Online Appendix:

Disagreement About the Term Structure of Inflation Expectations

Hie Joo Ahn[*]　　　　Leland E. Farmer[†]

July 17, 2026

Contents

[*]Ahn: Federal Reserve Board of Governors, hiejoo.ahn@frb.gov. The views expressed here are those of the authors and do not necessarily reflect the views and policies of the Board of Governors or the Federal Reserve System.

[†]Farmer: University of Virginia, lef2u@virginia.edu

A1 Literature Review

Our paper contributes to several literatures on expectation formation and the anchoring of inflation expectations. Recent studies model individual inflation expectations (e.g., Herbst and Winkler, 2026; Fisher, Melosi and Rast, 2025), but none jointly models short- and long-term expectations at the individual level with factors that correspond to the trend and transitory variation in inflation. Herbst and Winkler (2026) and Kim (2023) study short-term expectations only, and Fisher et al. (2025) focus on long-term expectations, modeling forecasters' expectations with a trend-cycle decomposition of SPF data. Research on the term structure of inflation expectations has centered on consensus paths: Aruoba (2020) models the aggregate term structure, and Crump, Eusepi, Moench and Preston (2023) estimate a multivariate trend-cycle model from consensus expectations across surveys, rather than individual responses.

We use a Nelson-Siegel model to parsimoniously characterize the individual-level term structure of inflation expectations. Methodologically, our paper is closest to Herbst and Winkler (2026) and Fisher et al. (2025), as these studies statistically characterize individual-level expectations and provide economic interpretations based on dispersed information models. In contrast to these studies, our model characterizes the complete term structure of inflation expectations by fitting the response of inflation expectations across all forecasting horizons. We also allow for heterogeneous responses to public information, unlike Herbst and Winkler (2026).[1] This lets us assess anchoring through long-run expectations and the transition to the long run, using both the consensus and disagreement, together with the roles of public and private information at each horizon.

Additionally, our paper contributes to the literature on noisy information and learning by providing a coherent empirical framework that disentangles the contributions of long-run beliefs, private information, and public information to distributional changes in inflation expectations across forecasting horizons. Previous studies have examined the roles of public and private information, long-run beliefs, or prior

[1] Fisher et al. (2025) allow for heterogeneous responses to public information.

beliefs as sources of expectation formation (e.g., Patton and Timmermann, 2010; Andrade, Crump, Eusepi and Moench, 2016; Lahiri and Sheng, 2008; Maćkowiak, Matějka and Wiederholt, 2023; Herbst and Winkler, 2026; Farmer, Nakamura and Steinsson, 2024; Fofana, Patzelt and Reis, 2025).[2] While earlier studies have focused on one or a subset of these three information sources individually, we develop a unified statistical framework to evaluate their relative importance in shaping disagreement. To the best of our knowledge, this paper is the first to recover the role of long-run beliefs, public information, and private information, in driving disagreement across forecasting horizons using a formal statistical model.[3]

This paper also contributes to the literature on anchored inflation expectations (e.g., Bundick and Smith, 2025; Carvalho, Eusepi, Moench and Preston, 2023; Castillo-Martinez and Reis, 2026; Coibion and Gorodnichenko, 2025). Reis (2020) develops a parsimonious structural model to characterize discrepancies in inflation expectations between financial markets and households, finding that inflation became increasingly unanchored from 2014 onward. Meanwhile, Bundick and Smith (2025) analyze the impact of monetary policy on inflation expectations, as measured by the Michigan Survey of Consumers and TIPS break-even rates, finding that inflation expectations became more anchored following the announcement of an explicit inflation target in 2012. Bundick and Smith assess the degree of anchoring by examining the path from near-term expectations to long-run expectations only, focusing on whether short-run deviations revert to the long-run target. Coibion and Gorodnichenko (2025) measure the degree of households' anchoring based on the association between short-run inflation expectations and long-run inflation expectations, similar in spirit to Bundick and Smith. Fisher et al. (2025) also studies anchored inflation expectations, but their main focus is on long-run expectations. Unlike these studies, we first demonstrate that disagreement can serve as an independent metric for assessing anchored expectations. Second, our model jointly characterizes the transition from short-term to long-term expectations and the level of long-run expectations, thereby capturing both path and level anchoring. Finally, our model uncovers the role of three distinct information sources in shaping both the path and long-run level of

[2] Lahiri and Sheng (2008) develop a Bayesian learning model in which experts' forecasts reflect long-term beliefs and interpretations of public information, and account for within-forecaster variability and between-forecaster disagreement in GDP projections across horizons. Patton and Timmermann (2010) trace persistent disagreement to heterogeneity in priors and show that differences in opinion move countercyclically. Farmer et al. (2024) document the importance of long-term beliefs in professional forecasters' expectations. Their empirical findings align with ours: private information is the main source of short-term disagreement, and long-term beliefs are the primary driver of long-run disagreement.

[3] We focus on inflation forecasts because our primary interest lies in measuring the degree of anchoring in inflation expectations and assessing its implications for monetary policy. Another reason is that the SPF provides more granular information on forecasting horizons for inflation forecasts, making it particularly well-suited for recovering an individual term structure. By contrast, the available horizon information is more limited for GDP and other macroeconomic variables. Nevertheless, this does not preclude the possibility that forecasts of other macroeconomic variables may also be informative for assessing the anchoring of inflation expectations (see, for example, Andrade et al. (2016)). Extending our analysis to a joint term structure of expectations across multiple macroeconomic variables and decomposing the contributions of the three information sources represents a promising direction for future research.

expectations, identifying which sources enhance or hinder the anchoring process.

Lastly, this paper relates to recent studies on the interaction between monetary policy and disagreement among economic agents (e.g., Andrade et al., 2016; Glas and Hartmann, 2016; Ehrmann, Gaballo, Hoffmann and Strasser, 2019; Barbera, Xia and Zhu, 2023).[4] Ehrmann et al. (2019) show that long-horizon time-contingent and state-contingent forward guidance effectively reduces disagreement, while short-horizon time-contingent forward guidance does not. Falck, Hoffmann and Hürtgen (2021) reveal that a price puzzle arises in states of high disagreement but disappears in states of low disagreement using state-dependent local projections. Dong, Liu, Wang and Wei (2024) empirically demonstrate that inflation disagreement weakens the power of forward guidance and conventional monetary policy, providing a structural model in which households hold heterogeneous beliefs about the inflation target of the central bank. Our research differs from recent studies by further identifying the sources of disagreement and explicitly showing that disagreement attributable to public information is the component that delays the effects of monetary policy and leads to a price puzzle.

A2 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.

A2.1 State Equation

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

$$\mathbf{x}_t := \left[\ L_t, S_t, \varepsilon^L_{1,t}, \varepsilon^S_{1,t}, \cdots, \varepsilon^L_{n,t}, \varepsilon^S_{n,t}\ \right]' \tag{A2.1}$$

The transition matrix and shock vector are

$$\mathbf{F} := \operatorname{diag}\left(\rho_L, \rho_S, \underbrace{\bar{\rho}_L, \bar{\rho}_S, \ldots, \bar{\rho}_L, \bar{\rho}_S}_{n \text{ forecaster blocks}}\right), \qquad \mathbf{u}_t := \left[\ u_{L,t}, u_{S,t}, u^L_{1,t}, u^S_{1,t}, \cdots, u^L_{n,t}, u^S_{n,t}\ \right]', \tag{A2.2}$$

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

with shock covariance matrix

$$\mathbf{Q} := \text{diag}\left(1, 1, \underbrace{\sigma_L^2, \sigma_S^2, \ldots, \sigma_L^2, \sigma_S^2}_{n \text{ forecaster blocks}}\right). \tag{A2.3}$$

Finally, we arrive at the state equation:

$$\mathbf{x}_t = \mathbf{F}\mathbf{x}_{t-1} + \mathbf{u}_t, \quad \mathbf{u}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}). \tag{A2.4}$$

A2.2 Measurement Equation

The measurement equation is of the form:

$$\mathbf{y}_t = \mu_y + \mathbf{H}\mathbf{x}_t + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}), \tag{A2.5}$$

where $\mathbf{y}_t$ stacks 20 survey questions for each of the n forecasters, μ_y is a vector of forecaster fixed effects, $\mathbf{H}$ maps the states to the observed forecasts, and $\mathbf{v}_t$ is a vector of measurement errors with diagonal covariance matrix $\mathbf{R} = \text{diag}(\sigma^2_{v,1}, \ldots, \sigma^2_{v,20}, \ldots, \sigma^2_{v,1}, \ldots, \sigma^2_{v,20})$.

Every row of the system has the same structure. Define the slope loading for a forecast of average inflation between horizons h_1 and h_2 as

$$f_S(h_1, h_2) = \frac{e^{-\lambda h_1} - e^{-\lambda h_2}}{\lambda(h_2 - h_1)}. \tag{A2.6}$$

The row of (A2.5) for question k of forecaster i, with forecast window $[t + h_{1,k},\ t + h_{2,k}]$, is

$$y_{it,k} = \underbrace{\alpha_{L,i} - \alpha_{S,i} f_S(h_{1,k}, h_{2,k})}_{\text{row of } \boldsymbol{\mu}_y} + \underbrace{\beta_{L,i} L_t - \beta_{S,i} f_S(h_{1,k}, h_{2,k})\, S_t + \varepsilon_{L,it} - f_S(h_{1,k}, h_{2,k})\, \varepsilon_{S,it}}_{\text{row of } \mathbf{H}\mathbf{x}_t} + v_{it,k}, \tag{A2.7}$$

with $v_{it,k} \sim N(0, \sigma^2_{v,k})$: each forecaster's rows load on the common factors and on that forecaster's own idiosyncratic factors, with weights 1 and $-f_S(h_{1,k}, h_{2,k})$, and on no other forecaster's states. The questions and their windows are:

k	Forecast object	Survey quarter	$(h_{1,k}, h_{2,k})$
1–4	fixed horizon, h quarters ahead, $h = 1,\ldots,4$	every quarter	$(h-1, h)$
5–8	next calendar year	Q1–Q4	$(4-q, 8-q)$
9–12	second calendar year ahead	Q1–Q4	$(8-q, 12-q)$
13–16	5-year average, adjusted	Q1–Q4	$(0, 20-q)$
17–20	10-year average, adjusted	Q1–Q4	$(0, 40-q)$

Here $q \in \{1,2,3,4\}$ is the quarter of the calendar year in which the survey is conducted. The first four questions are observed every period. Each fixed-event question is observed only in its survey quarter, so 12 of the 20 observables per forecaster are missing by construction in any given quarter.

The reported five- and ten-year forecasts are averages over calendar-year windows that include quarters already realized at the survey date. In survey quarter q, the reported objects are $\pi_{i\,t-q\to t+20-q|t}$ and $\pi_{i\,t-q\to t+40-q|t}$, and we convert them to the forward-looking windows in the table above by removing the realized quarters:

$$\pi_{i\,t\to t+20-q|t} = \frac{4}{20-q}\left(5\,\pi_{i\,t-q\to t+20-q|t} - \frac{1}{4}\sum_{j=1}^{q}\pi_{i\,t-j\to t-j+1|t}\right) \tag{A2.8}$$

$$\pi_{i\,t\to t+40-q|t} = \frac{4}{40-q}\left(10\,\pi_{i\,t-q\to t+40-q|t} - \frac{1}{4}\sum_{j=1}^{q}\pi_{i\,t-j\to t-j+1|t}\right). \tag{A2.9}$$

The $j=1$ term is the reported SPF nowcast and the $j=2$ term the reported SPF backcast; for $j=3,4$, which enter in the third and fourth survey quarters, we use realized inflation from the most recent CPI vintage available at the time of the survey.

A3 Gibbs Sampler

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

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

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

2. Sample $\boldsymbol{\theta}_2 = \rho_S$, the autoregressive coefficient of the common slope factor, conditional on remaining parameters. In the baseline specification the level factor is a random walk, so $\rho_L = 1$ is fixed and not sampled; only in the AR(1) robustness specification of Section A9 is ρ_L drawn in this block. Since L_t and S_t are observed, this is a standard regression model. We impose stationarity through the prior.

3. Sample $\boldsymbol{\theta}_3 = \text{vec}(\boldsymbol{B})$, where $\boldsymbol{B} = \text{diag}(\bar{\rho}_L, \bar{\rho}_S)$ collects the pooled autoregressive coefficients of the idiosyncratic components, conditional on remaining parameters. Since $\varepsilon_{L,it}$ and $\varepsilon_{S,it}$ are observed for every forecaster $i = 1, \ldots, n$, this is a standard multivariate regression model with known covariance matrix where we pool the data across forecasters, using the transitions for which forecaster i reports at least one forecast in the current quarter. We impose stationarity through the prior.

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

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

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

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

A4 Figures

Figure A1 summarizes the raw SPF forecasts. The upper panels show the cross-sectional mean (the consensus) of the 4-quarter-ahead (left) and 10-year-ahead (right) forecasts, together with the projections of two individual forecasters; the bottom panels show the cross-sectional standard deviation at the same horizons. Both short- and long-run forecasts trend down from about 4% to about 2.5% over the first ten to fifteen years of the sample, short-run forecasts are more volatile because they respond more to transitory shocks, and short-run dispersion typically exceeds long-run dispersion, with the early 1990s a notable exception. Figure A2 plots the reported forecasts of the same two forecasters across all horizons at three dates (2007Q3, 2009Q2, and 2022Q2), against the consensus (black line) and the shaded band between the 5th and 95th percentiles of the cross-sectional distribution. The three dates illustrate the variety of shapes the term structure of inflation expectations takes, and how far an individual forecaster can sit from the consensus in both level and trajectory.

Figure A1: SPF Forecast Summary Statistics and Examples

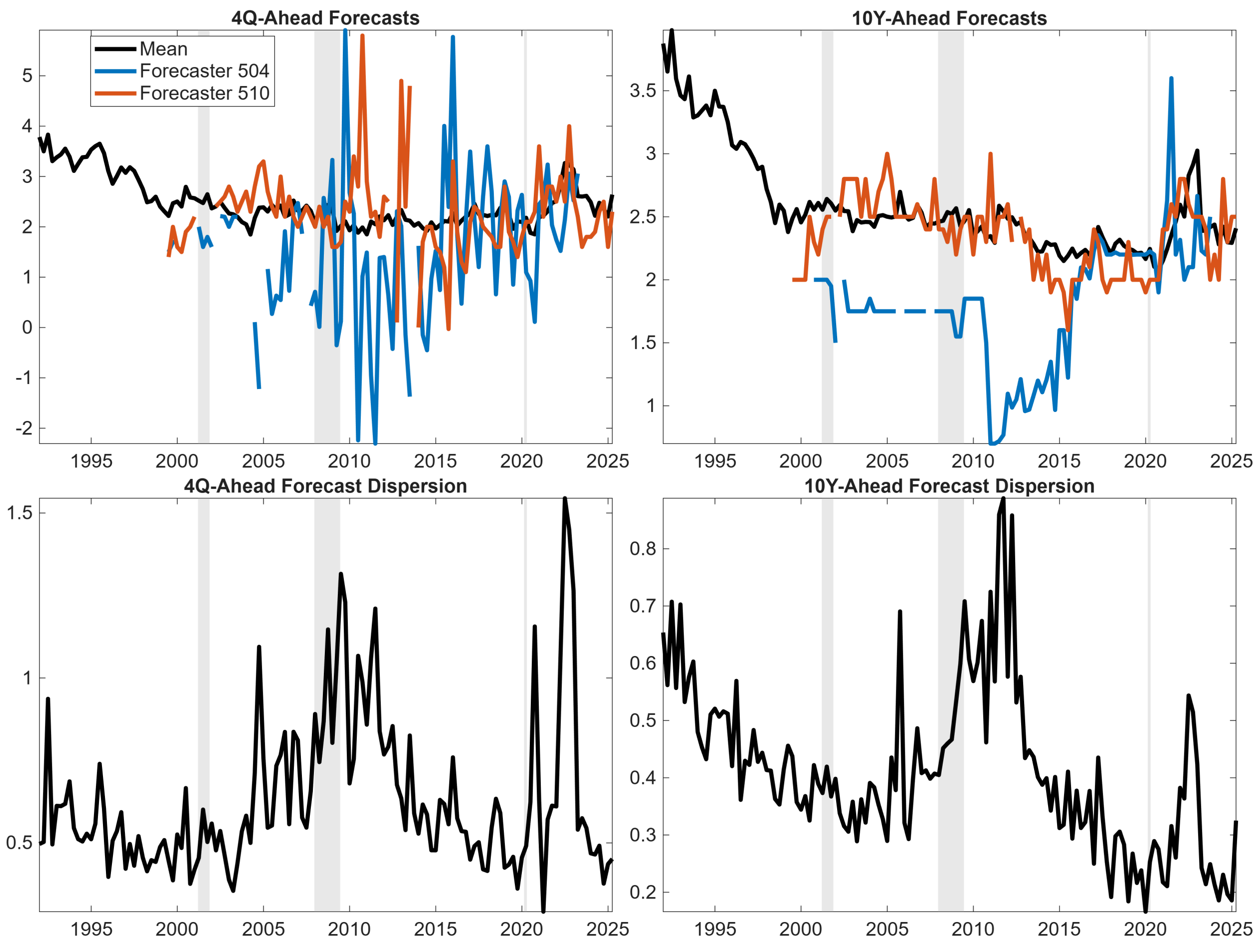

Notes: The black lines in the upper two panels show the average of inflation forecasts 4 quarters (left panel) and 10 years (right panel) ahead. The colored lines are the forecasts of two forecasters whose IDs are 504 and 510 in the survey. The bottom two panels report the cross-sectional standard deviation of the forecasts.

Sources: SPF and authors' calculation.

Figure A2: Three Term Structures of Observed SPF Forecasts

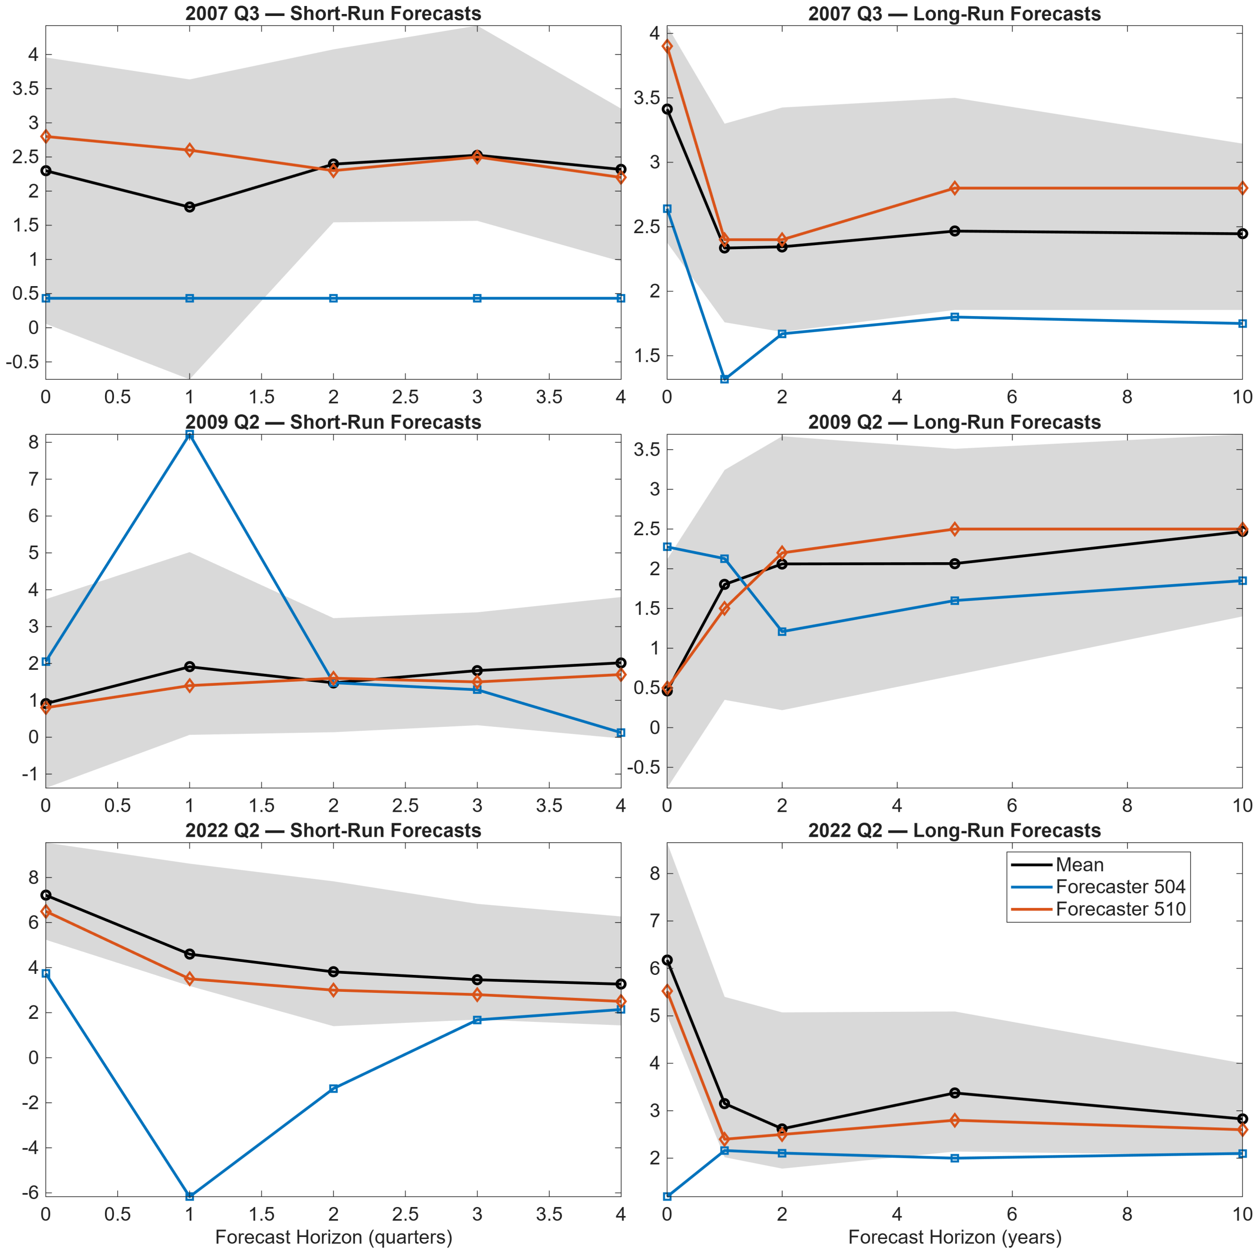

Notes: Each row plots a snapshot of SPF inflation forecasts at the origin date noted in the panel titles (2007Q3, 2009Q2, and 2022Q2). The left column shows short-horizon forecasts; the x-axis is the forecast horizon in quarters (0 to 4). The right column shows long-horizon forecasts; the x-axis is the forecast horizon in years (0 to 10), interpolated linearly between the SPF-observed horizons of 0, 1, 2, 5, and 10 years ahead. The black lines are consensus forecasts (cross-sectional means); the colored lines are the forecasts of two individual forecasters with IDs 504 and 510. The shaded gray bands span the 5th to 95th percentiles of the cross-sectional distribution at each horizon.

Sources: SPF and authors' calculation.

A5 Noisy Information Model Proofs

This section proves Proposition 1 of the main text. Section A5.1 fixes notation and states the maintained assumption. Section A5.2 proves three lemmas: the moving-average representation of the nowcast, a monotonicity property of truncated geometric sums, and a second-order moment bound. Section A5.3 proves the proposition. Every approximation carries an exact error bound with explicit constants.

A5.1 Setup and Notation

The trend and transitory components of the noisy information model have identical scalar structure, because every system matrix in Section 4.1 of the main text is diagonal. We therefore work with a generic scalar component and specialize at the end.

