for public and non-public information 8-quarter ahead. The dashed lines indicate the 95% posterior intervals.

Source: Authors' calculation

Figure E9: (ROBUSTNESS CHECK 4) PROPAGATION OF FED'S REACTIONS TO NEWS: LP-IV

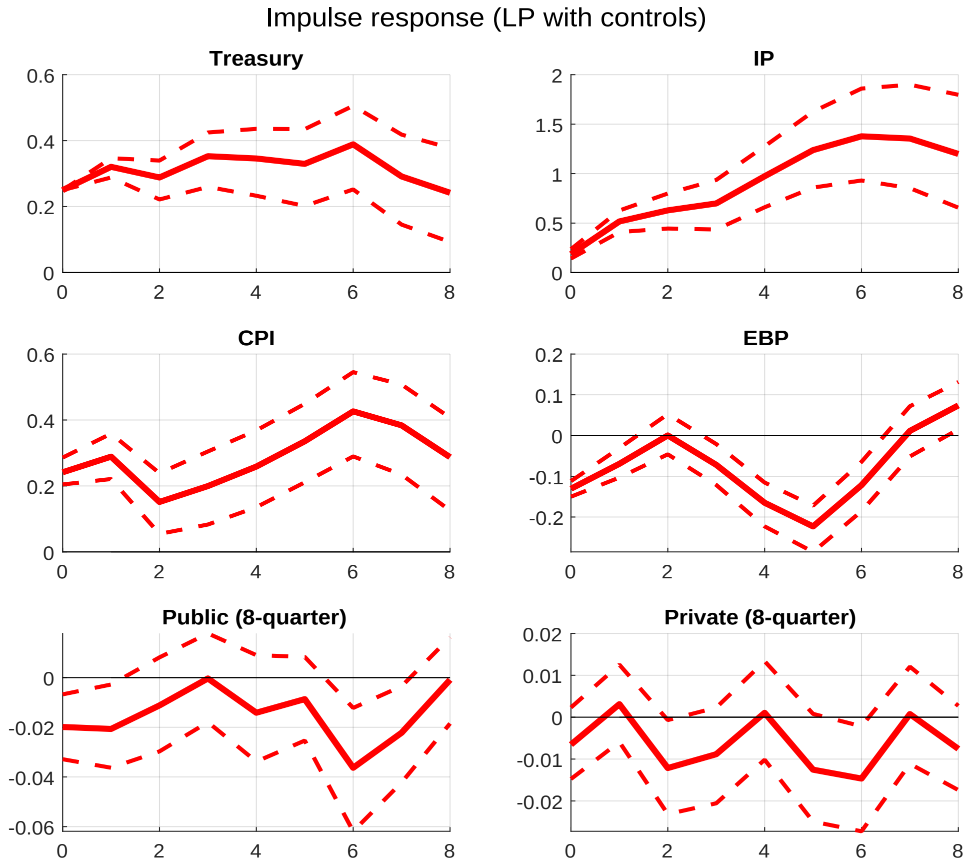

Notes: The figure shows the responses of four macroeconomic variables and the disagreement about 8-quarter-ahead inflation expectations driven by public and private information (including long-term beliefs) to Fed's response to news from Bauer and Swanson (2023). The LP-IV is employed for the estimation. The LP-IV includes four lags of dependent variables and four lags of external shocks in the first stage regression. The F-statistics is significantly larger than 10. The cumulative responses are reported for the IP growth and CPI inflation. The dashed line captures the 90 percent posterior intervals.

