The Macroeconomics of Lumpy Forecasts*

Isaac Baley†
UPF, CREI, BSE and CEPR

Javier Turén ‡
Pontificia Universidad Católica de Chile

February 14, 2025

Work in progress, for the latest version: Click here

Abstract

The recent surge in inflation has highlighted the need to understand how inflation volatility shapes the formation of expectations. While inflation forecasts are typically lumpy, we document that periods of high volatility lead to more frequent and larger expectation revisions, making expectations appear more flexible. Despite this increased responsiveness, forecast accuracy declines, suggesting agents react more often but with greater uncertainty. Using a lumpy forecast model, we show that volatility-driven expectation revisions accelerate the pass-through of inflation shocks. Our findings have important implications for monetary policy and inflation stabilization.

JEL: D80, D81, D83, D84, E20, E30

Keywords: forecasting, monetary policy, fixed horizon, learning, consensus, adjustment costs

*Baley acknowledges financial support from the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S) and ERC Grant 101041334. Turen acknowledges financial support from the Fondecyt Regular Grant 1231600. Felipe Bordini provided excellent research assistance.

†Universitat Pompeu Fabra, CREI, Barcelona School of Economics, and CEPR, isaac.baley@upf.edu.

‡Pontificia Universidad Católica de Chile, jturen@uc.cl.

1 Introduction

The recent surge in inflation, triggered by supply chain disruptions, fiscal expansions, and energy shocks, has reignited interest in how inflation expectations adjust in turbulent times. While central banks closely monitor expectations to gauge the persistence of inflation, a crucial yet overlooked question is how inflation volatility affects how expectations respond to shocks. This paper shows that when inflation volatility is high, professional forecasters adjust their forecasts more frequently and by larger amounts, yet these revisions do not necessarily improve accuracy. Instead, high volatility makes inflation forecasts more flexible but noisier.

Our key insight is the distinction between forecasts, reported in surveys, and underlying beliefs, which continuously adjust but do not always translate into forecast revisions. While beliefs adjust continuously, forecasts may remain sticky due to fixed revision costs and strategic concerns (Baley and Turen, 2024). Beliefs become more volatile in uncertain times as Bayesian forecasters weigh more on new, volatile information relative to their prior knowledge. This increased belief volatility triggers frequent forecast updates, reducing the gap between survey forecasts and internal beliefs. This mechanism provides a new perspective on why expectations appear more responsive in turbulent periods: it is not that forecasters become systematically better at predicting inflation, but rather that they update their forecasts more often as the relative cost of maintaining a stable forecast rises.

We classify tranquil periods (2010–2020) and turbulent periods (2008–2009, 2021–2023) based on monthly CPI inflation volatility in the United States. Using the Bloomberg's Economic Forecasts (ECFC) survey of professional forecasters, we identify distinct regimes of expectation formation. In tranquil periods, forecasts are lumpy and updated infrequently, even as beliefs fluctuate in the background. However, forecasts track beliefs much more closely in turbulent periods, with increased frequency and magnitude of revisions. The pass-through of inflation shocks is stronger in high-volatility environments, as forecast updates reflect the increased uncertainty in beliefs.

To explain these findings, we extend the model of lumpy forecasts developed by Baley and Turen (2024) to incorporate inflation volatility regimes and imperfect information about the true inflation process. Our framework introduces three key frictions: fixed costs of revising expectations, which create inaction in forecast adjustments; strategic concerns, where forecasters prefer to remain close to the consensus; and private signals, leading to heterogenous revisions. In this model, beliefs adjust continuously as agents receive new information, but forecasts remain unchanged unless the change in beliefs is significant enough to justify revision costs.

When volatility is low, beliefs fluctuate within a stable range, rarely crossing the revision threshold. However, when volatility increases, beliefs become more volatile, leading forecasts to adjust more frequently. Periods of large volatility push agents' beliefs out of their inaction regions more often, triggering more frequent updates (a volatility effect). At the same time, the higher volatility widens inaction regions; thus, conditional on revising, it is done for a larger amount (an option effect). We show that the volatility effect dominates, so forecasts are revised more often in uncertain times. This mechanism provides a unified explanation for why inflation expectations become more flexible yet less accurate when inflation volatility surges.