Generic component. Let ξ_t denote a scalar state with true law

$$\xi_t = a\xi_{t-1} + \sigma\varepsilon_t, \qquad \varepsilon_t \overset{\text{iid}}{\sim} N(0,1), \qquad a \in (0,1], \quad \sigma > 0. \tag{A5.1}$$

Forecaster $i \in \{1,\ldots,N\}$ perceives the persistence $a_i \in (0,1]$ and observes the signal

$$s_{it} = \xi_t + \theta z_t + \omega_i w_{it}, \tag{A5.2}$$

where $z_t \overset{\text{iid}}{\sim} N(0,1)$ is public noise, $w_{it} \overset{\text{iid}}{\sim} N(0,1)$ is private noise, $\theta \geq 0$, $\omega_i \geq 0$, and the families $\{\varepsilon_t\}$, $\{z_t\}$, $\{w_{1t}\},\ldots,\{w_{Nt}\}$ are mutually independent. The forecaster filters with the constant gain $g_i \in (0,1]$ and holds the long-term belief b_i, so that, as in Section 4.1 of the main text,

$$\mathbb{F}_{it} = b_i + \mathbb{F}^*_{it}, \qquad \mathbb{F}^*_{it} = g_i s_{it} + (1-g_i)\, a_i\, \mathbb{F}^*_{it-1}. \tag{A5.3}$$

The trend and transitory components correspond to the parameter values

$$(\xi_t, a, a_i, g_i, \sigma, \theta, \omega_i, b_i) = \begin{cases} (\tau_t,\, 1,\, 1,\, g_{i,\tau},\, \sigma_\tau,\, \theta_\tau,\, \omega_{i,\tau},\, b_i^\tau) & \text{trend} \\ (c_t,\, \rho,\, \rho_i,\, g_{i,c},\, \sigma_c,\, \theta_c,\, \omega_{i,c},\, b_i^c) & \text{transitory.} \end{cases}$$

Cross-sectional moments. For a collection $\{x_i\}_{i=1}^N$, write $\mathbb{E}_i[x] := N^{-1}\sum_{i=1}^N x_i$, $\text{Var}_i(x) := \mathbb{E}_i\left[(x-\mathbb{E}_i[x])^2\right]$, and $\text{Cov}_i(x,y) := \mathbb{E}_i\left[(x-\mathbb{E}_i[x])(y-\mathbb{E}_i[y])\right]$.

Belief persistence and benchmark parameters. Define, for each forecaster,

$$\varphi_i := (1-g_i)\,a_i, \qquad \kappa_i := \varphi_i/a, \tag{A5.4}$$

and the benchmark values

$$g := \mathbb{E}_i[g_i], \qquad \bar{a} := \mathbb{E}_i[a_i], \qquad \varphi := (1-g)\,\bar{a}, \qquad \kappa := \varphi/a. \tag{A5.5}$$

The benchmark φ is the belief persistence evaluated at the average gain and the average perceived persistence. It is not in general the cross-sectional average of φ_i: expanding the product gives the exact identity

$$\mathbb{E}_i[\varphi_i] = \bar{a} - \mathbb{E}_i[g_i a_i] = (1-g)\bar{a} - \mathrm{Cov}_i(g_i, a_i) = \varphi - \mathrm{Cov}_i(g_i, a_i), \tag{A5.6}$$

and $\left|\mathrm{Cov}_i(g_i, a_i)\right| \le \mathrm{Var}_i(g_i)^{1/2}\,\mathrm{Var}_i(a_i)^{1/2}$ by the Cauchy–Schwarz inequality. For the trend, $a_i = 1$ for every i, so $\mathrm{Cov}_i(g_i, a_i) = 0$ and $\varphi = \mathbb{E}_i[\varphi_i] = \mathbb{E}_i[\rho_{L,i}] = \bar{\rho}_L = 1 - g_\tau$ exactly. For the transitory component, $\varphi = (1-g_c)\rho_c$, while $\mathbb{E}_i[\varphi_i] = \mathbb{E}_i[\rho_{S,i}] = \bar{\rho}_S$; the two differ by $\mathrm{Cov}_i(g_{i,c}, \rho_i)$, which (A5.6) bounds by $\mathrm{Var}_i(g_{i,c})^{1/2}\mathrm{Var}_i(\rho_i)^{1/2}$.

Assumption 1. *(i) For every i: $g_i \in (0,1]$, $a_i \in (0,1]$, and $\omega_i \ge 0$; hence $\varphi_i = (1-g_i)a_i \in [0,1)$ and $\varphi \in [0,1)$. (ii) Define $\varphi_{\max} := \max\{\varphi, \max_i \varphi_i\} < 1$ and, when $a < 1$, $m := \max(a, \varphi_{\max}) < 1$. (iii) Results that additionally require $\kappa_i < 1$ for every i, that is, $\varphi_i < a$, state the condition explicitly; for the trend ($a = 1$) it holds automatically by (i).*

Initialization. Fix an initial date $T_0 < t$ and set $\xi_{T_0-1} = 0$ and $\mathbb{F}^*_{i,T_0-1} = 0$ for every i. All representations below are then finite sums over lags $k = 0, \ldots, K$ with $K := t - T_0$, and all identities and bounds hold for every K, uniformly in K. To lighten notation we write $\sum_{k\ge 0}$, with the convention that every sum terminates at $k = K$. Weight sequences that are geometrically bounded (all private- and public-noise weights, and the fundamental-shock weights when $a < 1$) converge in L^2 as $T_0 \to -\infty$, and second-moment statements refer to that limit. When $a = 1$ the fundamental-shock weights do not vanish at long lags; statements about the trend are exact finite-sum identities at a fixed T_0, and none of the bounds depend on T_0.

A5.2 Auxiliary Lemmas

Lemma 1 (Nowcast representation)**.** *Under Assumption 1, for every i and t,*

$$\mathbb{F}_{it} = \underbrace{b_i}_{\textit{long-term belief}} + \underbrace{\sum_{k\geq 0} \lambda_{ik}\,\varepsilon_{t-k} + \sum_{k\geq 0} \delta_{ik}\,z_{t-k}}_{\textit{common component } \mathbb{F}^C_{it}} + \underbrace{\sum_{k\geq 0} g_i\omega_i\,\varphi_i^k\,w_{it-k}}_{\textit{idiosyncratic component } v_{it}}, \tag{A5.7}$$

with weights

$$\lambda_{ik} = \sum_{j=0}^{k} \varphi_i^j\,g_i\sigma\,a^{k-j} = g_i\sigma\,a^k\,h_k(\kappa_i), \qquad \delta_{ik} = g_i\theta\,\varphi_i^k, \qquad h_k(x) := \sum_{j=0}^{k} x^j. \tag{A5.8}$$

Proof. Solving the linear recursion (A5.3) backward from T_0, with $\mathbb{F}^*_{i,T_0-1} = 0$,

$$\mathbb{F}^*_{it} = \sum_{k\geq 0} \varphi_i^k\,g_i\,s_{it-k}.$$

Substituting (A5.2) for s_{it-k} produces the public- and private-noise sums directly, with weights $g_i\theta\varphi_i^k$ and $g_i\omega_i\varphi_i^k$. For the fundamental shocks, the backward solution of (A5.1) from $\xi_{T_0-1} = 0$ is $\xi_{t-k} = \sum_{m\geq 0}\sigma a^m \varepsilon_{t-k-m}$, so the coefficient on ε_{t-k} collects all pairs of indices summing to k:

$$\lambda_{ik} = \sum_{j=0}^{k} \varphi_i^j\,g_i\,\sigma\,a^{k-j} = g_i\sigma a^k \sum_{j=0}^{k} (\varphi_i/a)^j = g_i\sigma\,a^k h_k(\kappa_i).$$

This establishes (A5.8). $\square$

The public- and private-noise sums in (A5.7) are the processes $\boldsymbol{u}_{it}$ and $\boldsymbol{v}_{it}$ of Section 4.1 of the main text, componentwise. The closed forms referenced there are the specializations

$$\lambda^\tau_{ik} = \sigma_\tau\left[1 - (1-g_{i,\tau})^{k+1}\right], \qquad \lambda^c_{ik} = g_{i,c}\sigma_c\,\rho^k\,h_k(\kappa_i) = g_{i,c}\sigma_c\,\rho^k\,\frac{1-\kappa_i^{k+1}}{1-\kappa_i}, \tag{A5.9}$$

where the trend form uses $g_{i,\tau}h_k(1-g_{i,\tau}) = 1-(1-g_{i,\tau})^{k+1}$ and the last expression is a finite-sum identity, equal to $(k+1)g_{i,c}\sigma_c\rho^k$ at $\kappa_i = 1$. For $a < 1$ the weights decay geometrically at rate $\max(a,\varphi_i)$; for the trend they increase to σ_τ.

Lemma 2 (Truncated geometric sums)**.** *Let $x, y \geq 0$ and $R_k := h_k(x)/h_k(y)$. Then: (i) $R_0 = 1$; (ii) R_k is nondecreasing in k if $x \geq y$ and nonincreasing in k if $x \leq y$; (iii) if $x, y \in [0,1)$, then $\lim_{k\to\infty} R_k = (1-y)/(1-x)$*

and

$$\left|\frac{h_k(x)}{h_k(y)}-1\right| \le \frac{|x-y|}{1-x} \qquad \textit{for every } k \ge 0.$$

Proof. (i) $h_0 \equiv 1$. (ii) The sign of $R_{k+1}-R_k$ is the sign of

$$h_{k+1}(x)h_k(y)-h_k(x)h_{k+1}(y) = x^{k+1}h_k(y)-y^{k+1}h_k(x) = \sum_{j=0}^{k} x^j y^j \left(x^{k+1-j}-y^{k+1-j}\right),$$

where the first equality subtracts the common term $h_k(x)h_k(y)$ and the second pairs the summands of $x^{k+1}h_k(y)$ and $y^{k+1}h_k(x)$ index by index. Every summand has the sign of $x-y$. (iii) For $x \in [0,1)$, $h_k(x) \to (1-x)^{-1}$, so $R_k \to (1-y)/(1-x)$. By (i) and (ii), R_k lies between $R_0 = 1$ and this limit for every k, so $|R_k-1| \le \left|\frac{1-y}{1-x}-1\right| = \frac{|x-y|}{1-x}$. □

Lemma 3 (Second-order moment bound)**.** *Let x be a random variable with values in $[0,q]$, $0 \le q < 1$, let $\bar{x} := \mathbb{E}[x]$, and let $n \ge 1$ be an integer. Then*

$$0 \;\le\; \mathbb{E}[x^n]-\bar{x}^n \;\le\; \frac{n(n-1)}{2}\, q^{n-2}\mathrm{Var}(x) \;\le\; \frac{\mathrm{Var}(x)}{(1-q)^3}.$$

The same holds with $\mathbb{E} = \mathbb{E}_i$, the empirical cross-sectional distribution.

Proof. For $u, v \ge 0$, applying the factorization $u^n - v^n = (u-v)\sum_{j=0}^{n-1} u^j v^{n-1-j}$ twice gives the exact identity

$$u^n - v^n - nv^{n-1}(u-v) = (u-v)\sum_{j=0}^{n-1} v^{n-1-j}\left(u^j - v^j\right) = (u-v)^2 \sum_{j=1}^{n-1}\sum_{l=0}^{j-1} u^l\, v^{n-2-l}.$$

The double sum has $n(n-1)/2$ terms, each a monomial of total degree $n-2$ in (u,v) and hence contained in $[0, q^{n-2}]$ when $u, v \in [0,q]$. Set $u = x$ and $v = \bar{x}$ and take expectations; the linear term vanishes, so $\mathbb{E}[x^n]-\bar{x}^n = \mathbb{E}\left[(x-\bar{x})^2 \cdot (\text{double sum})\right] \in \left[0, \frac{n(n-1)}{2} q^{n-2}\mathrm{Var}(x)\right]$. Finally, $n(n-1)q^{n-2} \le \sum_{j\ge 2} j(j-1)q^{j-2} = 2/(1-q)^3$. □

A5.3 Proof of Proposition 1

Proof of Proposition 1. Throughout, "exact" means the statement involves no approximation. Steps 1–3 treat the generic component; Steps 4 and 5 specialize to the trend and the transitory component; Step 6 collects the mapping.

Step 1: forecaster-specific entries (exact).

(i) *Idiosyncratic component.* By Lemma 1, v_{it} satisfies $v_{it} = \varphi_i v_{it-1} + g_i \omega_i w_{it}$: an AR(1) process with persistence φ_i and innovation standard deviation $g_i \omega_i$, independent across forecasters because the private noise is. Substituting the trend and transitory parameters gives, exactly,

$$\rho_{L,i} = 1 - g_{i,\tau}, \quad \sigma_{L,i} = g_{i,\tau}\omega_{i,\tau}, \quad \rho_{S,i} = (1 - g_{i,c})\rho_i, \quad \sigma_{S,i} = g_{i,c}\omega_{i,c}. \tag{A5.10}$$

(ii) *Loadings.* By (A5.8), $\lambda_{i0} = g_i \sigma$ and $\delta_{i0} = g_i \theta$, so the date-t innovation of $\mathbb{F}^C_{it}$ is $g_i(\sigma \varepsilon_t + \theta z_t)$. Define $u_t := (\sigma \varepsilon_t + \theta z_t)/(\sigma^2 + \theta^2)^{1/2}$, which is iid $N(0,1)$. The statistical model normalizes the common-factor innovation to unit variance, so the loading is the coefficient on u_t,

$$\beta_i = g_i \left(\sigma^2 + \theta^2\right)^{1/2}, \tag{A5.11}$$

with no approximation. The trend and transitory substitutions give $\beta_{L,i} = g_{i,\tau}(\sigma_\tau^2 + \theta_\tau^2)^{1/2}$ and $\beta_{S,i} = g_{i,c}(\sigma_c^2 + \theta_c^2)^{1/2}$. Bai and Ng (2013) show that separating a common factor from its loadings requires a scale restriction on the factor and a restriction on the loadings. The unit variance of u_t provides the first; averaging (A5.11) across forecasters provides the second, $\mathbb{E}_i[\beta_i] = g(\sigma^2 + \theta^2)^{1/2}$, which ties the average loading to the average gain.

(iii) *Fixed effects.* The constant b_i in (A5.7) is absorbed by forecaster i's fixed effect. The trend gives $\alpha_{L,i} = b_i^\tau$; the slope carries a sign flip, handled in Step 5(i).

Step 2: the benchmark factor and the residual. Evaluate the weight profiles (A5.8) at the benchmark parameters (A5.5) and define

$$\tilde{\lambda}_k := g\sigma\, a^k h_k(\kappa), \qquad \tilde{\delta}_k := g\theta\, \varphi^k, \qquad \tilde{F}_t := \sum_{k\geq 0} \tilde{\lambda}_k \varepsilon_{t-k} + \sum_{k\geq 0} \tilde{\delta}_k z_{t-k}, \qquad \tilde{\beta}_i := \frac{g_i}{g}. \tag{A5.12}$$

The one-factor representation replaces $\mathbb{F}^C_{it}$ with $\tilde{\beta}_i \tilde{F}_t$; after the unit-variance normalization of the date-t innovation of $\tilde{F}_t$, whose standard deviation is $g(\sigma^2 + \theta^2)^{1/2}$, the product $\tilde{\beta}_i \tilde{F}_t$ equals β_i times the normalized factor, consistent with (A5.11). Steps 4 and 5 characterize the dynamics of $\tilde{F}_t$. Define the residual

$$r_{it} := \mathbb{F}^C_{it} - \tilde{\beta}_i \tilde{F}_t = \sum_{k\geq 0} \left(\lambda_{ik} - \tilde{\beta}_i \tilde{\lambda}_k\right) \varepsilon_{t-k} + \sum_{k\geq 0} \left(\delta_{ik} - \tilde{\beta}_i \tilde{\delta}_k\right) z_{t-k}. \tag{A5.13}$$

Step 3: properties of the residual.

(a) *No date-t innovation (exact).* $\lambda_{i0} - \tilde{\beta}_i \tilde{\lambda}_0 = g_i \sigma - (g_i/g)\, g\sigma = 0$, and likewise $\delta_{i0} - \tilde{\beta}_i \tilde{\delta}_0 = 0$. The residual loads only on shocks dated $t-1$ and earlier. The loadings (A5.11), and with them Corollaries 1 and 2 of the main text, therefore involve no approximation.

(b) *Fundamental weights: exact ratio, and a uniform bound when $\kappa_i, \kappa < 1$.* Since $\tilde{\beta}_i \tilde{\lambda}_k = g_i \sigma a^k h_k(\kappa)$, equation (A5.8) gives, exactly and for every k,

$$\frac{\lambda_{ik}}{\tilde{\beta}_i \tilde{\lambda}_k} = \frac{h_k(\kappa_i)}{h_k(\kappa)}. \tag{A5.14}$$

If $\kappa_i < 1$ and $\kappa < 1$, Lemma 2(iii) applied to (A5.14) yields the uniform lag-by-lag bound

$$\left|\lambda_{ik} - \tilde{\beta}_i \tilde{\lambda}_k\right| \le \frac{|\kappa_i - \kappa|}{1 - \kappa_i}\, \tilde{\beta}_i \tilde{\lambda}_k \qquad \text{for every } k. \tag{A5.15}$$

For the trend, $\kappa_i = 1 - g_{i,\tau}$ and $\kappa = 1 - g_\tau$ are below one by Assumption 1(i), and the relative bound in (A5.15) equals $|g_{i,\tau} - g_\tau| / g_{i,\tau}$: at every lag, forecaster i's relative approximation error is at most their relative gain deviation. For the transitory component, $\kappa_i < 1$ is the requirement $\rho_{S,i} < \rho$ that belief persistence lie below the true persistence of the cycle, and the relative bound equals $|\rho_{S,i} - (1 - g_c)\rho_c| / (\rho - \rho_{S,i})$; by (A5.6), replacing the centering value $(1 - g_c)\rho_c$ with $\bar{\rho}_S$ changes the numerator by at most $\mathrm{Var}_i(g_{i,c})^{1/2} \mathrm{Var}_i(\rho_i)^{1/2}$.

(c) *Public-noise weights (L^2 bound, no condition on κ_i).* By (A5.8) and (A5.12), $\delta_{ik} - \tilde{\beta}_i \tilde{\delta}_k = g_i \theta\, (\varphi_i^k - \varphi^k)$. The mean value theorem applied to $x \mapsto x^k$ on $[0, \varphi_{\max}]$ gives $|\varphi_i^k - \varphi^k| \le k \varphi_{\max}^{k-1} |\varphi_i - \varphi|$, so, using $\sum_{k \ge 1} k^2 y^{k-1} = (1 + y)/(1 - y)^3$ with $y = \varphi_{\max}^2$,

$$\sum_{k \ge 0} \left(\delta_{ik} - \tilde{\beta}_i \tilde{\delta}_k\right)^2 \le g_i^2 \theta^2 \left(\varphi_i - \varphi\right)^2 \frac{1 + \varphi_{\max}^2}{\left(1 - \varphi_{\max}^2\right)^3}. \tag{A5.16}$$

(d) *Fundamental weights: L^2 bound without $\kappa_i < 1$ (requires $a < 1$).* Using $h_k(\kappa_i) - h_k(\kappa) = \sum_{j=1}^k (\kappa_i^j - \kappa^j)$, the mean value theorem with $\kappa^* := \max(\kappa_i, \kappa)$, and $a^k\, j\, (\kappa^*)^{j-1} = j\, a^{k-j+1} (a\kappa^*)^{j-1} \le j\, m^k$ (since $a\kappa^* = \max(\varphi_i, \varphi) \le \varphi_{\max}$ and $m = \max(a, \varphi_{\max})$),

$$\left|\lambda_{ik} - \tilde{\beta}_i \tilde{\lambda}_k\right| = g_i \sigma\, a^k\, |h_k(\kappa_i) - h_k(\kappa)| \le g_i \sigma\, |\kappa_i - \kappa|\, \frac{k(k+1)}{2}\, m^k.$$

Since $[k(k+1)/2]^2 \le k^4$ for $k \ge 1$ and $\sum_{k \ge 1} k^4 y^k \le 24/(1 - y)^5$ for $y \in [0, 1)$,

$$\sum_{k \ge 0} \left(\lambda_{ik} - \tilde{\beta}_i \tilde{\lambda}_k\right)^2 \le \frac{24\, g_i^2 \sigma^2\, (\kappa_i - \kappa)^2}{\left(1 - m^2\right)^5}. \tag{A5.17}$$

Summing (A5.16) and (A5.17): when $a < 1$, $\mathrm{Var}(r_{it})$ is bounded by the squared deviations $(\kappa_i - \kappa)^2$ and $(\varphi_i - \varphi)^2$ times constants that depend only on m. Since the date-t innovation variance of $\tilde{\beta}_i \tilde{F}_t$ alone is

$g_i^2(\sigma^2+\theta^2)$, the variance of the residual relative to the variance of the approximating component is of second order in the deviations. The residual's weights are of first order in the deviations, by (A5.15) and the two displays above; no condition on the sign of $\kappa_i - 1$ is needed here, because for $a<1$ every weight sequence decays geometrically at rate $\max(a,\varphi_i) \le m < 1$ regardless.

(e) *Consensus residual (second order, exact bound).* Since $\mathbb{E}_i[\tilde{\beta}_i] = 1$, the cross-sectional average of (A5.13) has lag-k weights $\mathbb{E}_i[\lambda_{ik}] - \tilde{\lambda}_k$ on ε_{t-k} and $\mathbb{E}_i[\delta_{ik}] - \tilde{\delta}_k$ on z_{t-k}. Both rest on the quantity $\Delta_j := \mathbb{E}_i[g_i\varphi_i^j] - g\varphi^j$: directly, $\mathbb{E}_i[\delta_{ik}] - \tilde{\delta}_k = \theta\Delta_k$, and from the middle expression in (A5.8), $\mathbb{E}_i[\lambda_{ik}] - \tilde{\lambda}_k = \sigma\sum_{j=0}^{k} a^{k-j}\Delta_j$. Note $\Delta_0 = 0$. For $j \ge 1$, decompose exactly, with $\bar{\varphi} := \mathbb{E}_i[\varphi_i]$:

$$\Delta_j = \mathrm{Cov}_i\Big(g_i, \varphi_i^j\Big) + g\left(\mathbb{E}_i[\varphi_i^j] - \bar{\varphi}^j\right) + g\left(\bar{\varphi}^j - \varphi^j\right). \tag{A5.18}$$

The three terms are bounded as follows, using $g \le 1$ throughout. First, by Cauchy–Schwarz and the mean value theorem, $|\mathrm{Cov}_i(g_i,\varphi_i^j)| \le \mathrm{Var}_i(g_i)^{1/2}\mathrm{Var}_i(\varphi_i^j)^{1/2} \le j\varphi_{\max}^{j-1}\mathrm{Var}_i(g_i)^{1/2}\mathrm{Var}_i(\varphi_i)^{1/2}$. Second, by Lemma 3 with $q = \varphi_{\max}$, $0 \le \mathbb{E}_i[\varphi_i^j] - \bar{\varphi}^j \le \frac{j(j-1)}{2}\varphi_{\max}^{j-2}\mathrm{Var}_i(\varphi_i)$. Third, by (A5.6) and the mean value theorem, $|\bar{\varphi}^j - \varphi^j| \le j\varphi_{\max}^{j-1}\big|\mathrm{Cov}_i(g_i,a_i)\big| \le j\varphi_{\max}^{j-1}\mathrm{Var}_i(g_i)^{1/2}\mathrm{Var}_i(a_i)^{1/2}$. Summing over j with $\sum_{j\ge1} j\varphi_{\max}^{j-1} = (1-\varphi_{\max})^{-2}$ and $\sum_{j\ge1}\frac{j(j-1)}{2}\varphi_{\max}^{j-2} = (1-\varphi_{\max})^{-3}$, and using $a \le 1$,

$$\left|\mathbb{E}_i[\lambda_{ik}] - \tilde{\lambda}_k\right| \le \sigma\left[\frac{\mathrm{Var}_i(g_i)^{1/2}\left(\mathrm{Var}_i(\varphi_i)^{1/2} + \mathrm{Var}_i(a_i)^{1/2}\right)}{(1-\varphi_{\max})^2} + \frac{\mathrm{Var}_i(\varphi_i)}{(1-\varphi_{\max})^3}\right] \tag{A5.19}$$

uniformly in k, and $|\mathbb{E}_i[\delta_{ik}] - \tilde{\delta}_k| = \theta|\Delta_k|$ obeys the same bound with σ replaced by θ (and decays geometrically in k). Every term on the right of (A5.19) is a product of two cross-sectional deviations. The residual's weights are of first order in the deviations; their cross-sectional average is of second order. This is the precise sense in which the residual averages out of the consensus forecast. The constants are controlled by the distance of the belief persistences from a unit root, the same object restricted by the clustering condition of the proposition.