Source: Authors' calculation

E.2 Nonlinear Effects of a Traditional Monetary Policy Shock

Figure E10: (ROBUSTNESS CHECK 1) PROPAGATION OF MONETARY POLICY SHOCKS: 10-YEAR AHEAD

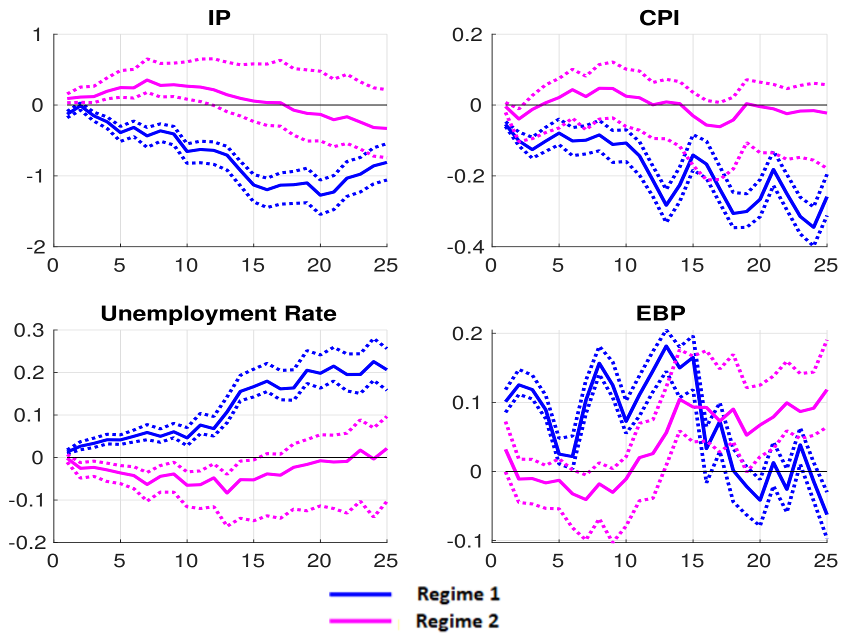

Notes: The figure reports the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2022). For the regime indicator, 10-year ahead disagreement shares of the previous quarter are considered. Panel A shows the impulse responses of regimes 1 and 2 scaled by 0.9 and 0.1, respectively. The blue lines represent the responses when disagreement is attributable non-public information (regime 1), while the magenta lines show the responses when disagreement is attributable to public information (regime 2). In each panel, 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 dashed lines represent the 90% posterior intervals.

Sources: Authors' calculation

Figure E11: (ROBUSTNESS CHECK 2) PROPAGATION OF MONETARY POLICY SHOCKS OF THE TWO REGIMES: 1-MONTH LAG FOR THE REGIME INDICATOR)

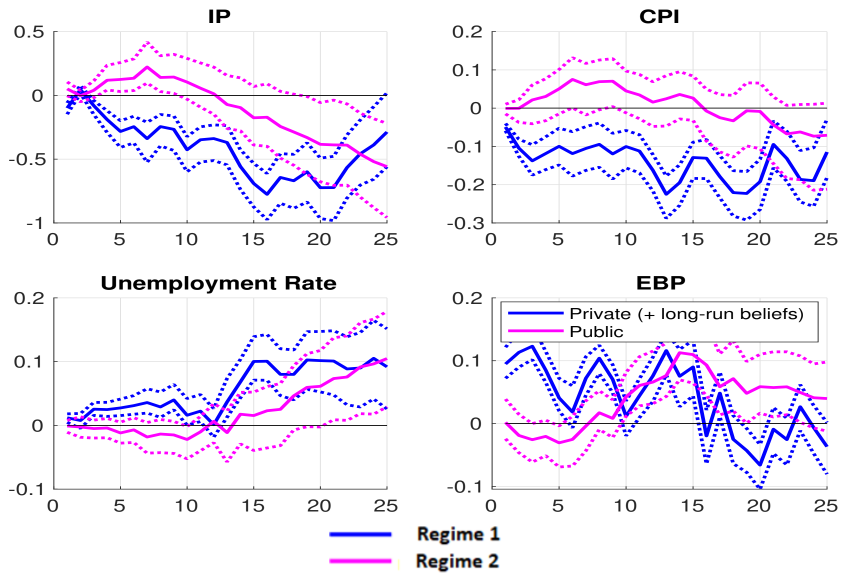

Notes: The figure reports the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2022). For the regime indicator, the one-month lags of 8-quarter-ahead disagreement shares are considered. The panels show the impulse responses of regimes 1 and 2 scaled by 0.9 and 0.1, respectively. The blue lines represent the responses when disagreement is attributable non-public information (regime 1), while the magenta lines show the responses when disagreement is attributable to public information (regime 2). In each panel, 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 dashed lines represent the 90% posterior intervals.

Sources: Authors' calculation

Figure E12: (ROBUSTNESS CHECK 3) PROPAGATION OF MONETARY POLICY SHOCKS OF THE TWO REGIMES: NO LAG FOR THE REGIME INDICATOR)