After calibrating the parameters that discipline the underlying frictions—the fixed revision costs, the strength of strategic complementarities, and the private signal noise—using data from tranquil periods, we conduct a series of exercises to assess the how forecast behavior responds to different macroeconomic conditions, emphasizing the state-dependent nature of expectation formation. We explore how forecasts and beliefs adjust when the volatility of inflation changes, considering different informational environments where forecasters may or may not be aware of these shifts. We then extend the analysis to study changes in the mean of inflation, distinguishing between inflationary and deflationary turbulent periods, and finally, assess how large inflation shocks influence the dynamics of forecast revisions. By systematically varying the state of the economy, we aim to disentangle the role of different frictions in expectation formation and understand how forecasters adapt to changing macroeconomic regimes.

When studying variation in inflation volatility, we consider three regimes to analyze how different information levels shape forecast dynamics. In the full information regime, all parameters of the inflation process are common knowledge, and forecasters adjust optimally when they change. In the unknown regime, forecasters do not observe changes in volatility and assume it remains constant, leading to policy inertia. In the learning regime, forecasters attempt to infer volatility changes over time, filtering new information to update their expectations dynamically (Baley and Veldkamp, 2025). These results highlight that volatility-driven expectation adjustments can accelerate inflation pass-through when uncertainty remains high.

Our analysis also examines how changes in the unconditional mean of inflation influence forecast behavior, linking to the debate on whether inflation shocks are transitory or persistent. By distinguishing between deflationary and inflationary turbulent periods, we assess how forecasters react when the long-run level of inflation shifts. If agents know whether the change is temporary or permanent, their revisions adjust accordingly, but if they must infer its persistence over time, their expectations evolve more gradually. This exercise helps disentangle whether the rise in forecast errors during turbulent periods is driven by volatility alone or by uncertainty about structural shifts in inflation dynamics.

Besides the role of changes in the underlying inflation process, we also study the effects of large inflation shocks. Inspired by Cavallo, Lippi and Miyahara (2024), we analyze how major shocks disrupt inaction thresholds and accelerate forecast revisions. We recover individual forecasts and consensus beliefs by feeding the model with actual inflation data. Mimicking results in price-setting; we find that large shocks travel fast in the context of expectation formation.

Our findings have direct implications for monetary policy. Standard stabilization tools assume that expectations adjust smoothly to shocks, but our results suggest that inflation volatility affects expectations' responsiveness, making pass-through faster in uncertain times. Moreover, forecasters revise expectations more frequently but with greater error, complicating monetary communication. Finally, the challenge of anchoring expectations grows in volatile environments, requiring stronger policy credibility. These insights suggest that volatility should be explicitly incorporated into macroeconomic expectations models, as uncertainty alters how inflation shocks propagate through the economy.

Contributions Past research has shown that heightened economic uncertainty leads firms to adjust prices more frequently, as seen in studies of menu costs and price-setting under aggregate or idiosyncratic volatility (Vavra, 2014; Baley and Blanco, 2019). We extend this idea to inflation expectations, demonstrating that forecasters revise their expectations more frequently and by larger amounts when inflation volatility is high, making expectations appear less lumpy. This increased responsiveness may accelerate the pass-through of inflation shocks to expectations, raising concerns about potential de-anchoring when inflation remains volatile for extended periods.

Our findings suggest that the responsiveness of expectations varies with the state of the economy, which has important implications for monetary policy. When expectations adjust more quickly in uncertain times but with greater error, policymakers may face new challenges in anchoring inflation expectations, particularly in periods of heightened volatility. By linking the macroeconomics of lumpy price adjustments to the formation of expectations, our paper provides a new perspective on how inflation shocks propagate through the economy.

This insight also offers a new explanation for why forecasts and beliefs converge in high-volatility periods, providing an alternative to attention-based models such as Pfäuti (2024), which suggests that inflation expectations become more sensitive when inflation exceeds a threshold. While rational inattention models posit that agents suddenly start processing information differently, our model provides a complementary mechanism: forecasters always track inflation, but the cost of revising and reporting forecasts changes with volatility.

The remainder of the paper is structured as follows. Section 2 describes the empirical data, the classification of tranquil and turbulent periods, and key stylized facts. Section 3 presents new evidence on forecast lumpiness across volatility regimes. Section 4 develops the theoretical model, extending the lumpy forecasts model in Baley and Turen (2024) to various regimes regarding knowledge about the inflation stochastic process. Section 5 conducts policy experiments, including analyzing large shocks and expectation pass-through. Section 6 concludes with policy implications and directions for future research.

2 Forecasting in Tranquil and Turbulent Times

We describe the data sources, tranquil and turbulent years samples, and the fixed-event forecasting framework.