Step 4: the trend ($a = 1$). The entries $\alpha_{L,i} = b_i^\tau$, $\beta_{L,i} = g_{i,\tau}(\sigma_\tau^2 + \theta_\tau^2)^{1/2}$, $\rho_{L,i} = 1 - g_{i,\tau}$, and $\sigma_{L,i} = g_{i,\tau}\omega_{i,\tau}$ follow from Step 1. It remains to establish $\rho_L = 1$ and to characterize the trend residual. Write $\varphi = 1 - g_\tau$ throughout this step.

(i) *The benchmark level factor is integrated of order one.* With $a = 1$, the benchmark weights in (A5.12) are $\tilde{\lambda}_k = \sigma_\tau(1-\varphi^{k+1})$ and $\tilde{\delta}_k = g_\tau\theta_\tau\varphi^k$. Splitting $\tilde{\lambda}_k$ into σ_τ and $-\sigma_\tau\varphi^{k+1}$ gives the exact decomposition

$$\tilde{F}_t^\tau = \tau_t + \eta_t, \qquad \eta_t := -\sigma_\tau\sum_{k\ge0}\varphi^{k+1}\varepsilon_{t-k} + g_\tau\theta_\tau\sum_{k\ge0}\varphi^k z_{t-k}, \tag{A5.20}$$

where $\tau_t = \sigma_\tau \sum_{k\geq 0} \varepsilon_{t-k}$ and η_t satisfies the AR(1) recursion $\eta_t = \varphi\eta_{t-1} - \varphi\sigma_\tau\varepsilon_t + g_\tau\theta_\tau z_t$, verified by matching moving-average coefficients. Differencing (A5.20) and applying $(1-\varphi L)$ to both sides, the ε_{t-1} terms cancel and $\sigma_\tau(1-\varphi) = g_\tau\sigma_\tau$, leaving

$$(1-\varphi L)\,\Delta\tilde{F}^\tau_t = g_\tau\sigma_\tau\,\varepsilon_t + g_\tau\theta_\tau\,(z_t - z_{t-1})\,. \tag{A5.21}$$

The increments of $\tilde{F}^\tau_t$ are therefore an exact ARMA(1,1) in the two common shocks, and $\tilde{F}^\tau_t$ is integrated of order one. Setting $\rho_L = 1$ matches this order of integration exactly; this is the sense in which the common level factor inherits the trend's unit root. The random-walk specification of L_t additionally treats ΔL_t as white noise, while (A5.21) has an autoregressive root $\varphi = 1 - g_\tau$ and a first-order moving-average part in z. This is an approximation to the time dynamics of the factor, distinct from the cross-sectional approximation bounded in Step 3. Normalizing the date-t innovation $g_\tau(\sigma_\tau\varepsilon_t + \theta_\tau z_t)$ to unit variance reproduces $\beta_{L,i}$ from Step 1(ii).

(ii) *The trend residual.* Substituting $\lambda_{ik} = \sigma_\tau(1-\varphi_i^{k+1})$ and $\tilde{\beta}_i\tilde{\lambda}_k = (g_{i,\tau}/g_\tau)\sigma_\tau(1-\varphi^{k+1})$ into (A5.13) and separating the constant part of each weight,

$$r^\tau_{it} = -\frac{g_{i,\tau}-g_\tau}{g_\tau}\,\tau_t + \tilde{r}_{it}, \qquad \tilde{r}_{it} \coloneqq \sum_{k\geq 0}\left[\frac{g_{i,\tau}}{g_\tau}\varphi^{k+1} - \varphi_i^{k+1}\right]\sigma_\tau\,\varepsilon_{t-k} + g_{i,\tau}\theta_\tau\sum_{k\geq 0}\left(\varphi_i^k - \varphi^k\right)z_{t-k}, \tag{A5.22}$$

where every weight of $\tilde{r}_{it}$ decays geometrically and vanishes at $g_{i,\tau} = g_\tau$, so each weight is of first order in $(g_{i,\tau} - g_\tau)$ by the mean value theorem. The unit-root part of the residual is the true trend scaled by the gain deviation, and $\mathbb{E}_i[g_{i,\tau} - g_\tau] = 0$, so it cancels from the consensus exactly. Its presence means that at long lags the constant-loading factor model overstates disagreement about the trend that originates in fundamental shocks: in the structural model this disagreement dies out because every forecaster's weight λ_{ik} converges to σ_τ as the filter absorbs a permanent shock, while the factor model keeps the responses proportional to the loadings. Per shock, the discrepancy is bounded by $\sigma_\tau|g_{i,\tau} - g_\tau|/g_\tau$; it does not affect the date-t innovation, by Step 3(a), and the uniform relative bound of Step 3(b) applies at every lag. In estimation it is absorbed by the persistent idiosyncratic components.

Step 5: the transitory component ($a = \rho < 1$).

(i) *Sign convention.* The Nelson-Siegel representation in equation (2) of the main text loads on the slope with a negative sign, so $S_{it} = -\mathbb{F}_{it}[c_t]$. Negating Lemma 1 applied to $\mathbb{F}_{it}[c_t]$ gives $S_{it} = -b_i^c - \mathbb{F}^C_{it} - v_{it}$. The slope fixed effect is $\alpha_{S,i} = -b_i^c$; the idiosyncratic component $-v_{it}$ has the same persistence $(1-g_{i,c})\rho_i$ and innovation variance $g^2_{i,c}\omega^2_{i,c}$, since both are invariant to the sign flip; and the common component is $-\mathbb{F}^{C}_{it} = -\tilde{\beta}_i \tilde{F}^c_t - r_{it}$, so the slope factor is the negative of the normalized benchmark factor while the loadings $\beta_{S,i} = g_{i,c}(\sigma_c^2 + \theta_c^2)^{1/2}$ keep their sign.

(ii) *The benchmark slope factor is an exact ARMA(2,1).* Write $\varphi = (1-g_c)\rho_c$ and $\kappa = \varphi/\rho$ for this step. Since $\kappa^j \rho^k = \varphi^j \rho^{k-j}$, the benchmark fundamental weights are the convolution of two geometric sequences,

$$\tilde{\lambda}_k = g_c \sigma_c \rho^k h_k(\kappa) = g_c \sigma_c \sum_{j=0}^{k} \varphi^j \rho^{k-j},$$

so the fundamental part $X_t := \sum_{k\ge 0} \tilde{\lambda}_k \varepsilon_{t-k}$ satisfies $(1-\rho L)(1-\varphi L) X_t = g_c \sigma_c \varepsilon_t$ exactly: a pure AR(2) with roots ρ and φ and no moving-average term. The moving-average term that generically appears when two AR(1) processes are added is absent here because the weights are exactly the convolution. The public-noise part $Y_t := \sum_{k\ge 0} \tilde{\delta}_k z_{t-k}$ is an AR(1), $(1-\varphi L) Y_t = g_c \theta_c z_t$. Applying $(1-\rho L)(1-\varphi L)$ to $\tilde{F}^c_t = X_t + Y_t$,

$$(1-\rho L)(1-\varphi L)\, \tilde{F}^c_t = g_c \sigma_c \varepsilon_t + g_c \theta_c \left(z_t - \rho z_{t-1}\right), \tag{A5.23}$$

an exact ARMA(2,1).

(iii) *The AR(1) coefficient* ρ_S. The statistical model specifies S_t as an AR(1). The AR(1) coefficient that matches the lag-one dynamics of the benchmark factor is the coefficient of the population projection of $\tilde{F}^c_t$ on $\tilde{F}^c_{t-1}$, which is its first-order autocorrelation ϱ_1; the sign flip in (i) leaves ϱ_1 unchanged. Because $\varepsilon \perp z$, $X_t \perp Y_t$ and $\varrho_1 = \left[\gamma_X(1) + \gamma_Y(1)\right] / \left[\gamma_X(0) + \gamma_Y(0)\right]$. Solving the Yule–Walker equations of the AR(2) in terms of its roots,

$$\gamma_X(0) = \frac{(1+\rho\varphi)\, g_c^2 \sigma_c^2}{(1-\rho\varphi)(1-\rho^2)(1-\varphi^2)}, \qquad \gamma_X(1) = \frac{\rho+\varphi}{1+\rho\varphi} \gamma_X(0),$$

while $\gamma_Y(0) = g_c^2 \theta_c^2/(1-\varphi^2)$ and $\gamma_Y(1) = \varphi \gamma_Y(0)$. Multiplying numerator and denominator by $(1-\varphi^2)/g_c^2$,

$$\varrho_1 = \frac{\dfrac{\sigma_c^2 (\rho+\varphi)}{(1-\rho\varphi)(1-\rho^2)} + \theta_c^2 \varphi}{\dfrac{\sigma_c^2 (1+\rho\varphi)}{(1-\rho\varphi)(1-\rho^2)} + \theta_c^2}. \tag{A5.24}$$

The proposition's entry

$$\rho_S = \frac{\sigma_c^2 \rho + \theta_c^2 (1-g_c)\rho_c}{\sigma_c^2 + \theta_c^2} \tag{A5.25}$$

follows from (A5.24) by two substitutions: the AR(2) part's lag-one autocorrelation $(\rho+\varphi)/(1+\rho\varphi)$ is replaced by its dominant root ρ, and the variance weights $\sigma_c^2(1+\rho\varphi)/[(1-\rho\varphi)(1-\rho^2)]$ and θ_c^2 are replaced by the innovation variances σ_c^2 and θ_c^2. Both (A5.24) and (A5.25) are closed forms, so the difference between them can be evaluated exactly at any parameter values. The two coincide as $\sigma_c \to 0$, where both equal φ; as $\theta_c \to 0$ the gap equals $\varphi(1-\rho^2)/(1+\rho\varphi)$, the amount by which the lag-one autocorrelation of an AR(2) with roots ρ and φ exceeds its dominant root. If forecasters are correct on average about the cycle's persistence, $\rho_c = \rho$, then (A5.25) reduces to $\rho_S = \rho\left[\sigma_c^2 + (1-g_c)\theta_c^2\right] / \left(\sigma_c^2 + \theta_c^2\right)$.

(iv) *The testable condition.* $\kappa < 1$ is equivalent to $\varphi < \rho$. For $\theta_c > 0$, (A5.25) is a strict convex combination of ρ and φ, so $\kappa < 1$ holds if and only if $\rho_S > \varphi$. The empirical counterpart replaces $\varphi = (1-g_c)\rho_c$ with $\bar{\rho}_S = \mathbb{E}_i[\rho_{S,i}]$, which by (A5.6) shifts the threshold by $\mathrm{Cov}_i(g_{i,c}, \rho_i)$, a product of cross-sectional deviations. The posterior estimates in Table 1 of the main paper satisfy $\rho_S > \bar{\rho}_S$, with 0.82 against 0.72. The check bears on the approximate mapping (A5.25): under the exact autocorrelation (A5.24), $\varrho_1 > \varphi$ holds for every parameter configuration, so the comparison is informative about the AR(1) coefficient interpreted through (A5.25) rather than about (A5.24) itself.

(v) *The shape parameter.* Under forecaster i's perceived law, $\mathbb{F}_{it}[c_{t+h}] = \rho_i^h \mathbb{F}_{it}[c_t]$ for every $h \geq 0$, and the perceived trend forecast is flat in h. The forecaster's point-forecast curve is therefore $h \mapsto L_{it} - \rho_i^h S_{it}$, with $S_{it} = -\mathbb{F}_{it}[c_t]$. The Nelson-Siegel curve for average inflation in equation (2) of the main text is the continuous average of the point-forecast curve $u \mapsto L_{it} - e^{-\lambda_i u} S_{it}$, since $h^{-1}\int_0^h e^{-\lambda u}\, du = (1-e^{-\lambda h})/(\lambda h)$. Matching the two point-forecast curves at every horizon requires $e^{-\lambda_i h} = \rho_i^h$ for all h, that is, $\lambda_i = -\log \rho_i$.

Step 6: the mapping. Collecting Steps 1, 4, and 5, the mapping from the noisy information model to the statistical model is

$$
\begin{aligned}
\alpha_{L,i} &= b_i^\tau & \qquad \alpha_{S,i} &= -b_i^c \\
\beta_{L,i} &= g_{i,\tau}\left(\sigma_\tau^2 + \theta_\tau^2\right)^{1/2} & \beta_{S,i} &= g_{i,c}\left(\sigma_c^2 + \theta_c^2\right)^{1/2} \\
\rho_{L,i} &= 1 - g_{i,\tau} & \rho_{S,i} &= (1-g_{i,c})\rho_i \\
\sigma_{L,i}^2 &= g_{i,\tau}^2 \omega_{i,\tau}^2 & \sigma_{S,i}^2 &= g_{i,c}^2 \omega_{i,c}^2 \\
\rho_L &= 1 & \rho_S &= \frac{\sigma_c^2\rho + \theta_c^2(1-g_c)\rho_c}{\sigma_c^2 + \theta_c^2} \\
\lambda_i &= -\log \rho_i
\end{aligned}
$$

The forecaster-specific entries are exact, by Steps 1, 4, and 5(i). The common-factor entries in the fifth row are approximate in two documented respects: $\rho_L = 1$ matches the exact order of integration of the benchmark level factor but treats its increments as white noise rather than the ARMA(1,1) in (A5.21), and ρ_S replaces the exact first-order autocorrelation (A5.24) with (A5.25). In addition, the one-factor representation itself carries the residual (A5.13), which has no date-t innovation, weights of first order in the cross-sectional deviations, and a cross-sectional average of second order, by Step 3. □

Remark 1 (Absorption in estimation)**.** *The residual r_{it} is a common process scaled by forecaster i's deviations; for the unit-root part of the trend residual this is exact, by* (A5.22)*. It enters cross-sectional covariances only through terms proportional to squared deviations, and in estimation it is absorbed as weak cross-sectional correlation among the idiosyncratic components, in the spirit of an approximate factor structure (Chamberlain and Rothschild, 1983).*

Remark 2 (Pooled idiosyncratic parameters)**.** *The baseline statistical model pools $\rho_{L,i}$, $\rho_{S,i}$, $\sigma^2_{L,i}$, and $\sigma^2_{S,i}$ across forecasters. Pooling the variances restricts only the products $g_{i,\tau}\omega_{i,\tau}$ and $g_{i,c}\omega_{i,c}$ in* (A5.10)*; since the private-noise scales $\omega_{i,\tau}$ and $\omega_{i,c}$ are free forecaster-specific parameters, this is consistent with any gain heterogeneity. Pooling the persistences evaluates the belief persistences at the average gains: the pooled values correspond to $1-g_\tau$, which equals $\mathbb{E}_i[\rho_{L,i}]$ exactly, and $(1-g_c)\rho_c$, which differs from $\mathbb{E}_i[\rho_{S,i}]$ by* $\mathrm{Cov}_i(g_{i,c},\rho_i)$*, by* (A5.6)*. The error from replacing each $\rho_{L,i}$ or $\rho_{S,i}$ by the common value is of first order in the cross-sectional dispersion of the belief persistences, the same dispersion the proposition's clustering condition restricts and the bounds of Step 3 control. In this setting $\alpha_{L,i}$ is forecaster i's belief about long-run inflation, their perception of the inflation target. The slope fixed effect $\alpha_{S,i} = -b^c_i$ is a persistent tilt of the near-term term structure whose effect decays with the forecast horizon and leaves the long-run forecast unchanged.*

A6 Additional Distributional Results

This section shows figures capturing moments of the forecast distribution that are not included in the main text.

Figure A3: DISAGREEMENT (STANDARD DEVIATION)

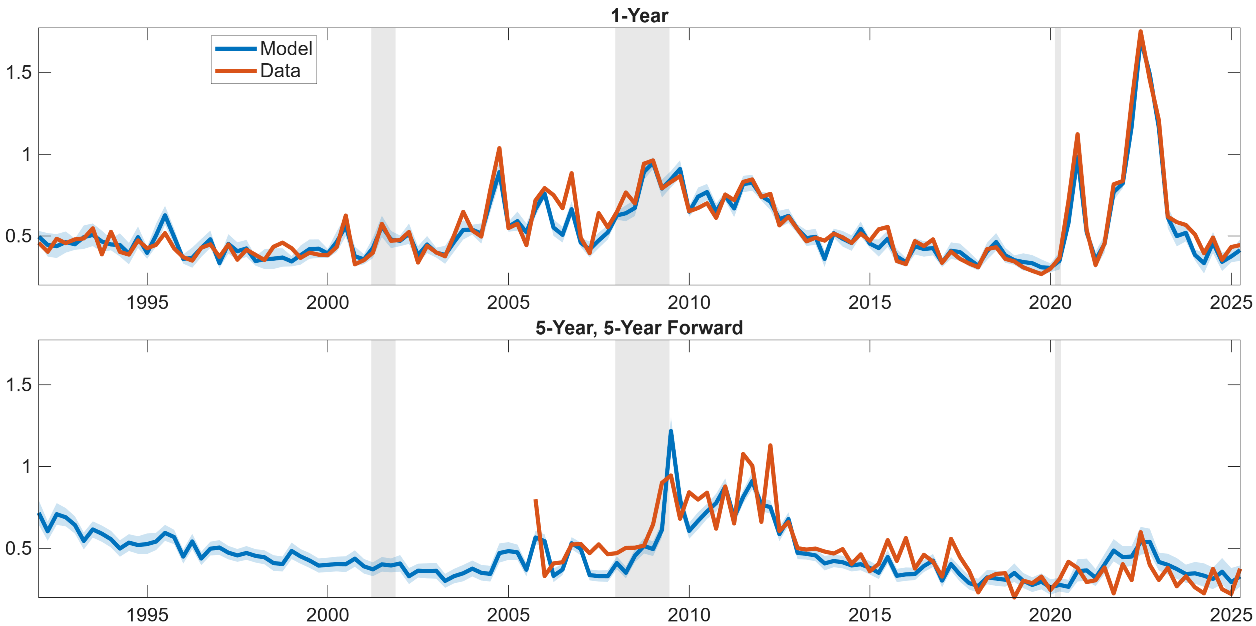

Notes: This figure compares the model-implied posterior median of cross-sectional dispersion (blue line, with the 95% credible band shaded) with the dispersion computed directly from the raw SPF responses (orange line). The top panel is the 1-year-ahead horizon (the average of the next four quarterly survey responses); the bottom panel is the 5-year, 5-year forward horizon (computed as $(40 \times 10\text{Y} - 20 \times 5\text{Y})/20$). The two series track each other closely at both horizons, so at horizons SPF observes directly, the term-structure model essentially reproduces the raw cross-sectional standard deviation. The model adds value in the decomposition (Figure 5 of the main paper) and the smooth term-structure objects (Figures 4 and 6 of the main paper), which raw forecasts cannot deliver. The shaded gray regions denote NBER recessions.

Sources: Authors' calculation

Figure A4: DISAGREEMENT (INTERQUARTILE RANGE)

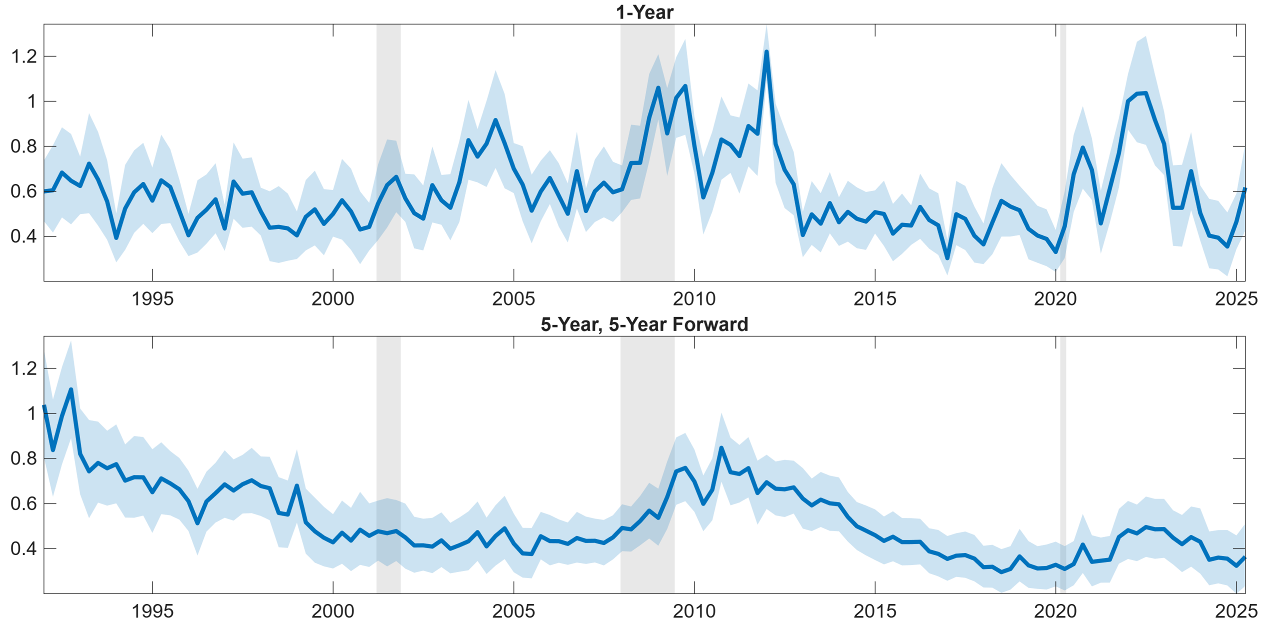

Notes: This figure plots the posterior median (blue line) and the shaded 95-percent credible band of forecast disagreement, measured by the cross-sectional interquartile range and recovered from the individual-level term structure model. The two panels correspond to the 1-year and 5-year, 5-year forward horizons. The shaded gray regions denote NBER recessions.

Sources: Authors' calculation

Figure A5: SKEWNESS

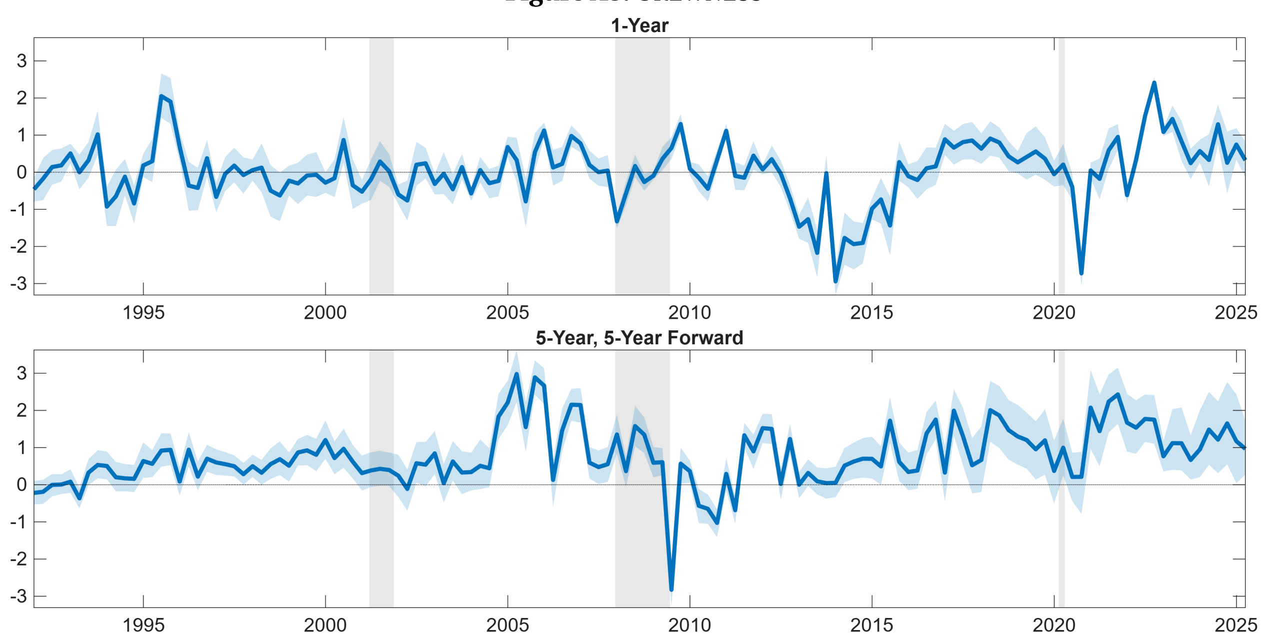

Notes: This figure plots the posterior median (blue line) and the shaded 95-percent credible band of forecast skewness recovered from the individual-level term structure model. The two panels correspond to the 1-year and 5-year, 5-year forward horizons. The shaded gray regions denote NBER recessions.