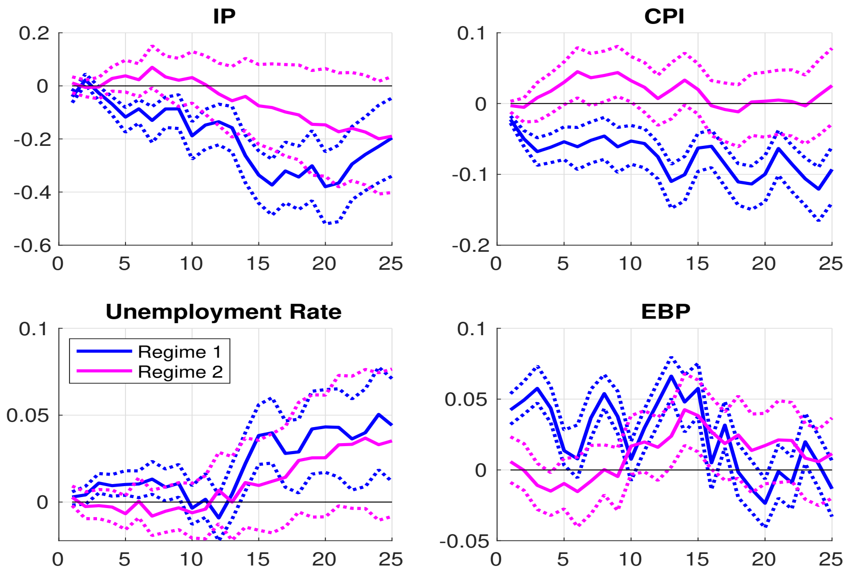

Notes: The figure reports the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2022). For the regime indicator, 8-quarter-ahead disagreement shares are considered without a lag. The panels show the impulse responses of regimes 1 and 2 scaled by 0.9 and 0.1, respectively. The blue lines represent the responses when disagreement is attributable non-public information (regime 1), while the magenta lines show the responses when disagreement is attributable to public information (regime 2). In each panel, 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 dashed lines represent the 90% posterior intervals.

Sources: Authors' calculation

Figure E13: (ROBUSTNESS CHECK 4) PROPAGATION OF MONETARY POLICY SHOCKS OF THE TWO REGIMES: ALTERNATIVE CONTROLS (1)

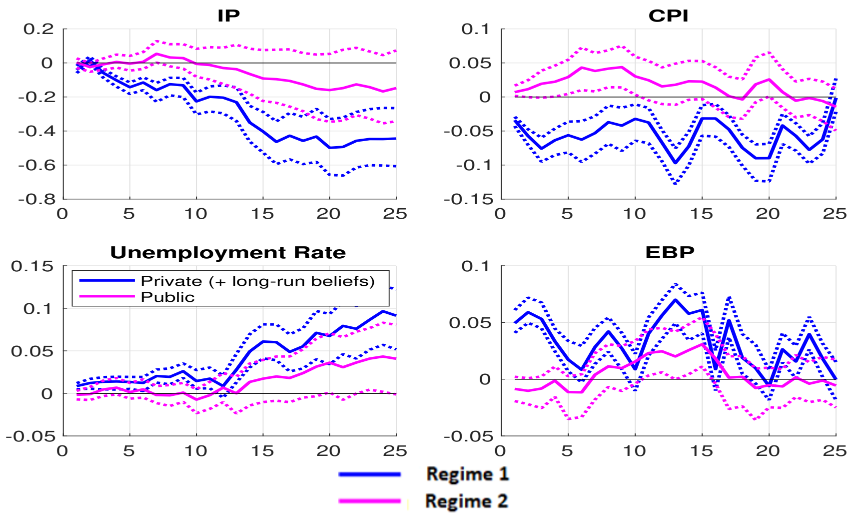

Notes: The figure reports the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2022). For the regime indicator, 8-quarter ahead disagreement shares of the previous quarter are considered. As the controls, the uncertainty measure from Binder (2017) and the consensus expectations for the next year are additionally included. The panels show the impulse responses of regimes 1 and 2 scaled by 0.9 and 0.1, respectively. The blue lines represent the responses when disagreement is attributable non-public information (regime 1), while the magenta lines show the responses when disagreement is attributable to public information (regime 2). In each panel, 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 dashed lines represent the 90% posterior intervals.

Sources: Authors' calculation

Figure E14: (ROBUSTNESS CHECK 4) PROPAGATION OF MONETARY POLICY SHOCKS OF THE TWO REGIMES: ALTERNATIVE CONTROLS WITH MONETARY POLICY UNCERTAINTY (2)