2.1 Inflation

We construct annual inflation using the Consumer Price Index (CPI) index. Let cpi_j be the CPI measured j months before the end of year and $\overline{cpi}_\tau = \frac{1}{12}\sum_{j=0}^{11} cpi_j$ be the average CPI in year τ. The annual inflation rate π_τ in any year τ is calculated as

$$\pi_\tau = \frac{\overline{cpi}_\tau - \overline{cpi}_{\tau-1}}{\overline{cpi}_{\tau-1}} \tag{1}$$

Following Giacomini, Skreta and Turen (2020), we can alternatively express the annual inflation rate using the CPI in month h, cpi_h, as follows:

$$\pi_\tau \cong \sum_{h=0}^{11} x_h, \quad \text{with} \quad x_h = \frac{1}{12}\left(\frac{cpi_h - cpi_{h+12}}{cpi_{h+12}}\right), \quad \forall h = 11, \ldots, 0. \tag{2}$$

Figure I – Inflation in the US

(a) Average Montly Inflation **(b)** Volatility of Montly Inflation

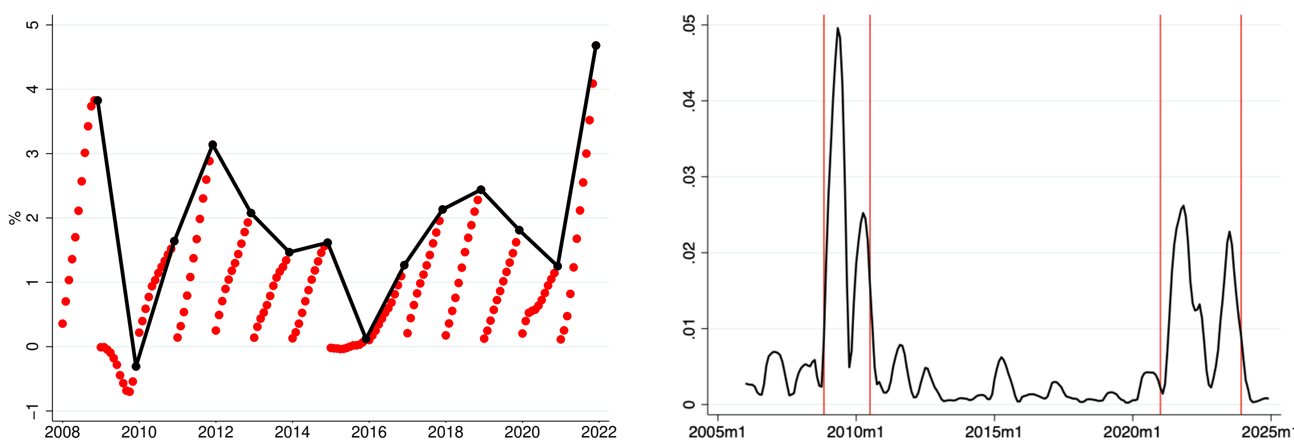

Notes: Bloomberg data. tranquil years = 2010-2019. Turbulent years = 2008-2009 and 2021-24. Panel (a) is the rolling mean of monthly inflation series x_h, and Panel (b) is the standard deviation. Shaded areas represent periods with heightened volatility.

Inflation regimes Figure I plots the dynamics of monthly inflation x_h. Panel (a) shows the average, and Panel (b) shows the standard deviation. We split the sample between tranquil (2010-

2020) and turbulent years (2008-2009 & 2021-2023) according to the volatility of inflation, which is significantly heightened in those years. Turbulent years consist of the Great Recession and the sequel to the COVID-19 pandemic, periods in which the economic conditions were far from ordinary.

Table I – Two Inflation Regimes

		All years	**Tranquil**	**Turbulent I**	**Turbulent II**
Periods		2008–2024	2010–2021	2008 – 2009	2021–2023
Average Inflation	$\mathbb{E}[\pi]$	2.45	1.72	1.75	5.6
Inflation Volatility	$\sigma(\pi)$	0.07	0.06	0.13	0.11
Observations	N	17,663	9,256	3,363	5,044

Table I provides summary statistics across the two inflation regimes. Inflation volatility is significantly higher in turbulent periods, reaching 0.13 in Turbulent I (2008–2009) and 0.011 in Turbulent II (2021–2023), compared to only 0.06 in tranquil times. While both turbulent periods share high volatility, their average inflation levels differ. Turbulent I (2008–2009) is characterized by relatively low inflation at 1.75 percent, reflecting deflationary pressures in the aftermath of the financial crisis. In contrast, Turbulent II (2021–2023) exhibits significantly higher inflation, averaging 5.6 percent. This distinction highlights that although both episodes were marked by volatility, they represent different macroeconomic conditions—one dominated by deflationary risk and the other by strong inflationary pressures.