Sources: Authors' calculation

Figure A6: Kurtosis

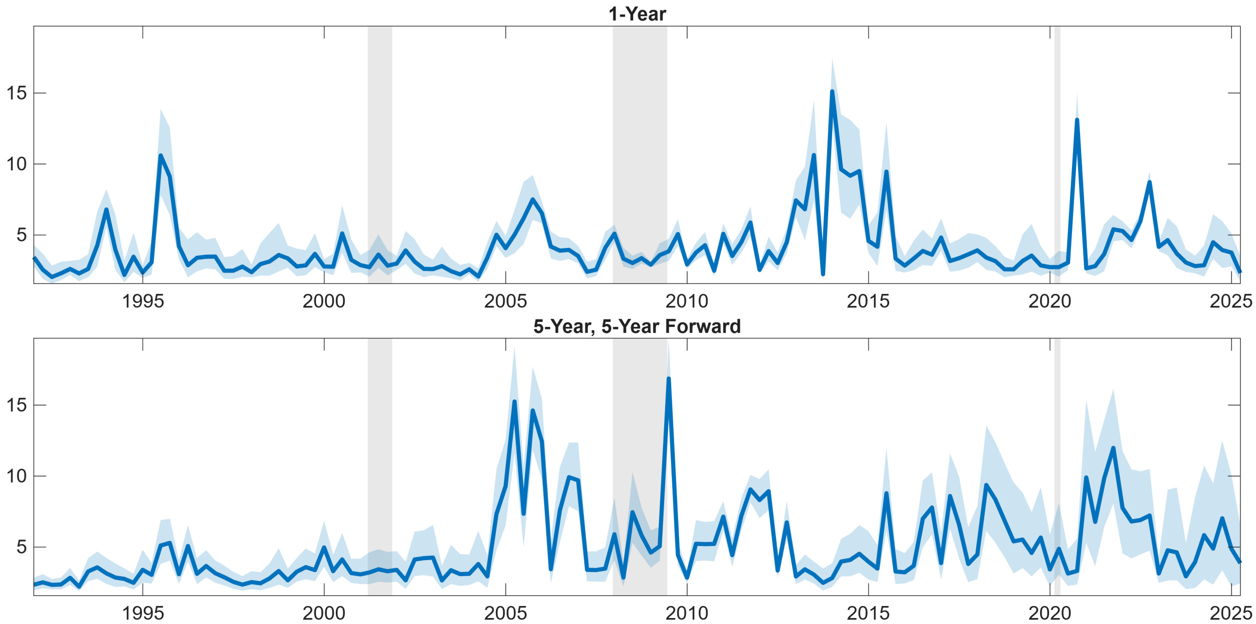

Notes: This figure plots the posterior median (blue line) and the shaded 95-percent credible band of forecast kurtosis recovered from the individual-level term structure model. The two panels correspond to the 1-year and 5-year, 5-year forward horizons. The shaded gray regions denote NBER recessions.

Sources: Authors' calculation

Figure A7: Consensus: Cross-Sectional Mean versus Median

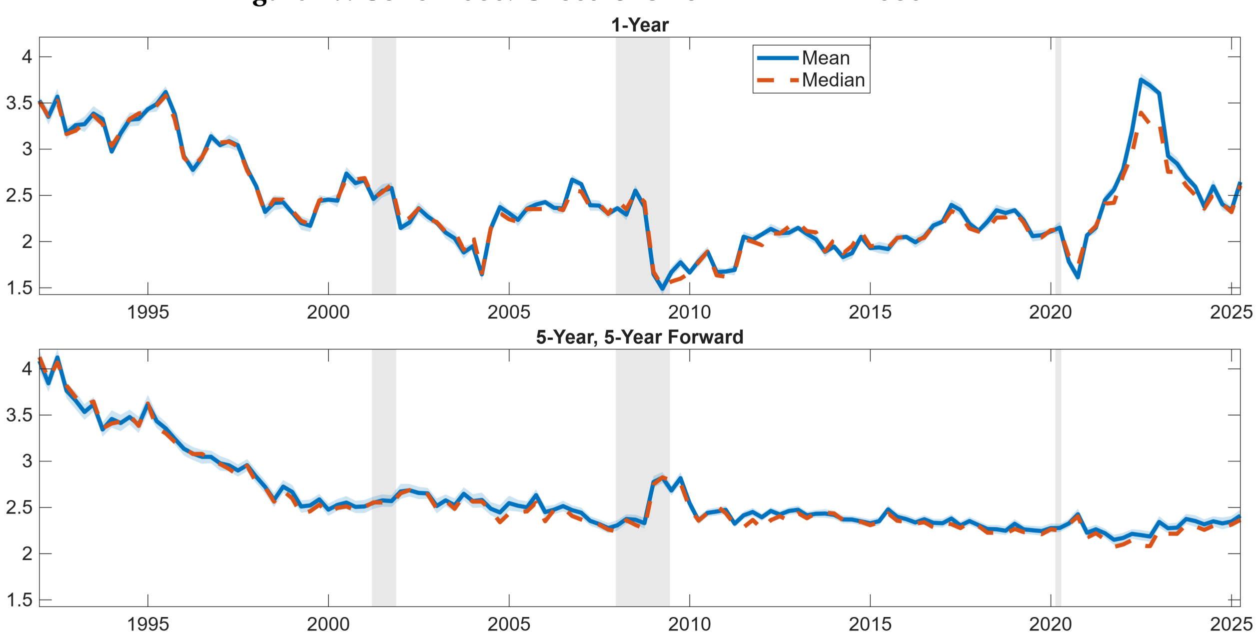

Notes: This figure plots the consensus inflation forecast computed as the cross-sectional mean across forecasters (solid blue line, with the shaded 95-percent posterior credible band) against the same object computed with the cross-sectional median (dashed orange line); both lines are posterior medians. The two panels correspond to the 1-year and 5-year, 5-year forward horizons. The two measures are visually indistinguishable over most of the sample. The visible exception is the 1-year horizon during the 2021–22 inflation surge, when the mean sits above the median. The shaded gray regions denote NBER recessions.

Sources: Authors' calculation

Figure A8: COMPARISON WITH ARUOBA (2020) ATSIX FACTORS

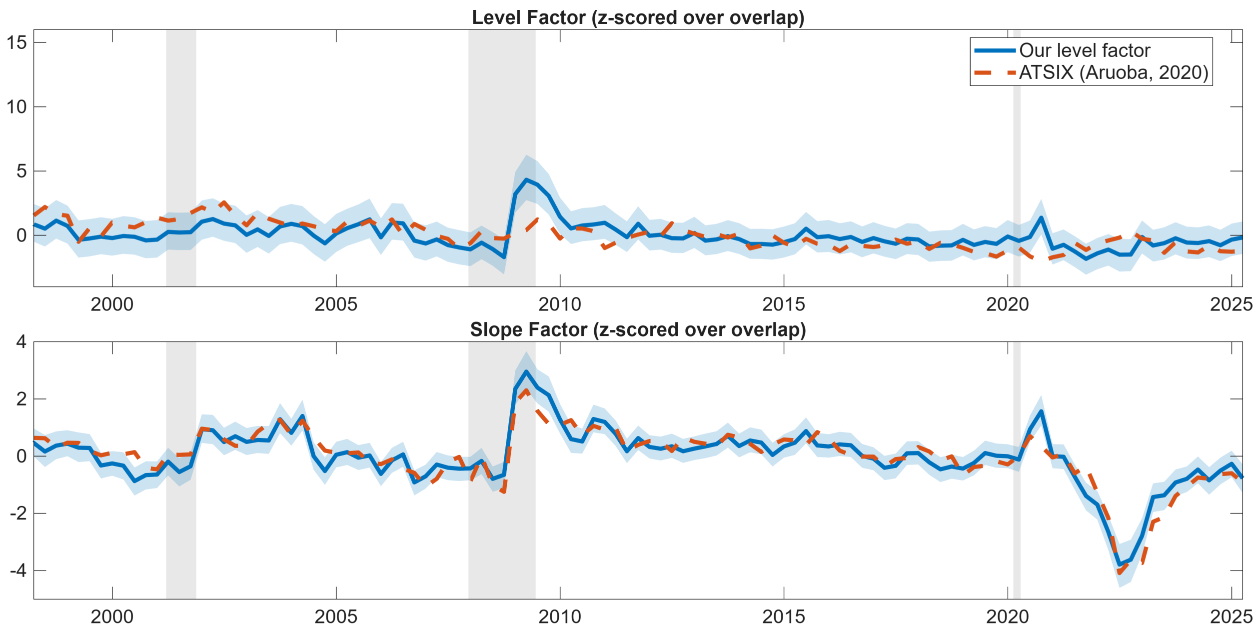

Notes: The panels compare our level (top) and slope (bottom) factors to the corresponding factors from the Aruoba Term Structure of Inflation Expectations (ATSIX; Aruoba, 2020). Our series (solid blue) is the posterior median of the common factor from the baseline specification, with the 95% posterior credible band shaded. For comparability, both series are z-scored over the overlap window (1998:Q1 onwards), which is also the plotted range. ATSIX is monthly; we align timing by taking ATSIX's value in the second month of each quarter (Feb/May/Aug/Nov), when SPF respondents typically submit their forecasts. The shaded gray regions denote NBER recessions.

Sources: Authors' calculation; Philadelphia Fed, ATSIX.

Figure A9: RESIDUAL FIT DIAGNOSTICS

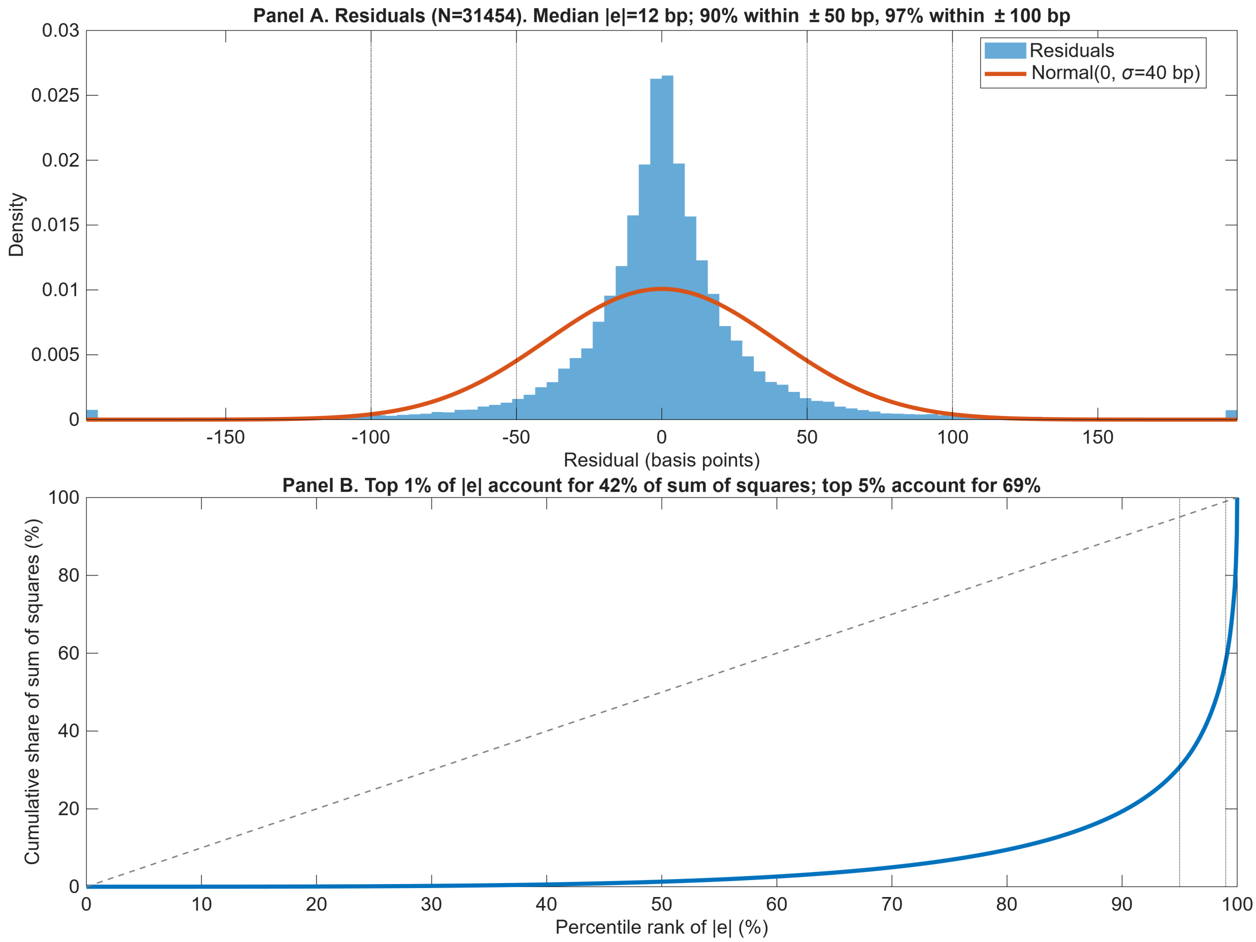

Notes: In-sample residuals $\hat{e}_{it,k} = y_{it,k} - \hat{\pi}_{it\to t+h_k|t}$ at the posterior median of the state vector and parameters. Panel A plots the pooled histogram of residuals (in basis points) with a Normal density fit to the pooled sample standard deviation overlaid in orange. Dotted vertical lines mark ±50 bp; dash-dotted lines mark ±100 bp. Residuals are clipped to $\pm 5\sigma$ for display; all observations are included in the pooled summary statistics. Panel B plots the cumulative share of the residual sum of squares $\sum_{i,t,k} \hat{e}^2_{it,k}$ as a function of the percentile rank of $|\hat{e}|$ (observations sorted ascending by $|\hat{e}|$). The grey 45-degree line is the reference for a uniform contribution of all observations.

Sources: Authors' calculation

Table A1: PER-QUESTION GOODNESS OF FIT

k	N	$\sigma_{v,k}$ (bp)	Med. $\lvert e \rvert$ (bp)	IQR (bp)	MAD (bp)	$\sigma_{v,k}$/MAD	Within 50bp (%)
1	4603	66.5	25.4	50.5	37.5	1.77	75.4
2	4603	46.6	14.8	29.5	21.8	2.14	88.5
3	4586	43.8	12.9	25.7	19.0	2.30	89.9
4	4492	43.7	15.5	30.9	23.0	1.90	88.4
5	1049	26.6	12.0	23.9	18.1	1.47	95.2
6	1061	25.6	10.3	20.3	15.2	1.69	96.9
7	1035	21.3	7.8	15.7	11.6	1.84	98.6
8	1162	9.7	2.1	4.2	3.1	3.10	100.0
9	544	30.0	11.3	21.7	16.1	1.86	91.9
10	561	32.0	13.3	26.3	19.0	1.68	92.0
11	567	29.5	13.7	27.1	20.4	1.45	91.7
12	628	32.7	15.3	31.6	23.1	1.42	90.3
13	654	18.1	6.4	12.7	9.3	1.94	98.5
14	640	19.9	8.0	15.7	11.6	1.72	97.5
15	639	23.0	8.9	18.1	13.0	1.76	96.2
16	648	25.6	11.3	22.6	16.8	1.52	94.1
17	1025	19.0	7.2	14.3	10.6	1.80	98.2
18	1000	18.8	7.6	15.0	11.1	1.69	98.7
19	953	17.4	7.9	15.8	11.6	1.50	99.4
20	1004	21.5	9.4	18.7	13.9	1.54	97.3
Pooled	31454	—	12.2	24.4	18.2	—	90.2

Notes: Per-question goodness-of-fit statistics for the in-sample residuals $\hat{e}_{it,k} = y_{it,k} - \hat{\pi}_{it\to t+h_k|t}$, evaluated at the posterior median of the state vector and parameters. $\sigma_{v,k}$ is the posterior median measurement-error standard deviation reported in Table 1 (here in basis points). MAD is the median absolute deviation scaled by 1.4826 for consistency with the Gaussian standard deviation under normality. IQR is the interquartile range of the residuals. The ratio $\sigma_{v,k}$/MAD quantifies how much the reported std is inflated by tail observations; values much above 1 indicate heavy-tailed residuals. The last column reports the fraction of observations whose absolute residual is within 50 basis points of the posterior-median fit.

Sources: Authors' calculation

Figure A10: PROPAGATION OF FED'S RESPONSE TO NEWS (NON-PUBLIC INFORMATION)

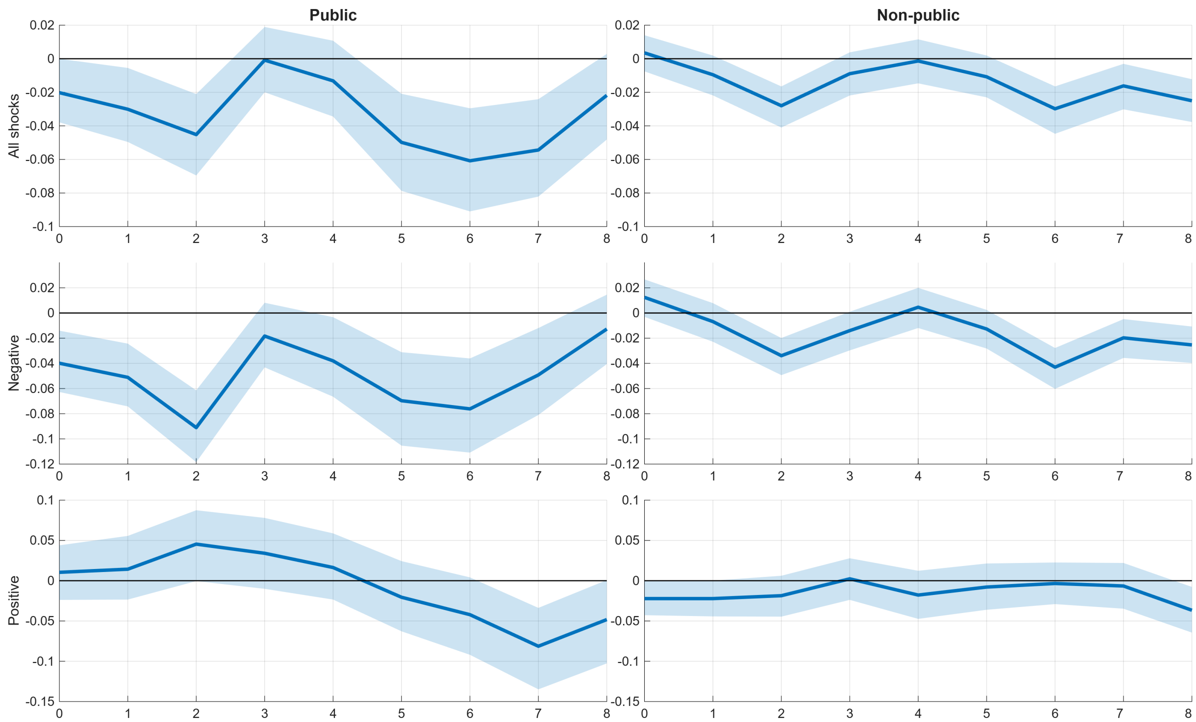

Notes: The figure shows the responses of disagreement about five-year five-year forward inflation expectations attributable to public information and to the combined non-public component (individual long-term beliefs plus private information) following the Fed's response to news from Bauer and Swanson (2023), estimated from the linear local projection model in the main text. The top row shows all shocks, the middle row restricts to negative shocks, and the bottom row restricts to positive shocks. The shaded areas indicate the 95% posterior intervals.

Source: Authors' calculation

A7 Additional Results on Monetary Policy

This section reports additional results on the transmission of monetary policy shocks. Figure A10 replaces the private component of the news local projection with the combined non-public component, individual long-term beliefs plus private information, and splits the shock by sign. The non-public measure declines at a few horizons, by roughly half as much as the public component; since the private component alone does not respond (Figure 7 of the main paper), the difference reflects the long-term-belief component and its comovement with private information. Figure A11 re-estimates the news local projection with the individual-level uncertainty measure of Binder (2017) and the SPF next-year consensus added to the controls. Figure A12 re-estimates the two-regime local projection for the orthogonalized monetary policy shock of Bauer and Swanson (2023) with the Binder (2017) uncertainty measure, the monetary policy uncertainty measure of Husted, Rogers and Sun (2020), and the SPF next-year consensus added to the controls.

Figure A11: PROPAGATION OF FED'S REACTIONS TO NEWS (WITH INDIVIDUAL UNCERTAINTY)

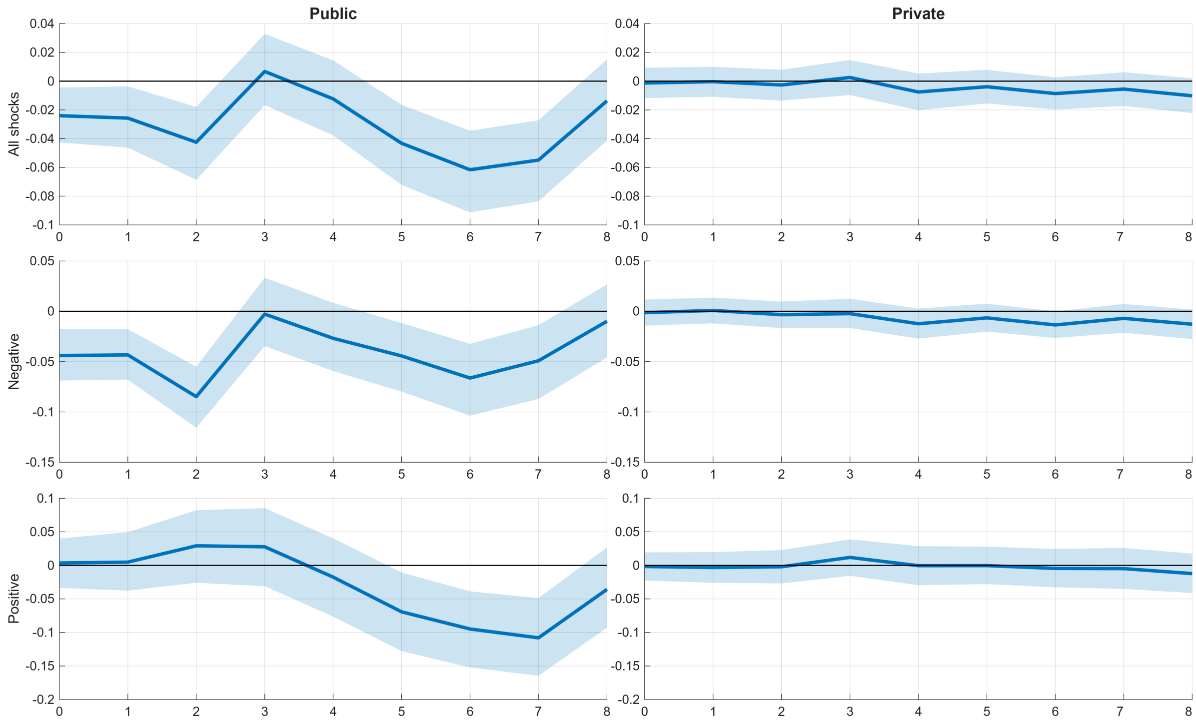

Notes: The figure shows the responses of disagreement about five-year five-year forward inflation expectations attributable to public information and private information following the Fed's response to news from Bauer and Swanson (2023), estimated from the linear local projection model in the main text. The individual-level uncertainty measure from Binder (2017) and the consensus expectations for the next year from the Survey of Professional Forecasters are additionally included as macro controls. The top row shows all shocks, the middle row restricts to negative shocks, and the bottom row restricts to positive shocks. The shaded areas indicate the 95% posterior intervals.