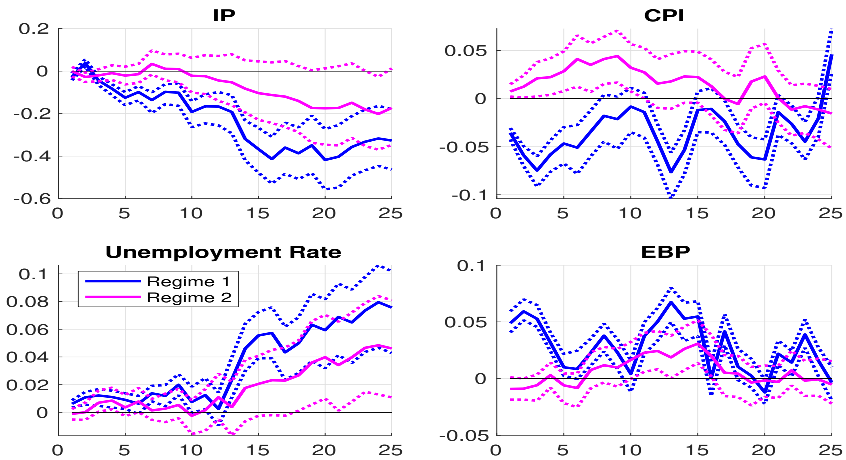

Notes: The figure reports the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2022). For the regime indicator, 8-quarter ahead disagreement shares of the previous quarter are considered. As the controls, the uncertainty measure from Binder (2017), the measure of monetary policy uncertainty from Husted et al. (2020), and the consensus expectations for the next year are additionally included. The panels show the impulse responses of regimes 1 and 2 scaled by 0.9 and 0.1, respectively. The blue lines represent the responses when disagreement is attributable non-public information (regime 1), while the magenta lines show the responses when disagreement is attributable to public information (regime 2). In each panel, 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 dashed lines represent the 90% posterior intervals.

Sources: Authors' calculation

Figure E15: (ROBUSTNESS CHECK 5) PROPAGATION OF MONETARY POLICY SHOCKS OF THE TWO REGIMES WITH RECESSION WEIGHTS (8-QUARTER AHEAD)

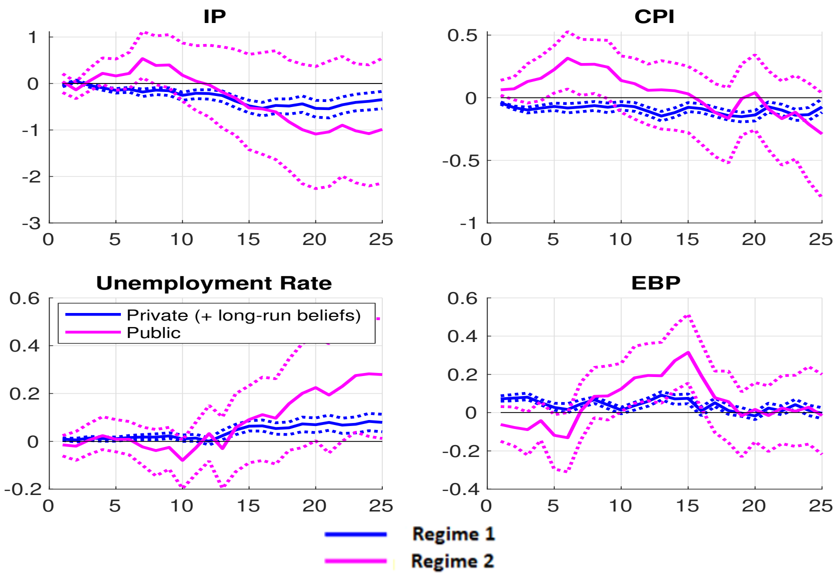

Notes: The figure reports the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2022). For the regime indicator, 8-quarter ahead disagreement shares of the previous quarter are considered. The figure shows the impulse responses of regimes 1 and 2 scaled by 0.6 and 0.4, respectively. These weights represent the disagreement shares of non-public and public information during an average economic downturn of the period of high inflation uncertainty. The blue lines show the responses when non-public information is the source of disagreement (regime 1), while the magenta lines represent 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 dashed lines represent the 90% posterior intervals.

Sources: Authors' calculation

Here, we further examine the responses during periods of economic recessions and heightened inflation risks. The responses of regimes 1 and 2 are scaled by 0.6 and 0.4, reflecting the average disagreement shares of non-public information and public information during these periods, respectively. Compared to normal times, the responses in regime 2 become dominant. Overall, the effects of monetary policy are not statistically significant. For unemployment rate and the EBP, the statistically significant effects appear but with substantial delays.