Source: Authors' calculation

Figure A12: Propagation of monetary policy shocks of the two regimes (with individual uncertainty and monetary policy uncertainty)

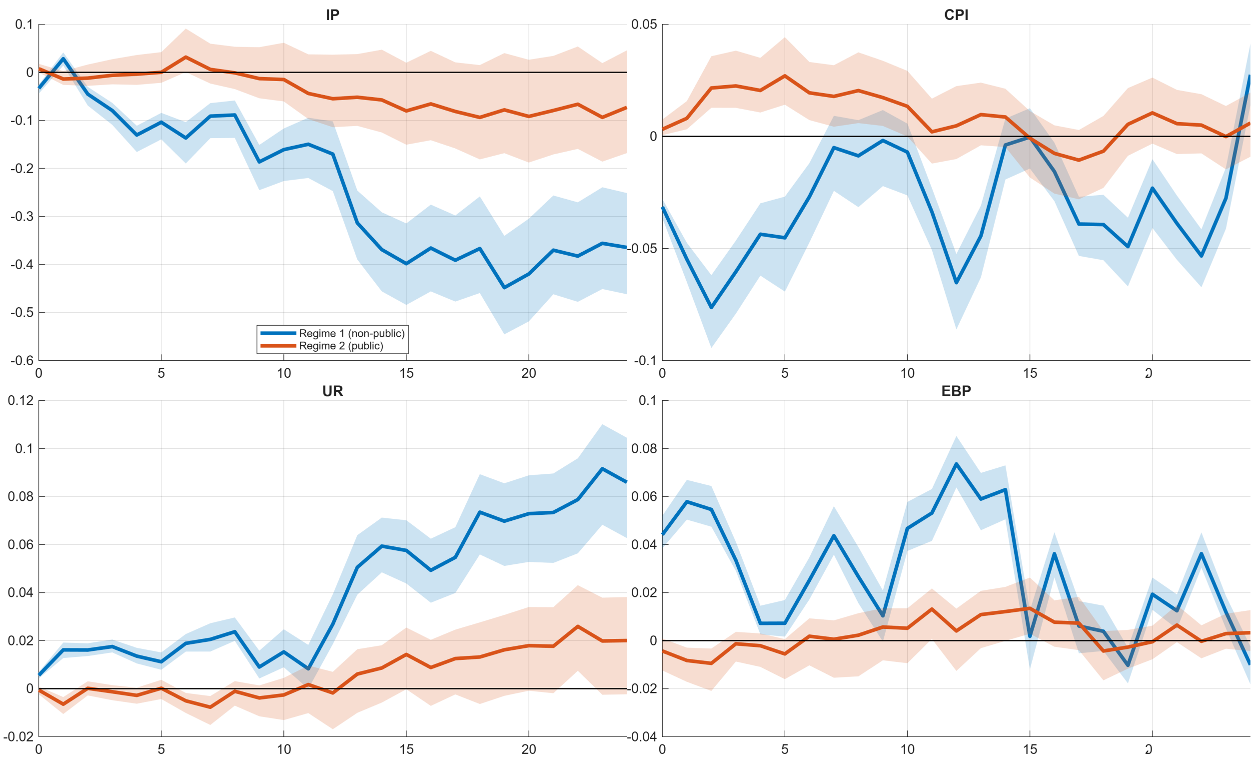

Notes: The figure reports the responses of four macroeconomic variables to the orthogonalized monetary policy shock from Bauer and Swanson (2023), with the uncertainty measure of Binder (2017), the monetary policy uncertainty measure of Husted et al. (2020), and the SPF next-year consensus added to the controls. The impulse responses of Regimes 1 and 2 are scaled by 0.9 and 0.1, respectively: the blue line shows the responses when non-public information is the source of disagreement (Regime 1) and the orange line the responses when public information is the source (Regime 2). Panels show the cumulative responses of industrial production growth (upper left) and CPI inflation (upper right), and the responses of the unemployment rate (bottom left) and the excess bond premium (bottom right). The shaded areas represent the 90% posterior intervals.

Sources: Authors' calculation

A8 Discussion of Modeling Choices

This section discusses alternative modeling choices. Section A8.1 discusses fixed factor loadings and stochastic volatility in the dynamic factor model.

A8.1 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 change 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 sensitive the estimates are to the COVID shock, we compared the model estimates through 2019:Q4 with those from the sample available at the time of the comparison, 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.

Table A2: POSTERIOR PARAMETER DISTRIBUTION STATISTICS (AR(1) SPECIFICATION)

Parameter	Median	95% CI	Parameter	Median	95% CI
$\alpha_{L,i}$	2.3483	[2.2466, 2.4574]	$\sigma_{v,6}$	0.2558	[0.2419, 0.2705]
$\alpha_{S,i}$	0.1846	[0.0288, 0.3152]	$\sigma_{v,7}$	0.2132	[0.1992, 0.2280]
$\beta_{L,i}$	0.1697	[0.1338, 0.2162]	$\sigma_{v,8}$	0.0967	[0.0808, 0.1140]
$\beta_{S,i}$	0.4080	[0.3472, 0.4794]	$\sigma_{v,9}$	0.3000	[0.2814, 0.3207]
ρ_L	0.9150	[0.8227, 0.9865]	$\sigma_{v,10}$	0.3204	[0.3010, 0.3416]
ρ_S	0.8193	[0.6738, 0.9531]	$\sigma_{v,11}$	0.2951	[0.2771, 0.3150]
$\bar{\rho}_L$	0.9430	[0.9243, 0.9606]	$\sigma_{v,12}$	0.3271	[0.3086, 0.3474]
$\bar{\rho}_S$	0.7221	[0.6866, 0.7565]	$\sigma_{v,13}$	0.1811	[0.1687, 0.1947]
σ_L	0.1185	[0.1104, 0.1267]	$\sigma_{v,14}$	0.1992	[0.1862, 0.2134]
σ_S	0.4316	[0.4138, 0.4495]	$\sigma_{v,15}$	0.2296	[0.2150, 0.2452]
λ	0.1671	[0.1600, 0.1745]	$\sigma_{v,16}$	0.2558	[0.2408, 0.2721]
$\sigma_{v,1}$	0.6652	[0.6498, 0.6810]	$\sigma_{v,17}$	0.1901	[0.1785, 0.2020]
$\sigma_{v,2}$	0.4660	[0.4553, 0.4771]	$\sigma_{v,18}$	0.1875	[0.1761, 0.1995]
$\sigma_{v,3}$	0.4378	[0.4279, 0.4480]	$\sigma_{v,19}$	0.1741	[0.1629, 0.1859]
$\sigma_{v,4}$	0.4374	[0.4278, 0.4473]	$\sigma_{v,20}$	0.2146	[0.2023, 0.2275]
$\sigma_{v,5}$	0.2656	[0.2521, 0.2798]			

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

Sources: Authors' calculation

A9 Level Factor as AR(1): Robustness

This section presents the model estimated under the alternative specification in which the common level factor follows an AR(1) process,

$$L_t = \rho_L L_{t-1} + u_{L,t},$$

$$S_t = \rho_S S_{t-1} + u_{S,t},$$

with ρ_L estimated freely. All other model components are identical to the random-walk specification used in the main text. Table A2 reports the posterior parameter distribution under the AR(1) specification; the posterior median of ρ_L is 0.92, with a 95-percent credible interval of $[0.82, 0.99]$. The figures below mirror the corresponding paper and online-appendix figures under the AR(1) specification. The substantive findings of the paper carry through: the cross-sectional heterogeneity in factor loadings, the variance decomposition of disagreement into long-term beliefs, public information, and private information, and the time-series behavior of disagreement are all quantitatively similar to the random-walk baseline.

Figure A13: Raw SPF Inflation Forecasts (AR(1) specification)

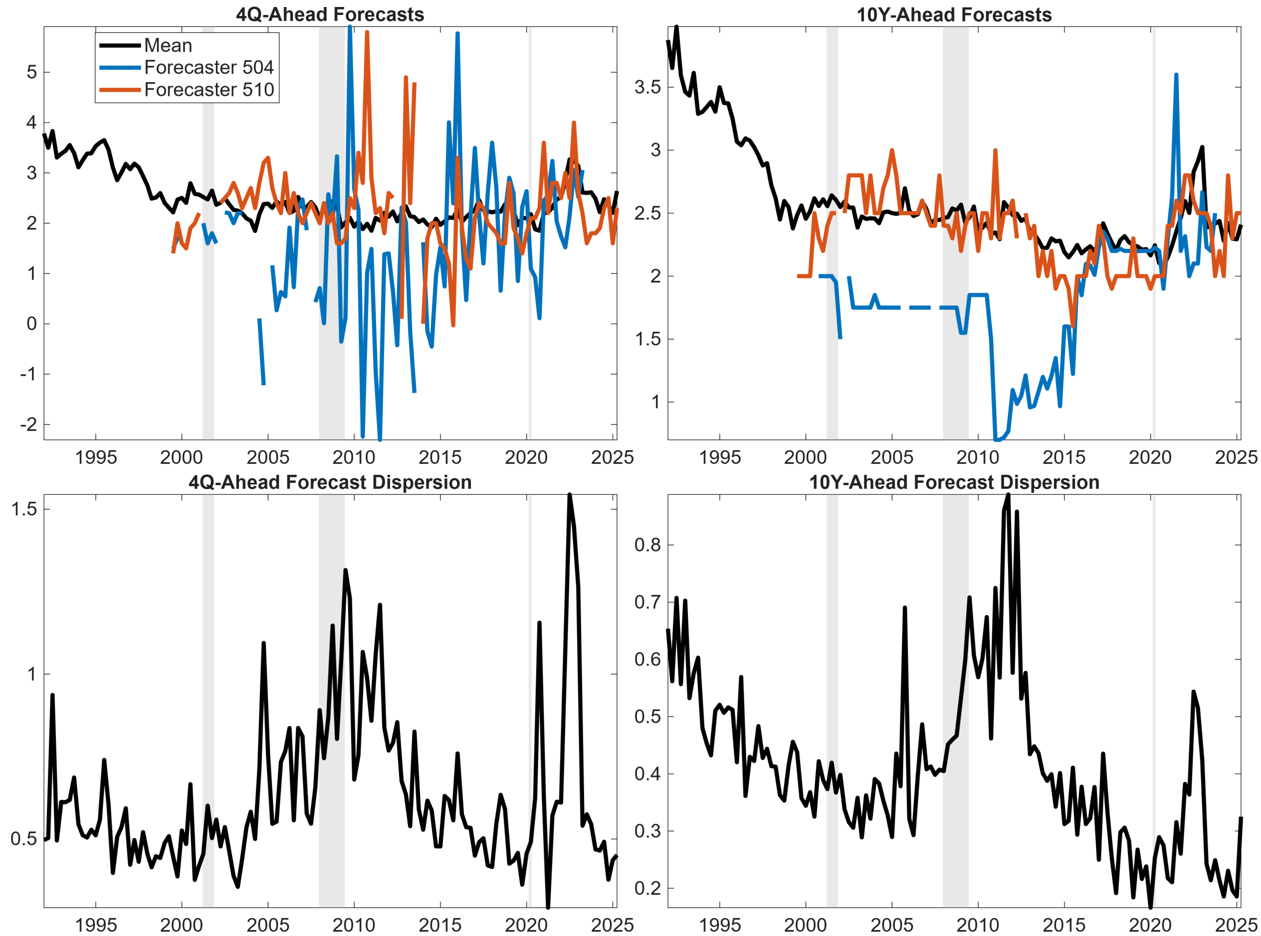

Figure A14: Term Structure of Individual Inflation Forecasts (AR(1) specification)

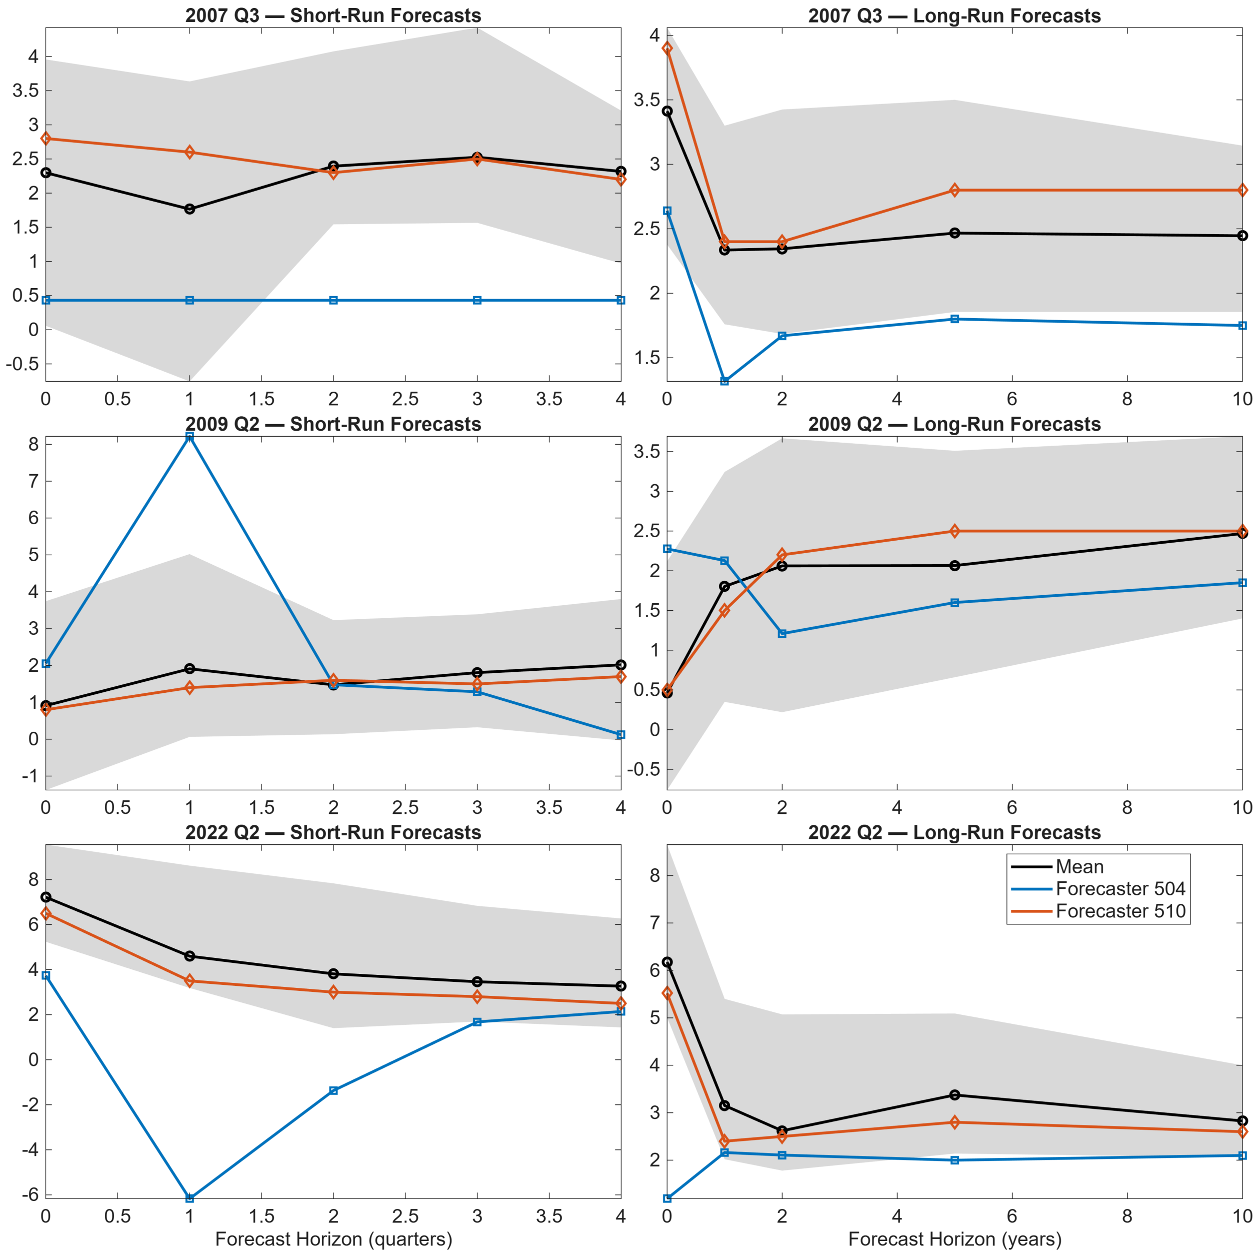

Figure A15: Smoothed Common Factors (AR(1) specification)

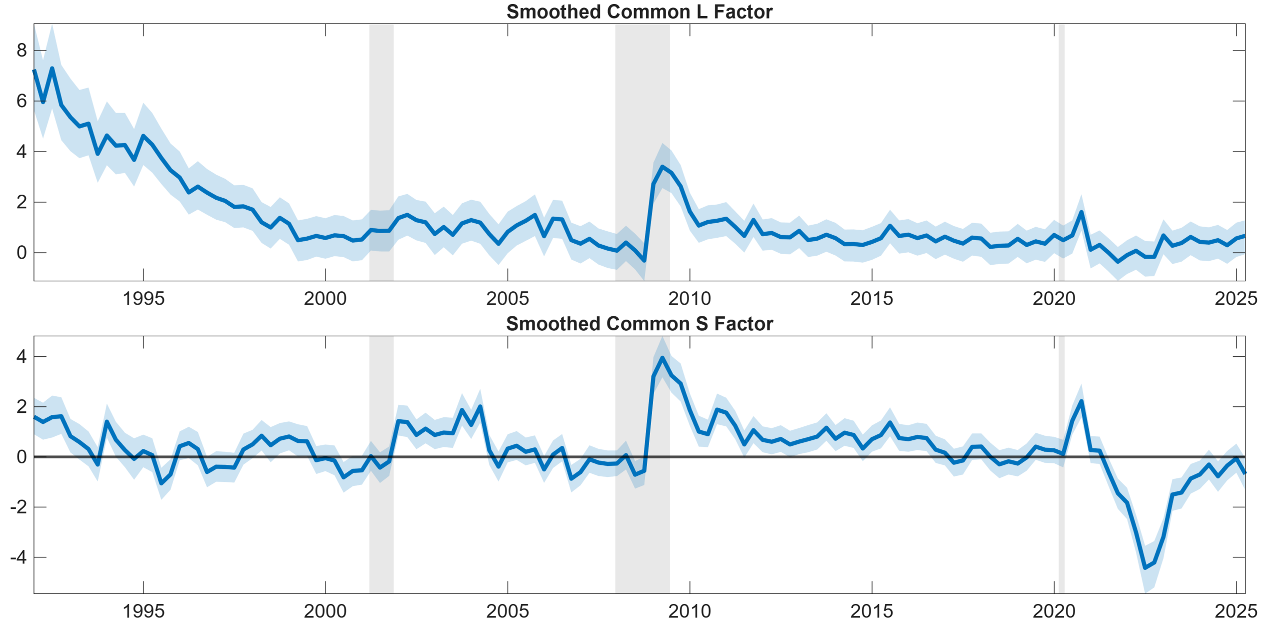

Figure A16: Consensus Inflation Forecasts (AR(1) specification)

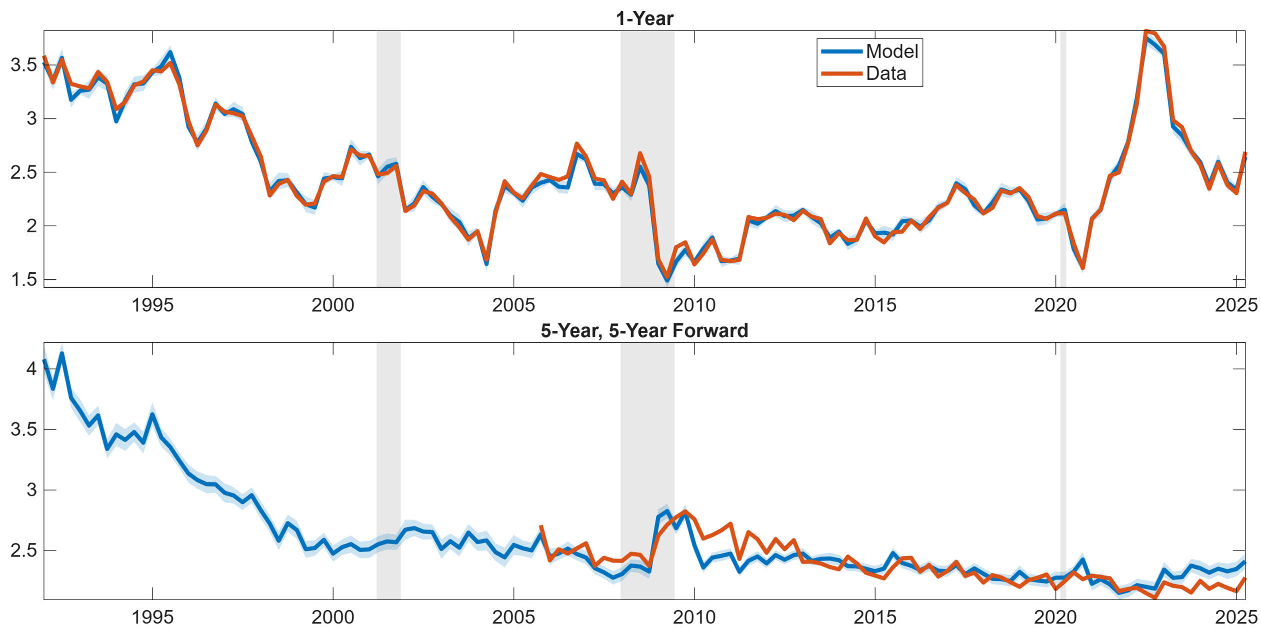

Figure A17: Posterior Distributions of Individual Factors (AR(1) specification)

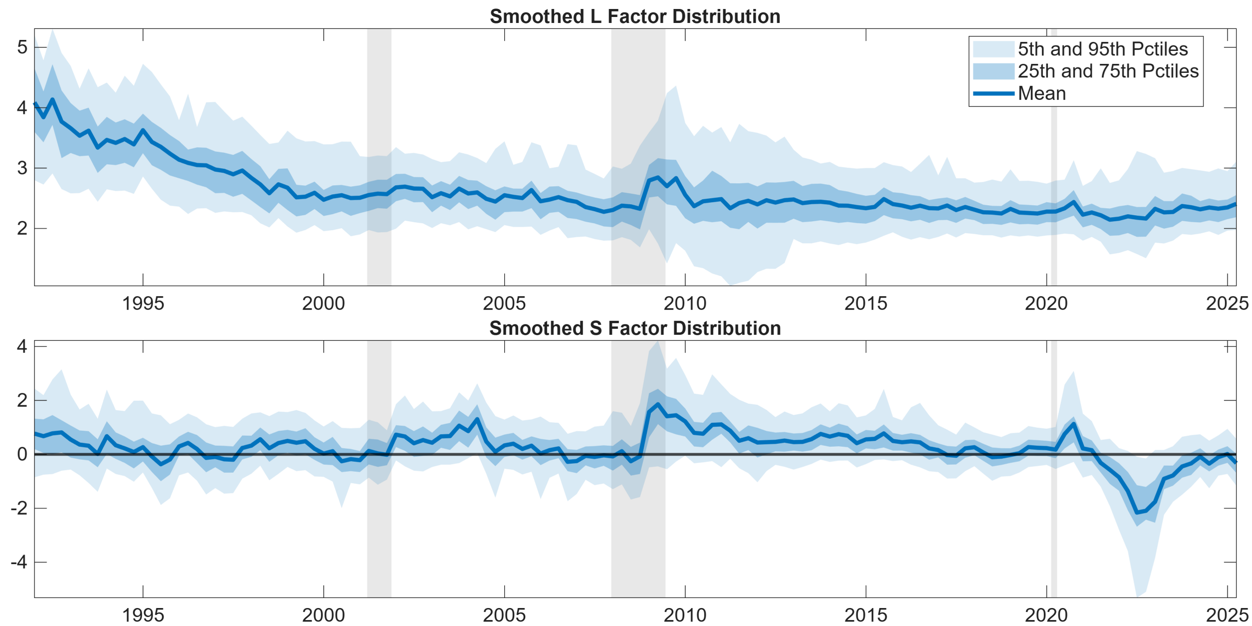

Figure A18: Posterior Distributions of Inflation Forecasts (AR(1) specification)

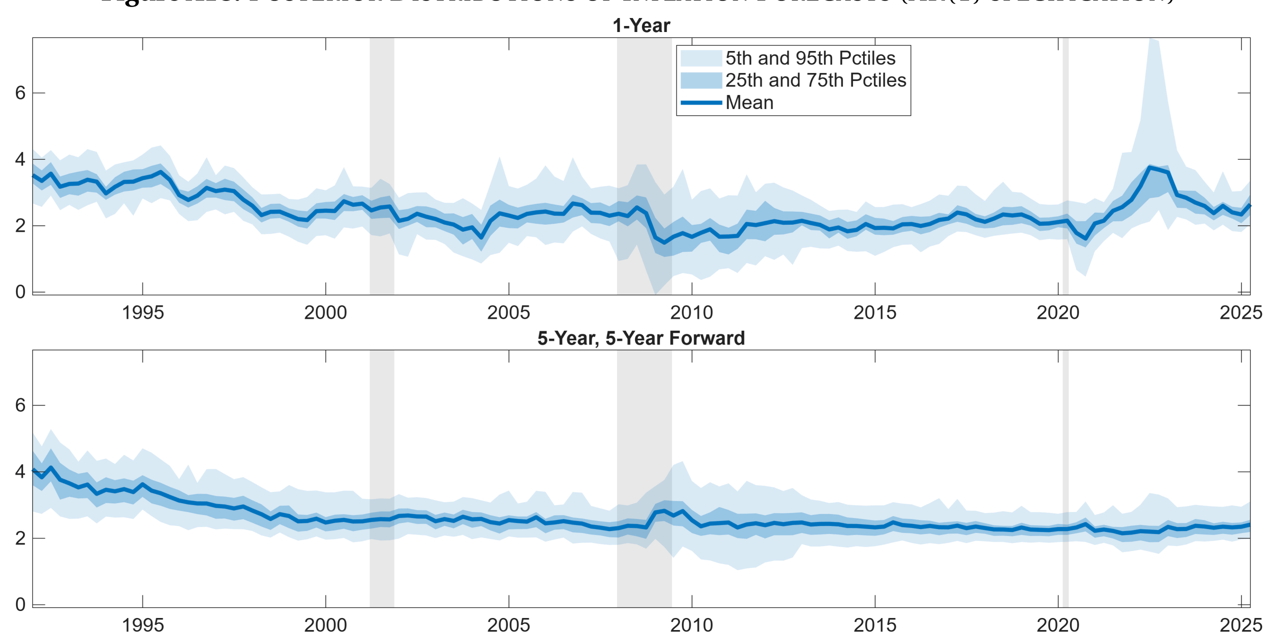

Figure A19: POSTERIOR DISTRIBUTIONS OF CONSTANTS AND FACTOR LOADINGS (AR(1) SPECIFICATION)

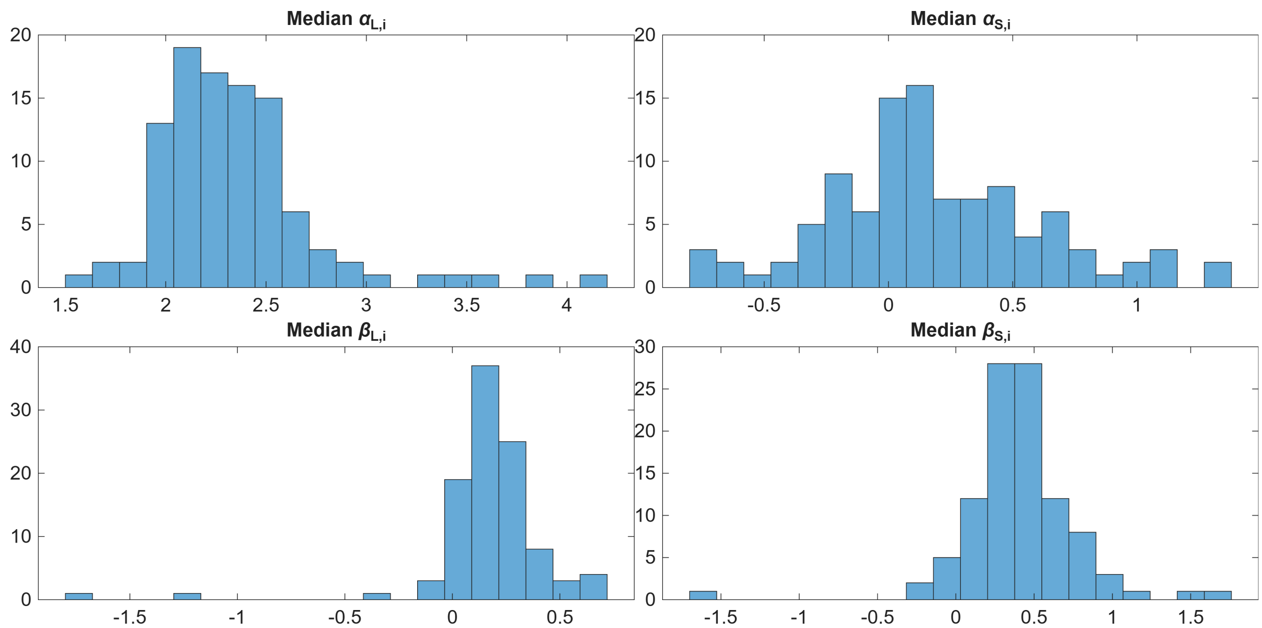

Figure A20: TERM STRUCTURE OF INFLATION FORECASTS AT THREE DATES (AR(1) SPECIFICATION)

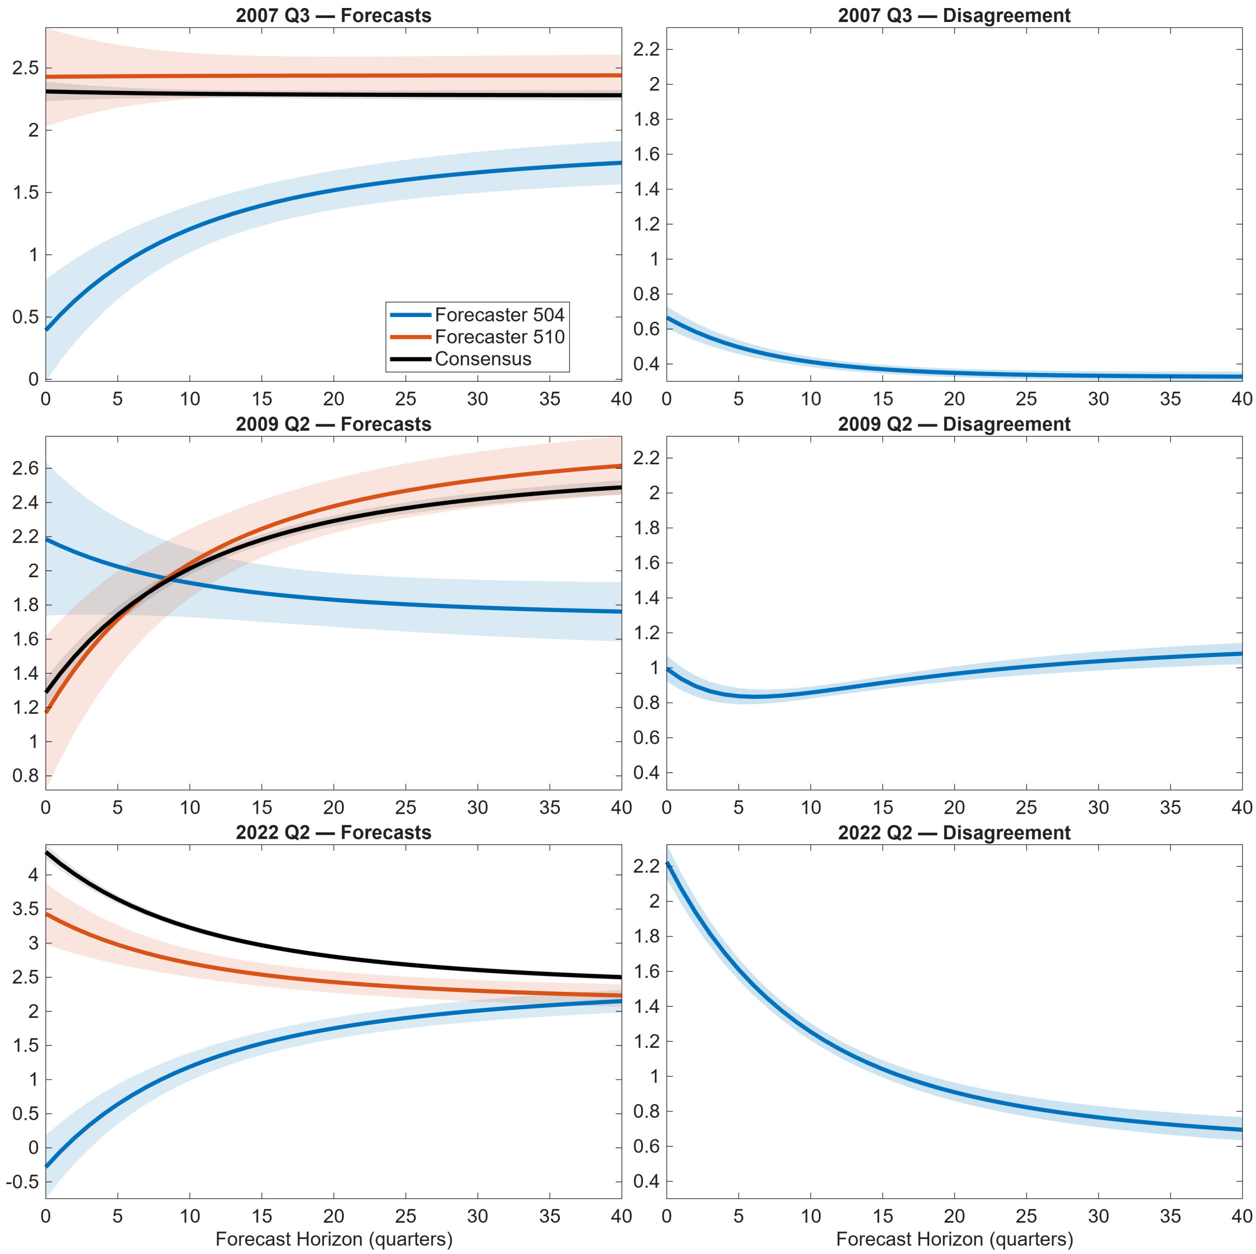

Figure A21: Variance Share of the Common Component (AR(1) specification)

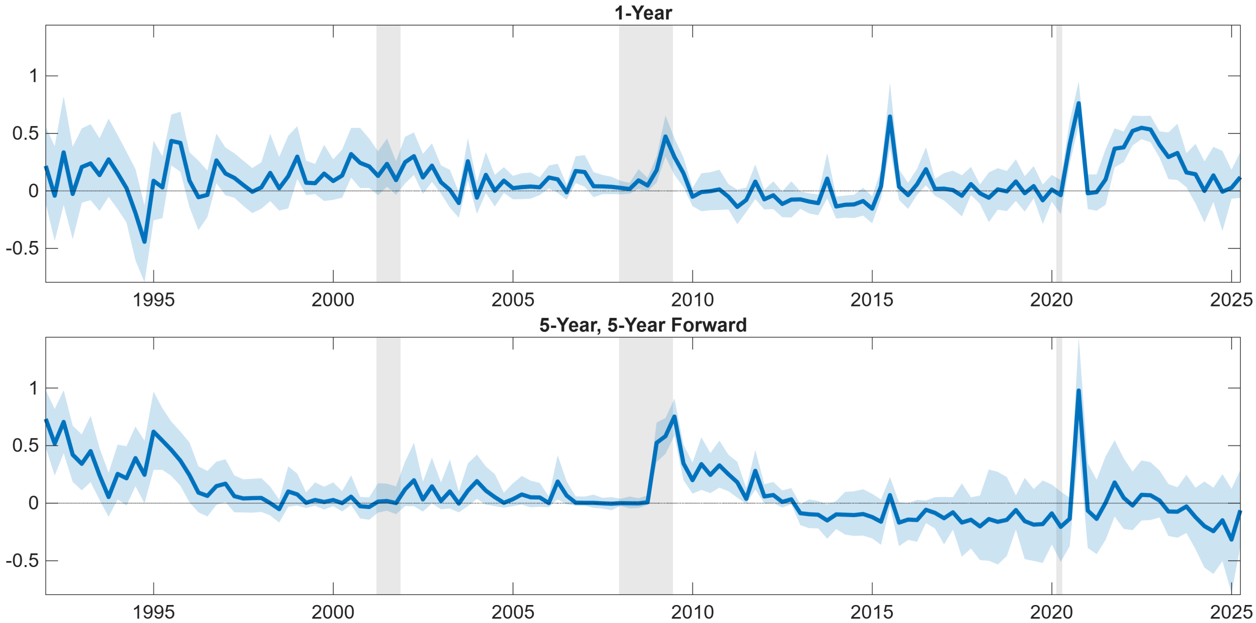

Figure A22: Term Structure of Variance Shares (AR(1) specification)

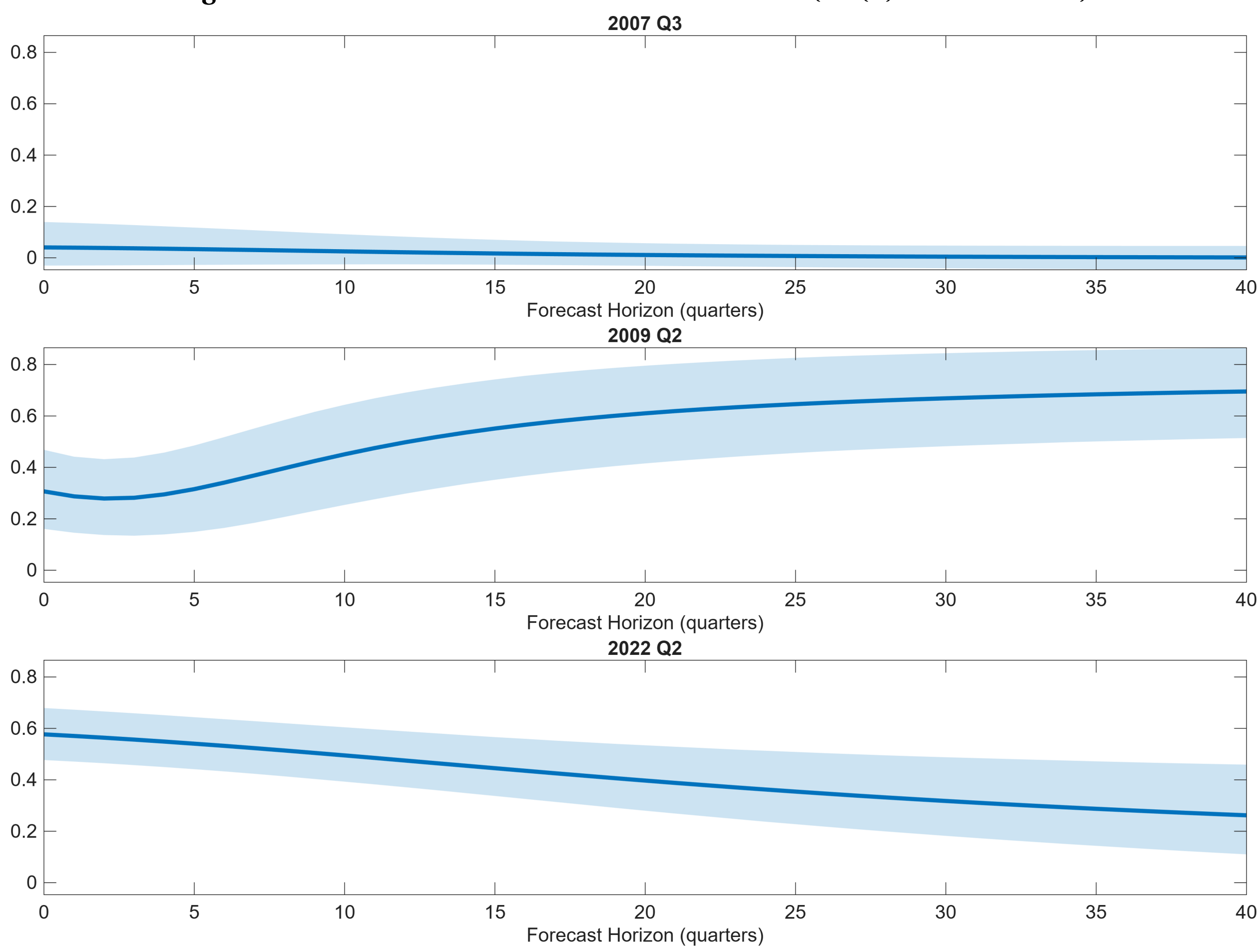

Figure A23: EFFECT OF THE FED'S RESPONSE TO NEWS ON DISAGREEMENT (AR(1) SPECIFICATION)

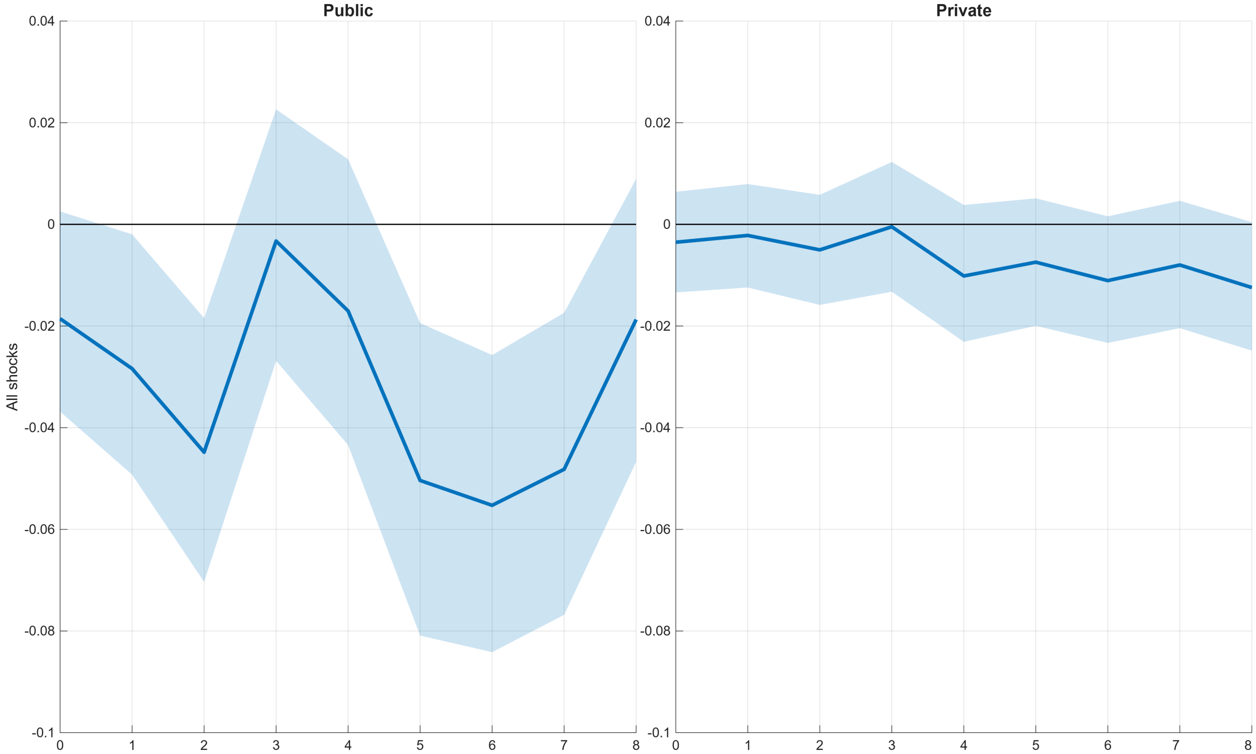

Figure A24: DISAGREEMENT AND MONETARY POLICY EFFECTIVENESS (AR(1) SPECIFICATION)

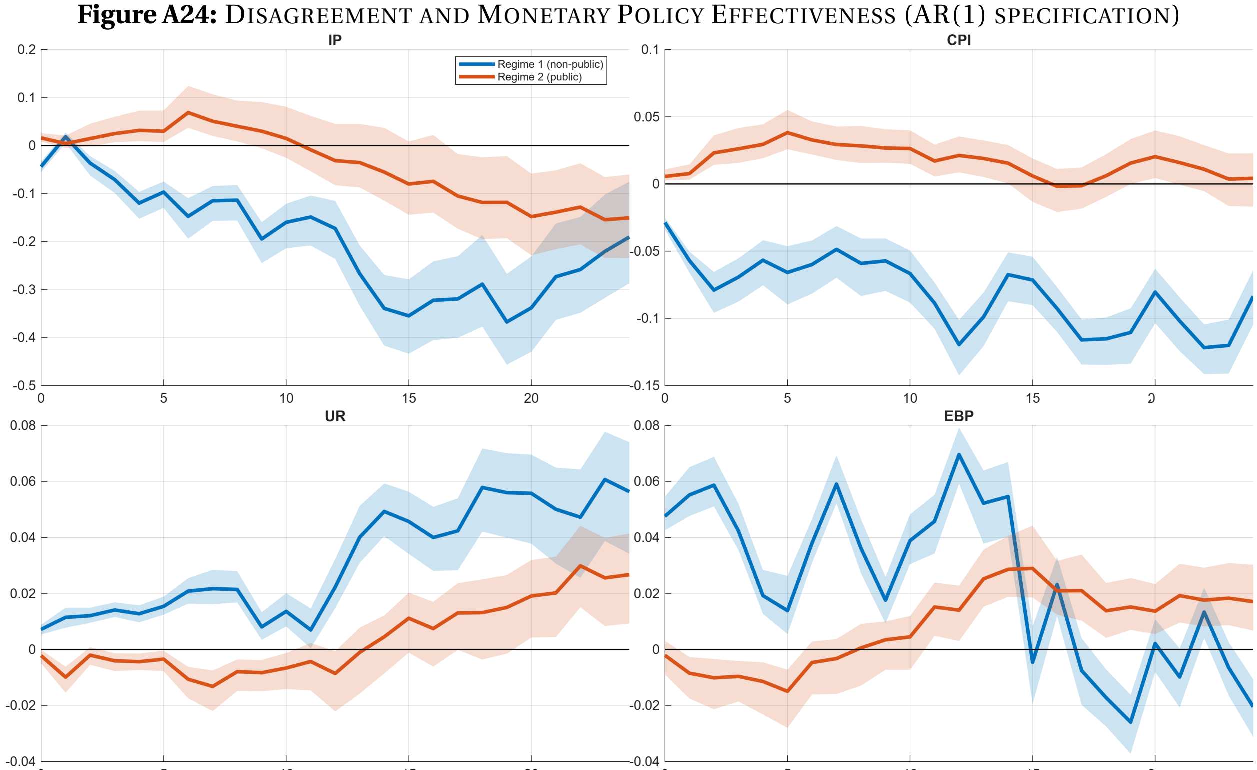

Figure A25: CROSS-SECTIONAL STANDARD DEVIATION OF INFLATION FORECASTS (AR(1) SPECIFICATION)