Figure E16: (ROBUSTNESS CHECK 6) PROPAGATION OF MONETARY POLICY SHOCKS: HIGH AND LOW DISAGREEMENT REGIMES

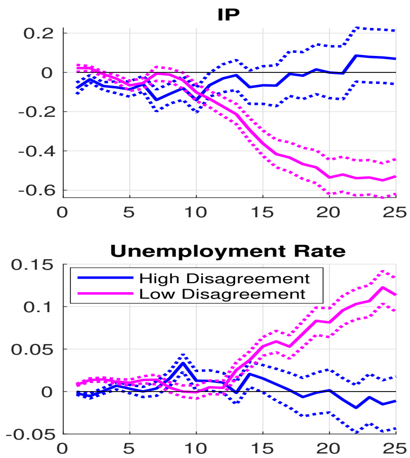

Notes: The figure presents the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2022). The responses to a one-standard-deviation innovation shock are reported. The panels show the impulse responses of the low-disagreement and high-disagreement regimes scaled by 0.35 and 0.65, respectively, representing the average regime probabilities. Each regime probability is computed as follows. After standardizing the level of 4-quarter-ahead disagreement computed with the variance of forecasts, we plug in the standardized value into a logit function with the smoothing parameter 5 to produce the probability of low-disagreement regime. The probability of high-disagreement regime is one minus the probability of low-disagreement regime. For the controls of local projection, we consider the individual-level uncertainty from Binder (2017) and next year's consensus inflation expectations in addition to the baseline controls. The blue lines represent the responses of the high-disagreement regime, while the magenta lines show the responses of the low-disagreement regime. In each panel, the upper left figure shows the cumulative response of percent changes in industrial production; the upper right figure shows the cumulative response of percent changes in the CPI; the bottom left figure displays the response of the unemployment rate; and the bottom right figure presents the response of the excess bond premium (EBP). The dashed lines represent the 95% posterior intervals.

Sources: Authors' calculation

Appendix F Literature Review of Related Theory

This section discusses implications of our empirical findings for theoretical models.

To the best of our knowledge, we contribute to the theoretical literature in the following ways. First, the core innovation that we bring to the literature is to uncover “heterogeneous responses" to public information as an important source of disagreement. Second, we found the differential contributions of three distinct information sources to changes in the term structure of inflation expectations. Specifically, the individual long-term beliefs is the main contributor to the long-run disagreement followed by private information, while private information is the main factor driving short-run disagreement. Last, heterogeneous responses to public information are the key driver of increased disagreement about both short-run and long-run inflation forecasts in times of economic downturns or high inflation uncertainty. We review the features of existing theoretical models and identify the features that could be added to make the models account for our empirical findings.

We discuss the literature on the sticky information model, the noisy information model, and disagreement about means and long-run priors. Table F1 summarizes the main features of each model in the context of our empirical findings. Note that there is no scope for disagreement in the full-information rational expectation (FIRE) model, where economic agents are ex ante identical and efficiently process all available information. Therefore, we focus on the remaining models.

To begin with, the importance of individual prior beliefs in long-run disagreement aligns with the findings of Patton and Timmermann (2010) and Farmer et al. (2021), which emphasize the significance of beliefs about long-run means and individual forecasters' priors for their long-term macroeconomic forecasts. It is worth noting that our statistical model captures a forecaster-specific long-run prior through individual fixed effects. In contrast to these studies, we also uncover both private and public information as additional important factors that contribute to disagreement about the long run.

Second, disagreement driven by private information can be accounted for by the noisy information model. The noisy information model Woodford (2001) can generate disagreement among forecasters, but is limited in characterizing countercyclical disagreement or increased disagreement in response to a large shock. In the model, forecasters are ex-ante identical with time-invariant information precision which is the same across forecasters, but they are faced with idiosyncratic signals which are uncorrelated over time. This feature generates time-varying disagreement, but is limited in generating the countercyclicality seen in disagreement attributable to public information. Thus, the noisy information model can account for disagreement attributable to private information but not the portion attributable to public information.

Both the noisy information model with heterogeneity and the sticky information model offer insights into characterizing disagreement driven by heterogeneous responses to public information. To begin with, the noisy information model with heterogeneous information precision across forecasters, as demonstrated by Coibion and Gorodnichenko (2012b), can generate increased disagreement to economic shocks compared to the model without such heterogeneity. Notably, in our statistical model, the differences in factor loadings on the level and slope can be interpreted as reflecting heterogeneous information precision, thereby contributing to disagreement arising from public information.

Additionally, the sticky information model is capable of generating disagreement at all times and capturing increased disagreement in times of large shocks under specific conditions. In this model, economic agents update their information set periodically due to the cost of information acquisition, as explained by Mankiw and Reis (2002). Consequently, disagreement arises because only a fraction of forecasters update their forecasts in response to macroeconomic news, while others do not. While this model can capture increased disagreement in response to macroeconomic news, the increased disagreement tends to dissipate quickly over time as more forecasters update their information set. This feature allows the model to characterize the limited role of public information in normal times, but the model is limited in capturing the increased and persistent importance of public information in disagreement during an economic

Table F1: DISAGREEMENT IN THE MODELS OF EXPECTATION FORMATION

	FIRE	Sticky Information	Noisy info. (Same)	Noisy info. (Different)	Disagreement about means
Scope of disagreement	X	✓	✓	✓	✓
Long-term beliefs (heterogeneity)	X	X	X	X	✓
Changing idiosyncratic disagreement	X	X	✓	X	X
Countercyclical common disagreement	X	✓	X	✓	X
Forecast-horizon differences	X	X	X	X	X

downturn.[37]

Lastly, while existing models can account for certain aspects of our empirical findings, they are constrained in their ability to fully capture the time-varying importance of public and private information, as well as long-run beliefs, in short- and long-term inflation forecasts. Note that the factor loadings on the level and slope factors play a crucial role in generating differential effects of public information on disagreement across forecasting horizons, which is the missing piece in the literature. The two sets of loadings suggest a need for two-dimensional heterogeneity in reactions to news: one related to the long run and another concerning the transition from the near term to the long run. All told, our empirical findings suggest directions for improving existing models of expectation formation and helping them capture the observed patterns of diasgreement.

[37] It's worth noting that if the frequency of economic shocks exceeds the frequency of information updating, the sticky information model can also generate disagreement at all times.

Appendix G Discussion of Modeling Choices

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

G.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.

G.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.

Appendix H Model with Curvature and AR(3) Dynamics

In this section, we present some results from a more generalized model that incorporates the curvature factor and an AR(3) process for factor dynamics. We provide the distribution of smoothed factors, the distribution of forecasts, the dispersion and skewness of forecasts, and the contribution of public information to disagreement.[38] Overall, the estimation results are very similar to those obtained from our baseline two-factor model with AR(1) dynamics.

Figure H17: SMOOTHED FACTOR DISTRIBUTIONS

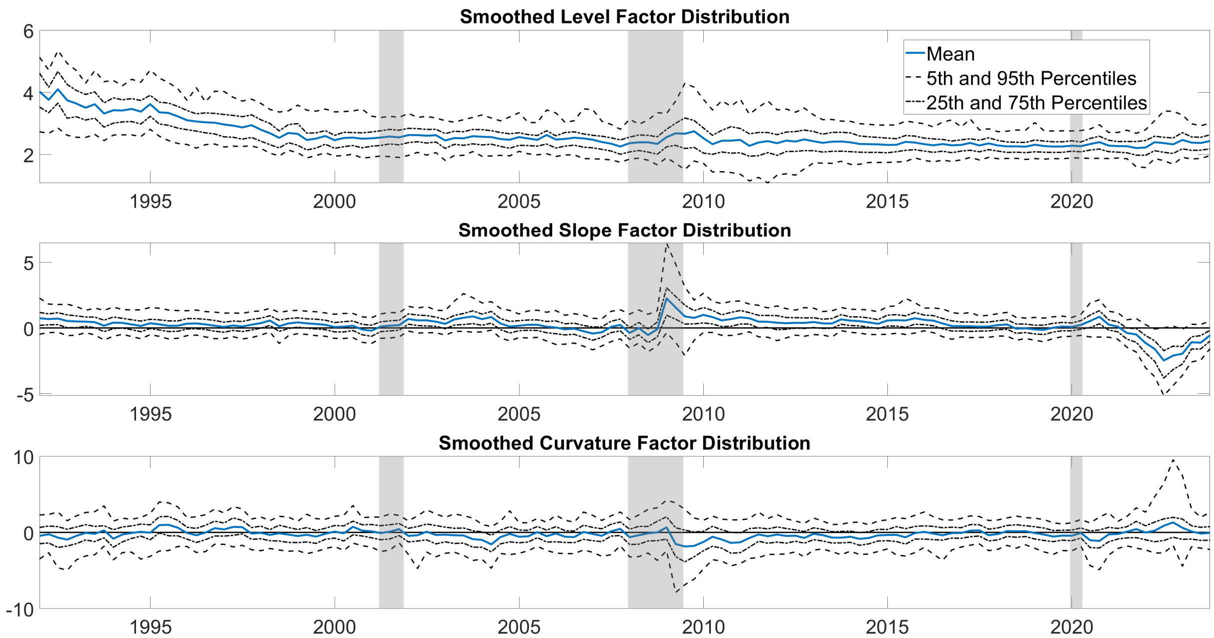

Notes: The figure shows the cross-sectional distributions of the individual level factors (upper panel) and individual slope factors (bottom panel). The solid blue line is the posterior median of the median factor across forecasters. The dashed-dotted lines depict the posterior medians of the 25th and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

Sources: Authors' calculation

[38] Additional results are available upon request.