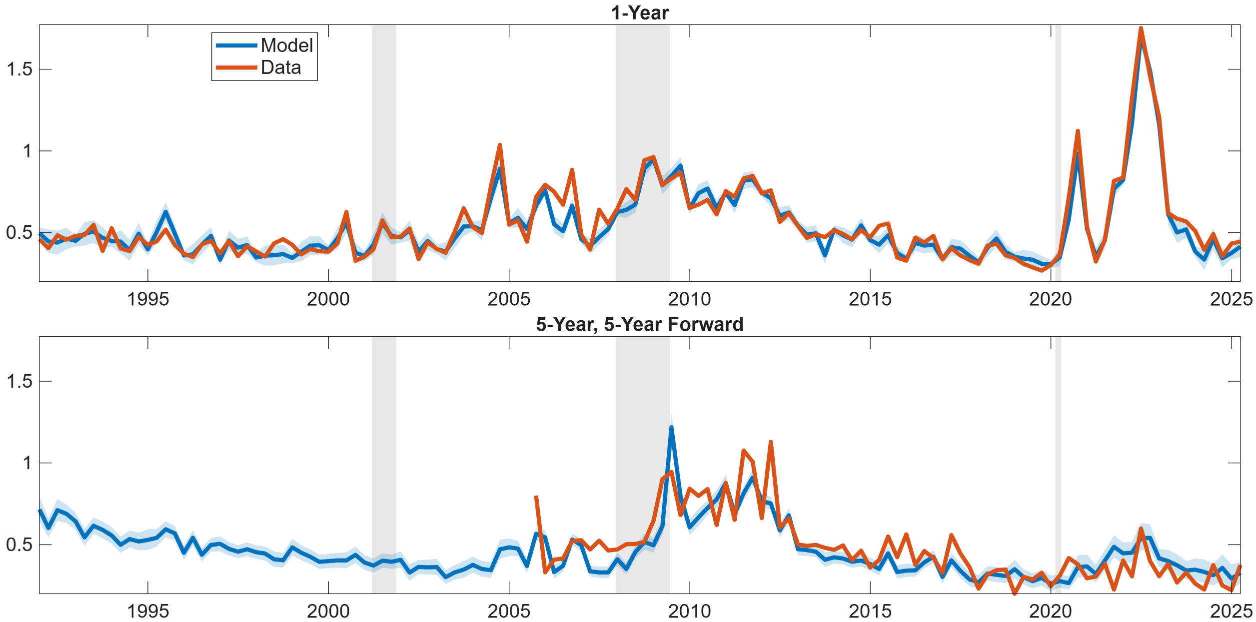

Figure A26: CROSS-SECTIONAL INTERQUARTILE RANGE OF INFLATION FORECASTS (AR(1) SPECIFICATION)

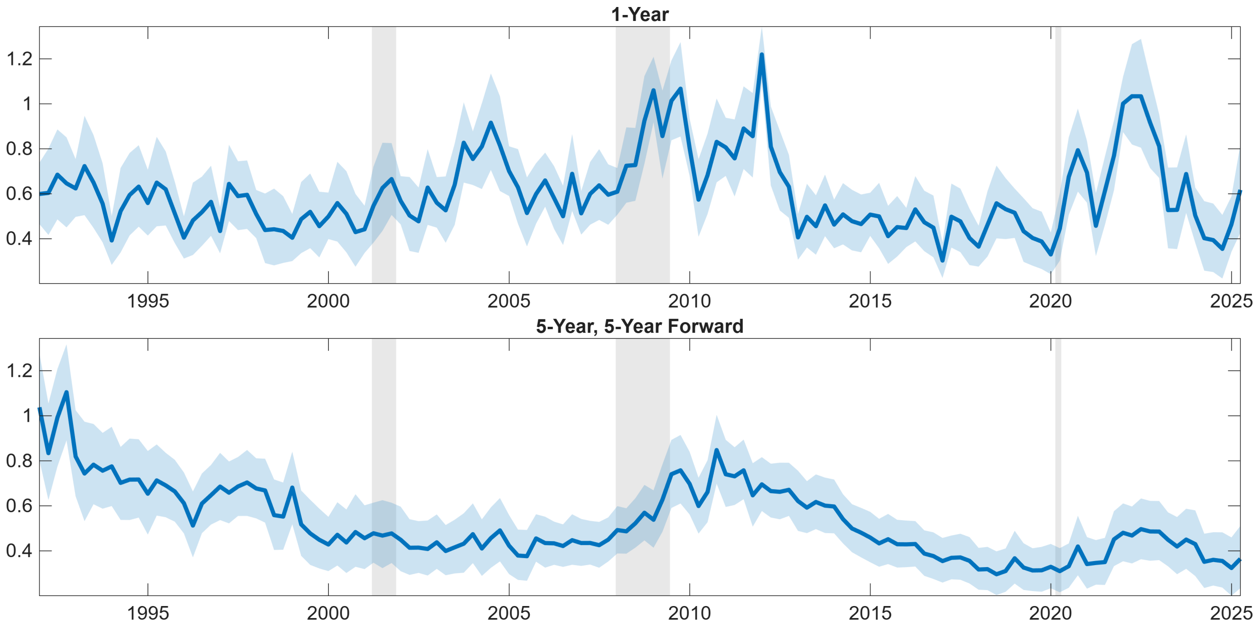

Figure A27: CROSS-SECTIONAL SKEWNESS OF INFLATION FORECASTS (AR(1) SPECIFICATION)

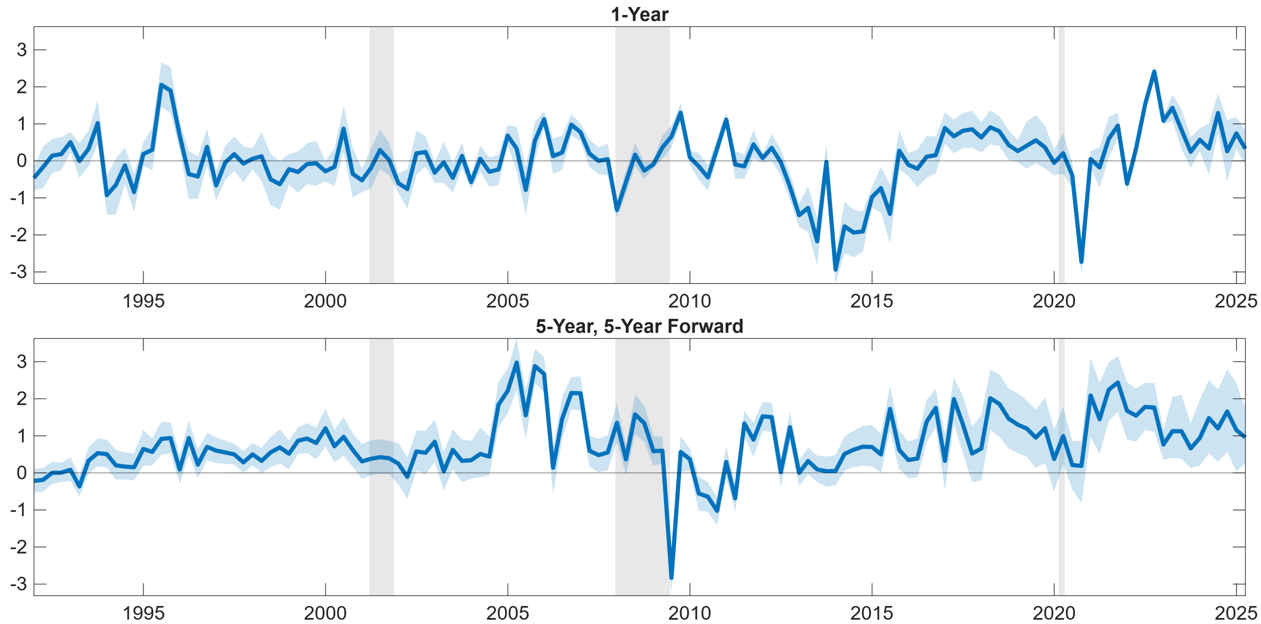

Figure A28: CROSS-SECTIONAL KURTOSIS OF INFLATION FORECASTS (AR(1) SPECIFICATION)

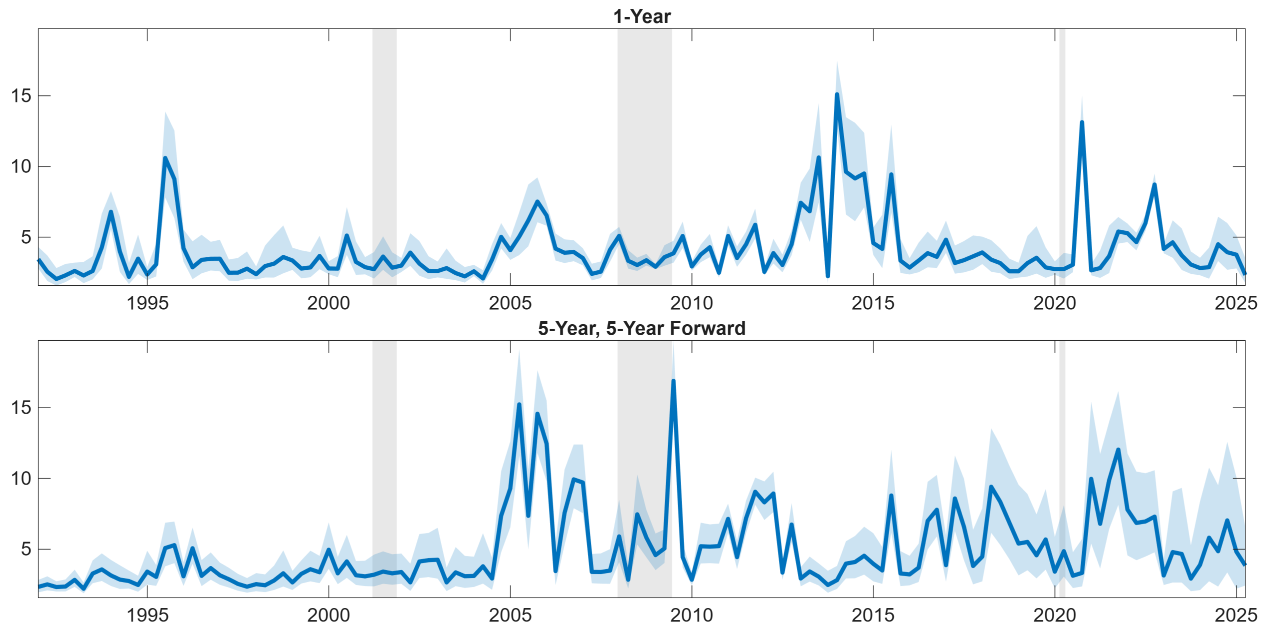

Figure A29: NEWS LOCAL PROJECTIONS, COMBINED NON-PUBLIC COMPONENT (AR(1) SPECIFICATION)

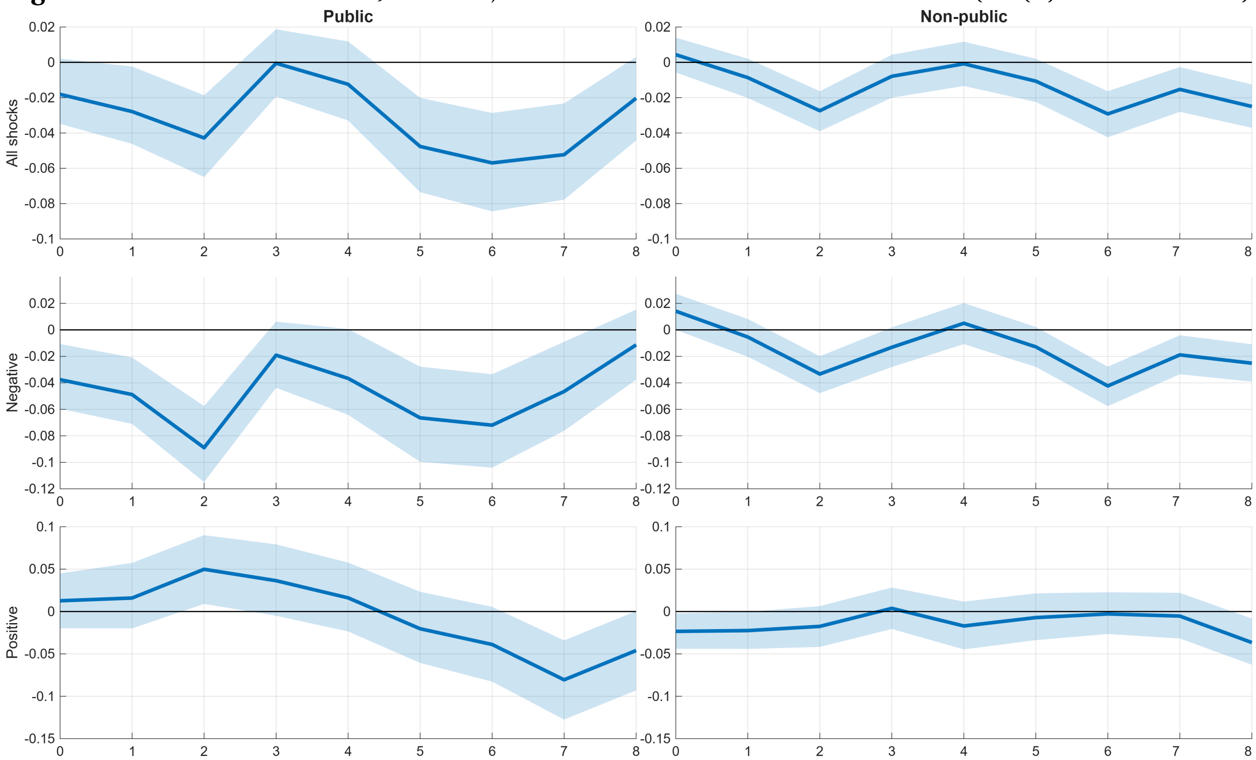

Figure A30: NEWS LOCAL PROJECTIONS WITH UNCERTAINTY CONTROLS (AR(1) SPECIFICATION)

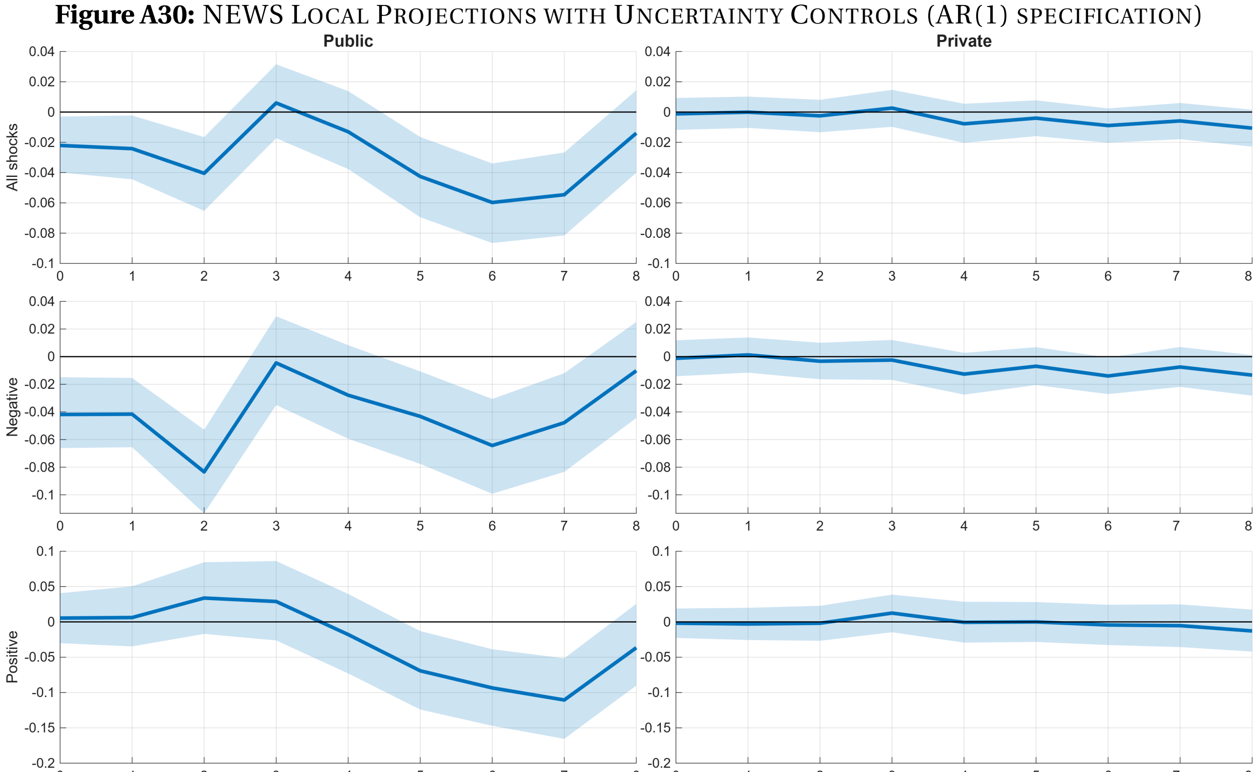

Figure A31: Nonlinear Local Projections with Uncertainty Controls (AR(1) specification)

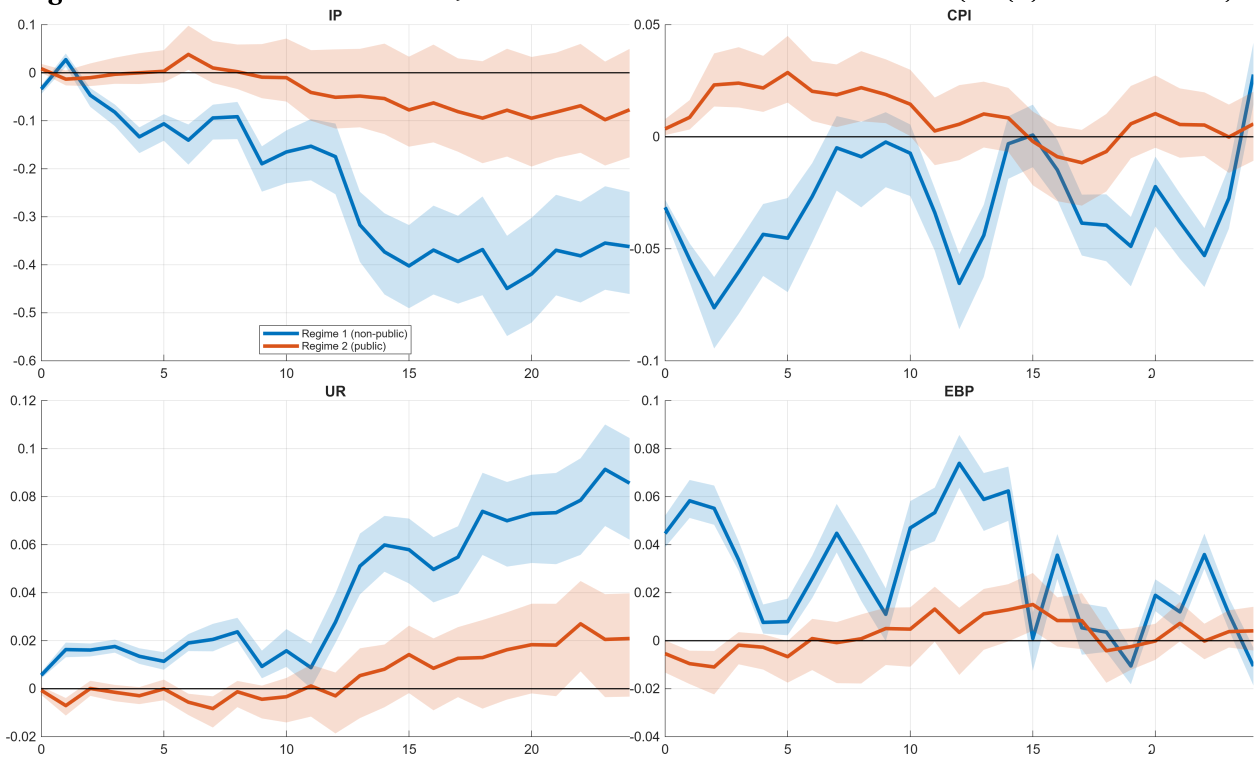

Figure A32: Comparison with Aruoba (2020) ATSIX Factors (AR(1) specification)

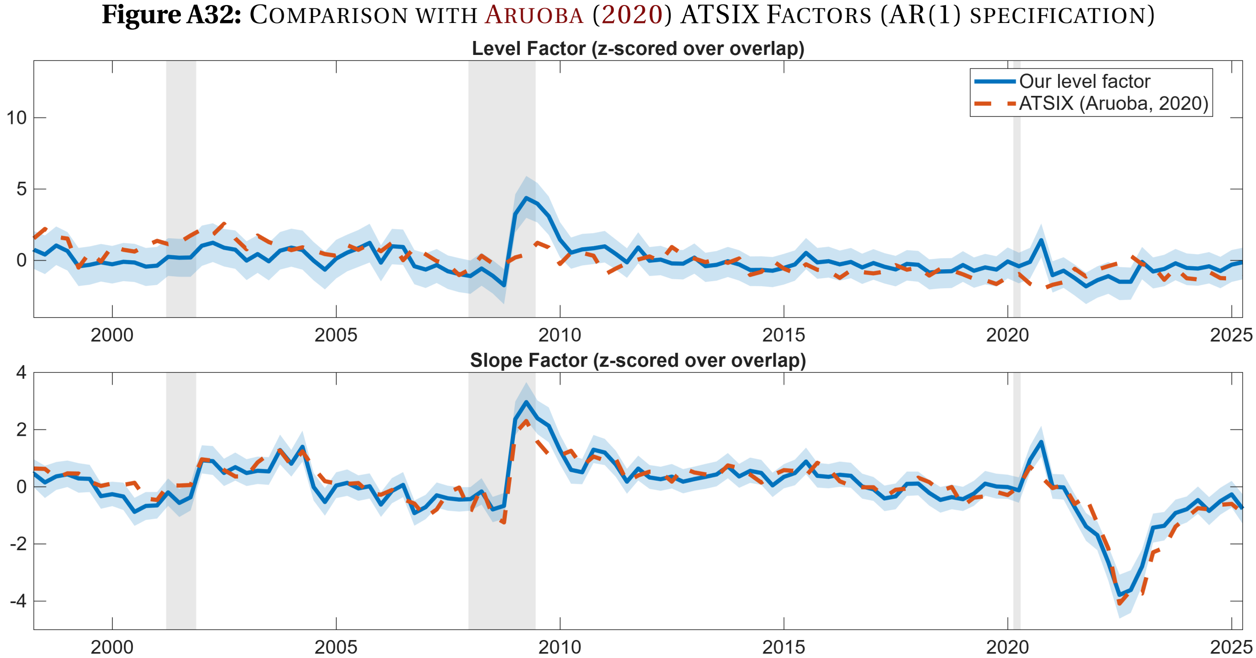

Figure A33: Residual Fit Diagnostics (AR(1) specification)

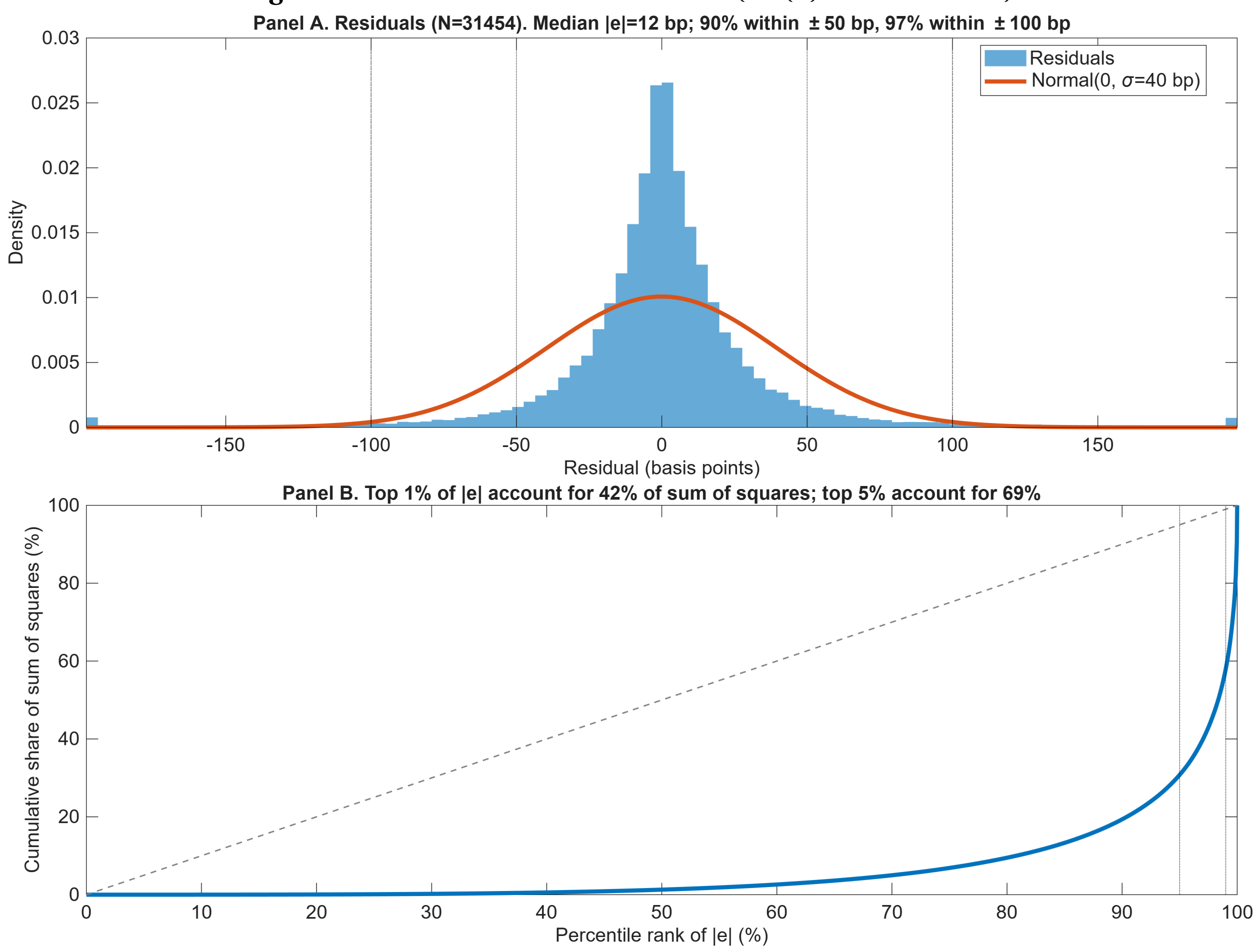

A10 General Idiosyncratic Heterogeneity

The baseline specification assigns a common autoregressive coefficient and shock variance to the idiosyncratic component of every forecaster, following Diebold, Li and Yue (2008). This appendix relaxes that assumption. We let each forecaster i have her own autoregressive coefficients $\rho_{L,i}$ and $\rho_{S,i}$ and her own shock variances $\sigma^2_{L,i}$ and $\sigma^2_{S,i}$ in the idiosyncratic level and slope processes. The Bayesian setup makes this practical despite the larger parameter count: we place hierarchical priors on the forecaster-specific parameters, drawing each autoregressive coefficient from a common normal distribution truncated to stationarity and each shock variance from a common inverse-gamma distribution, with the group-level means and scales themselves estimated under diffuse hyperpriors. The degree of pooling is therefore determined by the data rather than fixed in advance.