Figure H18: Distribution of Forecasts

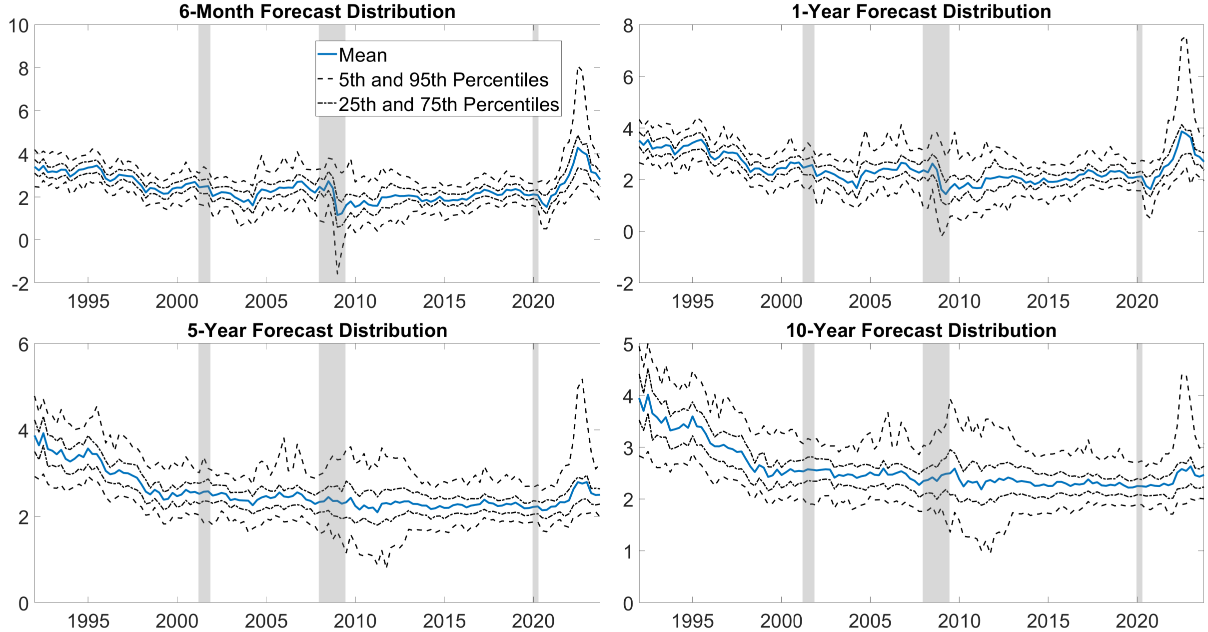

Notes: The figure shows the cross-sectional distribution of individual inflation forecasts at four different forecast horizons. The solid blue line is the posterior median of the mean forecast across forecasters. The dotted lines depict the posterior medians of the 25th and 75th percentiles. The dashed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

Sources: Authors' calculation

Figure H19: Disagreement About and Skewness of Forecasts

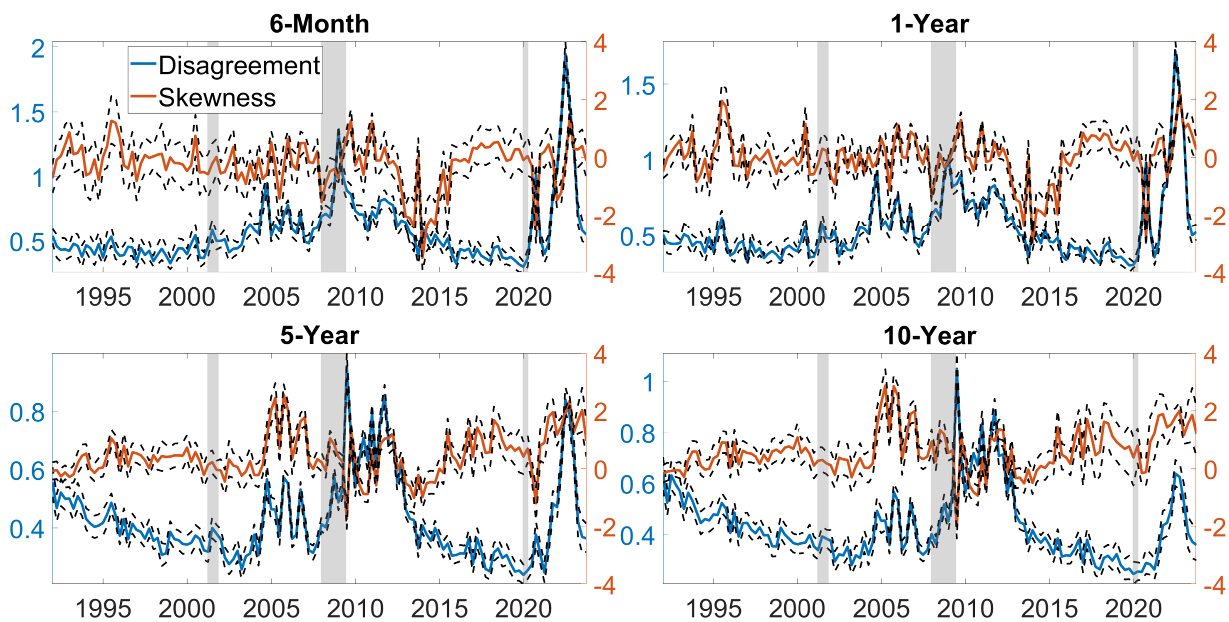

Notes: The figure shows the standard deviation and skewness of individual inflation forecasts at four different forecast horizons. The solid blue line is the posterior median of the disagreement across forecasters. The solid red line is the posterior median of the skewness across forecasters. The dashed lines depict the posterior 5th and 95th percentiles of the disagreement and skewness. The shaded areas denote NBER recessions.

Sources: Authors' calculation

Figure H20: FORECAST VARIANCE DECOMPOSITION

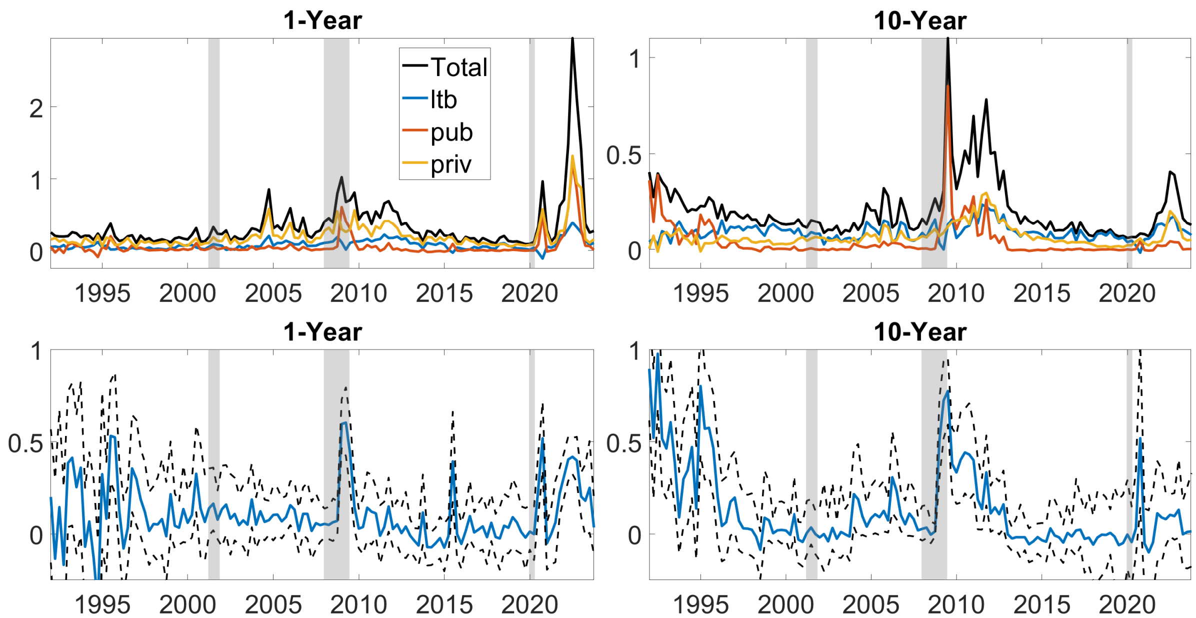

Notes: The top two panels show the decomposition of the cross-sectional variance of inflation forecasts (black line) into the components driven by individual long-term beliefs (denoted by ltb, blue line), heterogeneous responses to public information (denoted by pub, red line), and private information (denoted by priv, yellow line). Each line corresponds to the posterior median. The bottom panels show the variance share of public information, $\beta^{pub}_{h,t}$. The solid blue line corresponds to the posterior median and the dotted black lines correspond to pointwise 95% credible intervals. The left and right columns correspond to 1- and 10-year forecasting horizons respectively. The shaded areas denote NBER recessions. The bottom panels show the variance share of public information.

Sources: Authors' calculation