Table A3 reports the posterior of the cross-sectional means of the forecaster-specific parameters, together with the parameters the two specifications share, and Figure A34 plots priors against posteriors for the cross-sectional means and a selection of shared parameters. The posteriors move away from the priors and concentrate, so the data are informative about the idiosyncratic dynamics, and the estimated dispersion across forecasters is moderate.

Figure A35 reports the variance decomposition of disagreement under this specification, the counterpart of Figure 5 in the main paper. In normal times and around the Great Recession the decomposition is essentially the same as in the baseline: the rise in the share attributable to disagreement about public information in 2008–2009 is preserved, and at longer horizons it is if anything slightly larger.

The one visible difference is the early COVID period. There the share attributable to heterogeneous responses to public information is more muted than in the baseline, especially at the 5-year, 5-year-forward horizon, and part of what the baseline assigns to the common component is reallocated to the idiosyncratic component. This is what one would expect if the homogeneous specification were absorbing forecaster-specific variation during an unusual period.

The monetary-policy results in Section 6 of the main paper are unchanged under the heterogeneous specification. To conserve space we report the parameter table, the prior-posterior diagnostics, and the disagreement decomposition here; the full set of figures and tables is in the replication package.

Figure A34: Prior and Posterior Distributions of the Forecaster-Specific Idiosyncratic Parameters

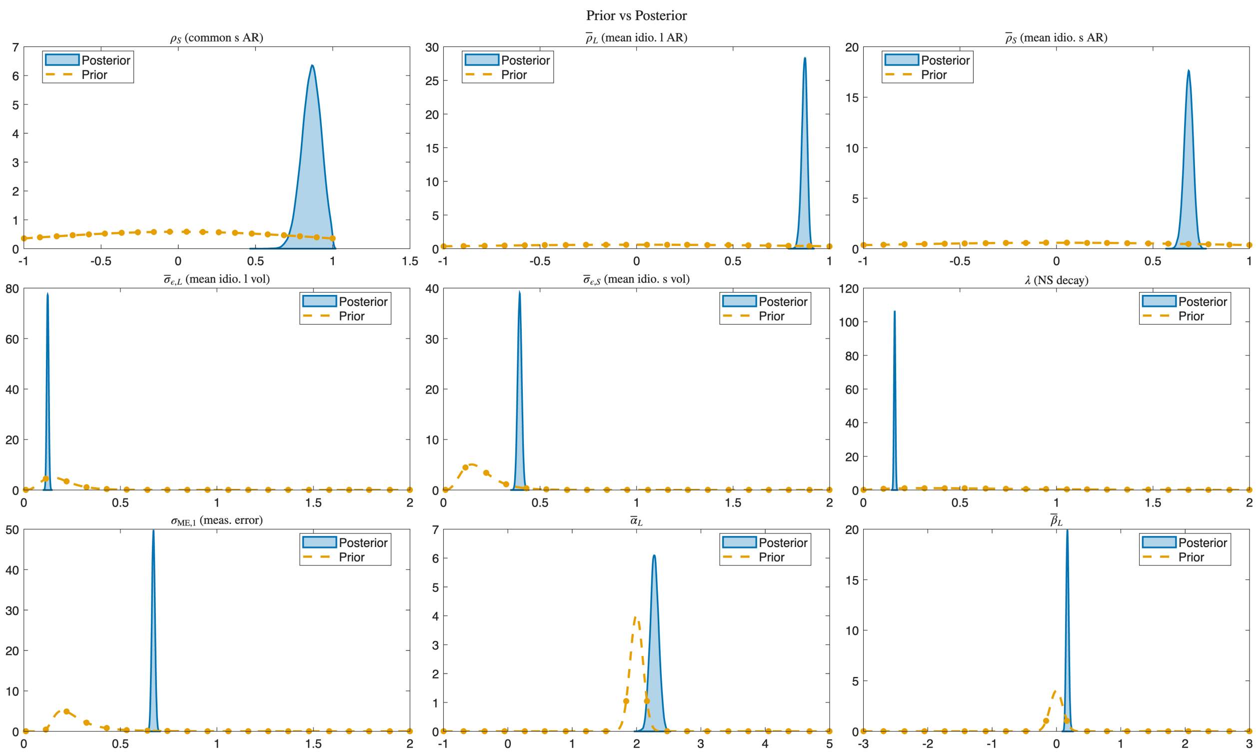

Notes: Prior (dashed) and posterior (shaded) distributions for selected parameters of the heterogeneous-idiosyncrasy specification: the common slope autoregressive coefficient ρ_S, the cross-sectional means of the forecaster-specific idiosyncratic autoregressive coefficients ($\bar{\rho}_L$, $\bar{\rho}_S$) and shock standard deviations ($\bar{\sigma}_{\varepsilon,L}$, $\bar{\sigma}_{\varepsilon,S}$), the Nelson-Siegel decay parameter λ, one representative measurement-error standard deviation ($\sigma_{\text{ME},1}$), and the cross-sectional means of the level fixed effects and loadings ($\bar{\alpha}_L$, $\bar{\beta}_L$). The forecaster-specific parameters carry hierarchical priors with estimated group-level means and scales; the dashed curves show the priors on the group-level parameters, which for the autoregressive coefficients are centered at zero and truncated to stationarity.

Sources: Authors' calculation

Table A3: POSTERIOR PARAMETER DISTRIBUTION STATISTICS (HETEROGENEOUS IDIOSYNCRATIC SPECIFICATION)

Parameter	Median	95% CI	Parameter	Median	95% CI
$\alpha_{L,i}$	2.2775	[2.1524, 2.4053]	$\sigma_{v,6}$	0.2555	[0.2417, 0.2702]
$\alpha_{S,i}$	0.1374	[-0.0109, 0.2827]	$\sigma_{v,7}$	0.2072	[0.1933, 0.2215]
$\beta_{L,i}$	0.1704	[0.1336, 0.2121]	$\sigma_{v,8}$	0.1140	[0.0971, 0.1308]
$\beta_{S,i}$	0.4163	[0.3534, 0.4862]	$\sigma_{v,9}$	0.3020	[0.2831, 0.3227]
ρ_S	0.8641	[0.7352, 0.9746]	$\sigma_{v,10}$	0.3190	[0.2995, 0.3405]
$\overline{\rho}_L$	0.8700	[0.8392, 0.8948]	$\sigma_{v,11}$	0.2930	[0.2750, 0.3126]
$\overline{\rho}_S$	0.6846	[0.6385, 0.7269]	$\sigma_{v,12}$	0.3268	[0.3079, 0.3473]
$\overline{\sigma}_L$	0.1236	[0.1141, 0.1337]	$\sigma_{v,13}$	0.1804	[0.1680, 0.1940]
$\overline{\sigma}_S$	0.3946	[0.3749, 0.4152]	$\sigma_{v,14}$	0.1931	[0.1803, 0.2071]
λ	0.1620	[0.1550, 0.1693]	$\sigma_{v,15}$	0.2276	[0.2133, 0.2430]
$\sigma_{v,1}$	0.6711	[0.6554, 0.6872]	$\sigma_{v,16}$	0.2496	[0.2344, 0.2658]
$\sigma_{v,2}$	0.4653	[0.4545, 0.4765]	$\sigma_{v,17}$	0.1745	[0.1629, 0.1867]
$\sigma_{v,3}$	0.4340	[0.4242, 0.4441]	$\sigma_{v,18}$	0.1720	[0.1604, 0.1843]
$\sigma_{v,4}$	0.4367	[0.4271, 0.4467]	$\sigma_{v,19}$	0.1565	[0.1448, 0.1687]
$\sigma_{v,5}$	0.2670	[0.2536, 0.2813]	$\sigma_{v,20}$	0.2009	[0.1878, 0.2145]

Notes: The “Median" column reports the posterior median of the corresponding parameter and the “95% CI" column reports the parameter’s 95-percent credible interval. For the parameters $\alpha_{L,i}$, $\beta_{L,i}$, $\alpha_{S,i}$, $\beta_{S,i}$, the table reports statistics for the average value across forecasters. Parameters with an overline ($\overline{\rho}$, $\overline{\sigma}$) report the cross-sectional mean of forecaster-specific estimates.

Sources: Authors’ calculation

Figure A35: DISAGREEMENT VARIANCE DECOMPOSITION (HETEROGENEOUS IDIOSYNCRATIC SPECIFICATION)

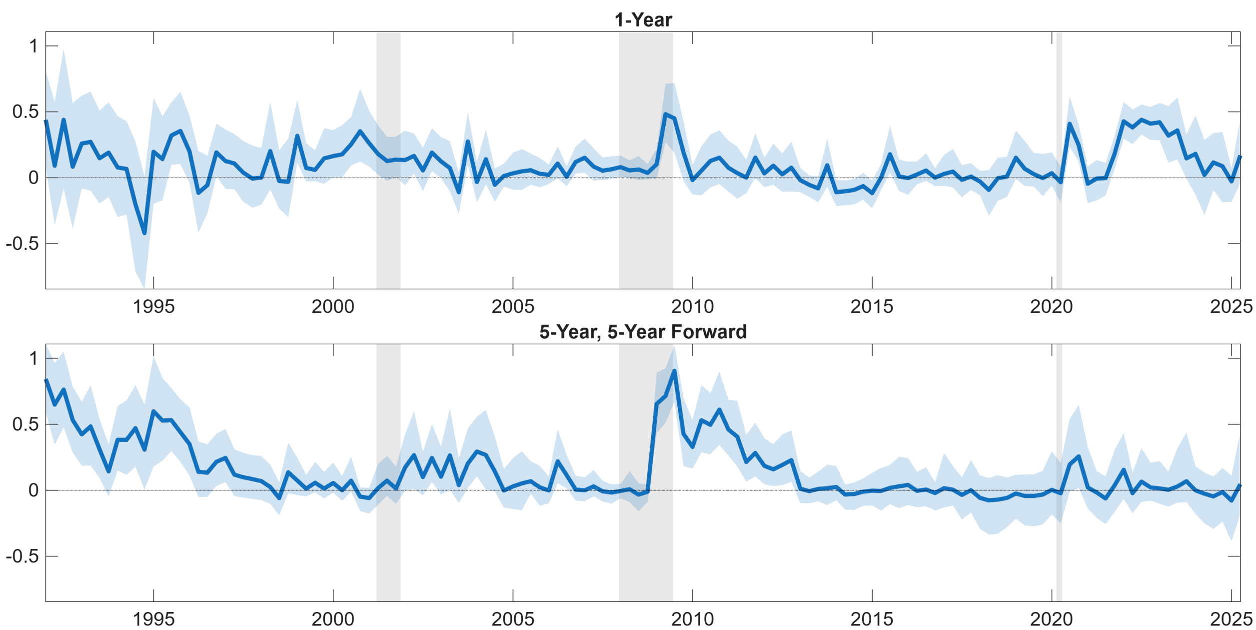

Notes: Counterpart of Figure 5 in the main paper under the heterogeneous-idiosyncrasy specification. See the notes to that figure for variable definitions.

Sources: Authors’ calculation

A11 Curvature Robustness

The baseline term-structure model uses two common factors, level and slope, and omits the curvature factor of the Nelson and Siegel (1987) representation. This appendix adds curvature as a third common factor, with its own loading and shape term, and re-estimates the model, so that each forecaster's term structure is described by level, slope, and curvature factors.

Table A4 reports the posterior parameter statistics. The curvature loadings are small: the average loading on the common curvature factor is -0.33 and the curvature fixed effect is -0.16, both modest relative to level and slope. Figure A36 reports the disagreement decomposition, the counterpart of Figure 5 in the main paper. It matches the baseline closely in normal times and around the Great Recession; as in the heterogeneous-idiosyncrasy specification of Section A10, the early-COVID spike in the public-information share at the 5-year, 5-year-forward horizon is smaller than in the baseline (about 0.4 against 1.0).

Adding curvature therefore does not materially change the results. Level and slope already describe the individual term structures well, in part because the slope is modeled with an exponential function that absorbs much of the nonlinearity across horizons. This is consistent with the discussion in the main text: curvature can emerge in the aggregate term structure from averaging heterogeneous individual level-and-slope shapes, even when each individual structure carries no separate curvature factor.

Table A4: POSTERIOR PARAMETER DISTRIBUTION STATISTICS (CURVATURE SPECIFICATION)

Parameter	Median	95% CI	Parameter	Median	95% CI
$\alpha_{L,i}$	2.2830	[2.1537, 2.3994]	$\sigma_{v,4}$	0.4014	[0.3922, 0.4108]
$\alpha_{S,i}$	0.2117	[0.0759, 0.3430]	$\sigma_{v,5}$	0.1889	[0.1753, 0.2030]
$\alpha_{C,i}$	-0.1598	[-0.2719, -0.0452]	$\sigma_{v,6}$	0.1864	[0.1725, 0.2008]
$\beta_{L,i}$	0.1786	[0.1397, 0.2288]	$\sigma_{v,7}$	0.1846	[0.1706, 0.1989]
$\beta_{S,i}$	0.3719	[0.3040, 0.4476]	$\sigma_{v,8}$	0.0699	[0.0601, 0.0812]
$\beta_{C,i}$	-0.3285	[-0.4117, -0.2524]	$\sigma_{v,9}$	0.2767	[0.2589, 0.2962]
ρ_S	0.8276	[0.6901, 0.9568]	$\sigma_{v,10}$	0.2860	[0.2677, 0.3063]
ρ_C	0.5816	[0.3680, 0.7763]	$\sigma_{v,11}$	0.2589	[0.2420, 0.2777]
$\bar{\rho}_L$	0.9377	[0.9175, 0.9561]	$\sigma_{v,12}$	0.2530	[0.2358, 0.2718]
$\bar{\rho}_S$	0.6320	[0.5817, 0.6793]	$\sigma_{v,13}$	0.1830	[0.1710, 0.1960]
$\bar{\rho}_C$	0.8068	[0.7742, 0.8368]	$\sigma_{v,14}$	0.2043	[0.1913, 0.2185]
σ_L	0.1168	[0.1084, 0.1254]	$\sigma_{v,15}$	0.2310	[0.2168, 0.2465]
σ_S	0.4610	[0.4383, 0.4837]	$\sigma_{v,16}$	0.2464	[0.2320, 0.2622]
σ_C	0.6700	[0.6289, 0.7116]	$\sigma_{v,17}$	0.1802	[0.1691, 0.1919]
λ	0.3194	[0.3071, 0.3321]	$\sigma_{v,18}$	0.1740	[0.1629, 0.1860]
$\sigma_{v,1}$	0.5499	[0.5348, 0.5651]	$\sigma_{v,19}$	0.1592	[0.1482, 0.1707]
$\sigma_{v,2}$	0.4591	[0.4486, 0.4699]	$\sigma_{v,20}$	0.2083	[0.1963, 0.2211]
$\sigma_{v,3}$	0.4292	[0.4195, 0.4392]			

Notes: The "Median" column reports the posterior median of the corresponding parameter and the "95**Sources**: Authors' calculation

Figure A36: DISAGREEMENT VARIANCE DECOMPOSITION (CURVATURE SPECIFICATION)

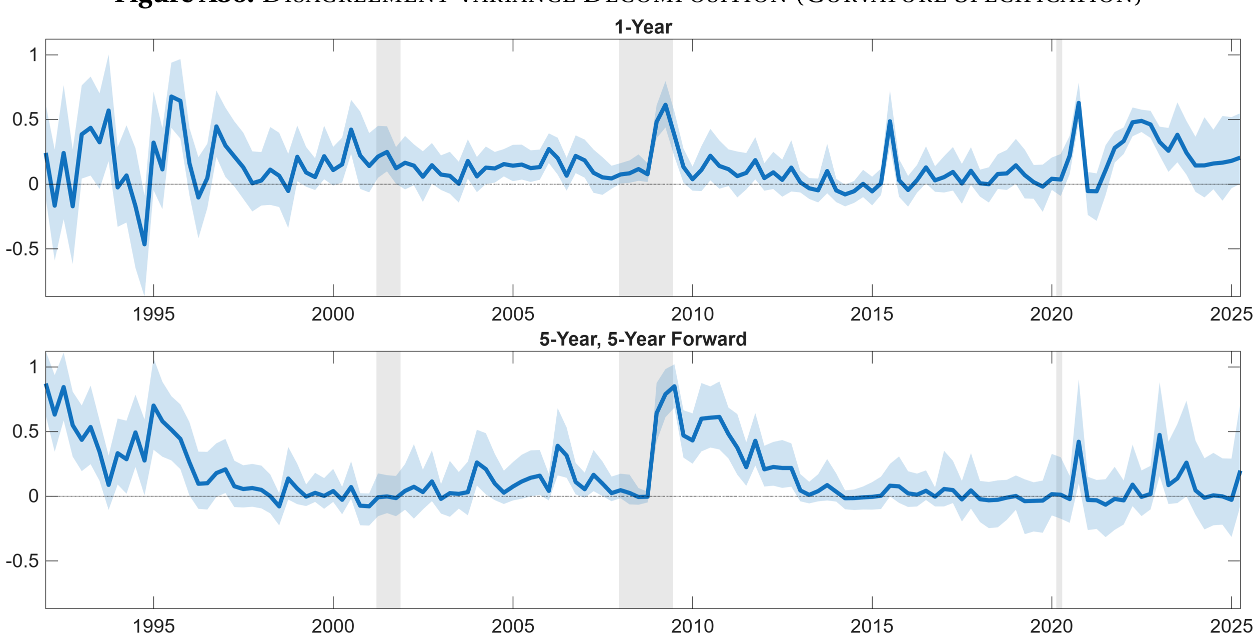

Notes: Counterpart of Figure 5 in the main paper under the three-factor specification that adds curvature. See the notes to that figure for variable definitions.

Sources: Authors' calculation

References

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

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

Bai, J., Ng, S., 2013. Principal components estimation and identification of static factors. Journal of Econometrics 176, 18–29.

Barbera, A., Xia, D., Zhu, S., 2023. The term structure of inflation forecasts disagreement and monetary policy transmission. BIS Working Papers 1114. Bank for International Settlements.

Bauer, M.D., Swanson, E.T., 2023. An alternative explanation for the "Fed information effect". American Economic Review 113, 664–700.

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

Bundick, B., Smith, A.L., 2025. Did the Federal Reserve Break the Phillips Curve? Theory and Evidence of Anchoring Inflation Expectations. The Review of Economics and Statistics 107, 1310–1326.

Carvalho, C., Eusepi, S., Moench, E., Preston, B., 2023. Anchored inflation expectations. American Economic Journal: Macroeconomics 15, 1–47.

Castillo-Martinez, L., Reis, R., 2026. How do central banks control inflation? A guide for the perplexed. Journal of Economic Literature 64, 195–245.

Chamberlain, G., Rothschild, M., 1983. Arbitrage, factor structure, and mean-variance analysis on large asset markets. Econometrica 51, 1281–1304.

Coibion, O., Gorodnichenko, Y., 2025. Inflation, Expectations and Monetary Policy: What Have We Learned and to What End? Working Paper 33858. National Bureau of Economic Research.

Crump, R.K., Eusepi, S., Moench, E., Preston, B., 2023. The term structure of expectations, in: Bachmann, R., Topa, G., van der Klaauw, W. (Eds.), Handbook of Economic Expectations. Academic Press, pp. 507–540.

Diebold, F.X., Li, C., Yue, V.Z., 2008. Global yield curve dynamics and interactions: a dynamic Nelson–Siegel approach. Journal of Econometrics 146, 351–363.

Dong, D., Liu, Z., Wang, P., Wei, M., 2024. Inflation Disagreement Weakens the Power of Monetary Policy. Working Paper 2024-27. Federal Reserve Bank of San Francisco.

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

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

Farmer, L., Nakamura, E., Steinsson, J., 2024. Learning About the Long Run. Journal of Political Economy 132, 3334–3377.

Fisher, J.D.M., Melosi, L., Rast, S., 2025. Long-Run Inflation Expectations. Working Paper 2025-03. Federal Reserve Bank of Chicago.

Fofana, S., Patzelt, P., Reis, R., 2025. Household disagreement about expected inflation, in: Ascari, G., Trezzi, R. (Eds.), Research Handbook on Inflation. Edward Elgar Publishing. chapter 15, pp. 335–357.

Glas, A., 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., Winkler, F., 2026. The Factor Structure of Disagreement. Journal of Business & Economic Statistics 44, 677–690.

Husted, L., Rogers, J., Sun, B., 2020. Monetary policy uncertainty. Journal of Monetary Economics 115, 20–36.

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

Lahiri, K., Sheng, X., 2008. Evolution of forecast disagreement in a Bayesian learning model. Journal of Econometrics 144, 325–340.

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

Nelson, C.R., Siegel, A.F., 1987. Parsimonious modeling of yield curves. Journal of Business 60, 473–489.

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

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