2.2 Forecasts

We analyze revisions of US annual (year-on-year) CPI inflation forecasts from the "Economic Forecasts ECFC" survey of professional forecasters conducted by Bloomberg. This survey is comparable to other surveys of professional forecasters regarding the number of participants and the background of forecasters, i.e., financial institutions and banks, consulting firms, universities, and research centers. Besides its similarities, one of the most appealing features of this particular dataset is that the most recent forecasts of any other forecaster, the date when each prediction was lastly updated, and the consensus forecast (the mean forecast) are visible to users of the Bloomberg terminal in *real time*.

Fixed-event forecasting We examine monthly fixed-event forecasts of annual US inflation. For each year, we consider participants who forecast inflation for all 12 months before the final figure (end-of-year inflation) is officially published. We remove forecasters that fail to provide at least one annual inflation revision. This leaves approximately 100 forecasters per year. The dataset contains the history of forecast updates for all forecasters over the 12-month horizon (h) before the end of each year.

Within each year, we denote with f_h^i the inflation forecast of forecaster i at horizon h (to save on notation, we do not explicitly use the year):

$$f_h^i, \quad i = 1...N, \quad h = 12, ...1. \tag{3}$$

We index the horizon backward so that the index $h = 12, ..., 1$ indicates that the forecast was produced h months before the end of each corresponding year (the fixed event). We also define the consensus forecast as the average of individual forecasts at each horizon:

$$F_h \equiv \frac{1}{N} \sum_{i=1}^{N} f_h^i \, \mathrm{d}i. \tag{4}$$

Bloomberg reports the consensus F_h in the terminal and is available in real time.

2.3 Forecasts revisions

At any given year, we define the forecast revision at horizon h, denoted by Δf_h^i, as the one-period difference between the forecast in two consecutive horizons (measured in percentage points):

$$\Delta f_h^i \equiv f_h^i - f_{h+1}^i. \tag{5}$$

Table II summarizes forecast revisions. Revisions are larger and more dispersed in turbulent periods (mean: 0.028, size: 0.453, variance: 0.227) than in tranquil periods (mean: −0.013, size: 0.247, variance: 0.055). Forecasters revise more frequently in turbulent times (5.60 vs. 5.06 revisions per year). The inaction rate drops sharply in turbulence (39.2% vs. 56.9%).

Table II – Summary Statistics of Forecast Revisions

		All	**Tranquil**	**Turbulent**
Average	$\mathbb{E}(\Delta f)$	−0.002	−0.013	0.028
Size	$\mathbb{E}(abs(\Delta f)\|\Delta f \neq 0)$	0.307	0.247	0.453
Dispersion	$\mathbb{V}ar(\Delta f)$	0.104	0.055	0.227
Number of revisions	$count(\Delta f \neq 0)$	5.204	5.059	5.602
Duration (months)	$\mathbb{E}(\tau)$	1.497	1.594	1.231
Inaction rate	$\Pr(\Delta f = 0)$	0.523	0.569	0.392
Frequency	$\Pr(\Delta f \neq 0)$	0.444	0.427	0.492
Observations	N	12,619	9,256	3,363

Notes: Bloomberg data. Tranquil = 2010-2020. Turbulent = 2008-09, 2021-23. Averages across years and horizons.

Next, we examine the term structure of forecast revisions and errors, analyzing how they evolve across forecasting horizons. Figure II presents two key patterns: Panel (a) shows the frequency of forecast revisions, while Panel (b) depicts the absolute size of revisions. Both panels reveal that as the forecasting horizon h shrinks, revisions become less frequent and smaller in magnitude. This pattern is consistent with forecasters making larger and more frequent adjustments for distant horizons, where uncertainty is greater, whereas near-term forecasts require fewer updates. However, a key distinction emerges across different volatility regimes. In turbulent years, forecast revisions are systematically more frequent and larger at every horizon compared to tranquil times. The gap between tranquil and turbulent periods is particularly pronounced at longer horizons.

Figure II – Forecast revisions across horizons

(a) Frequency of Revisions

(b) Size of Revisions

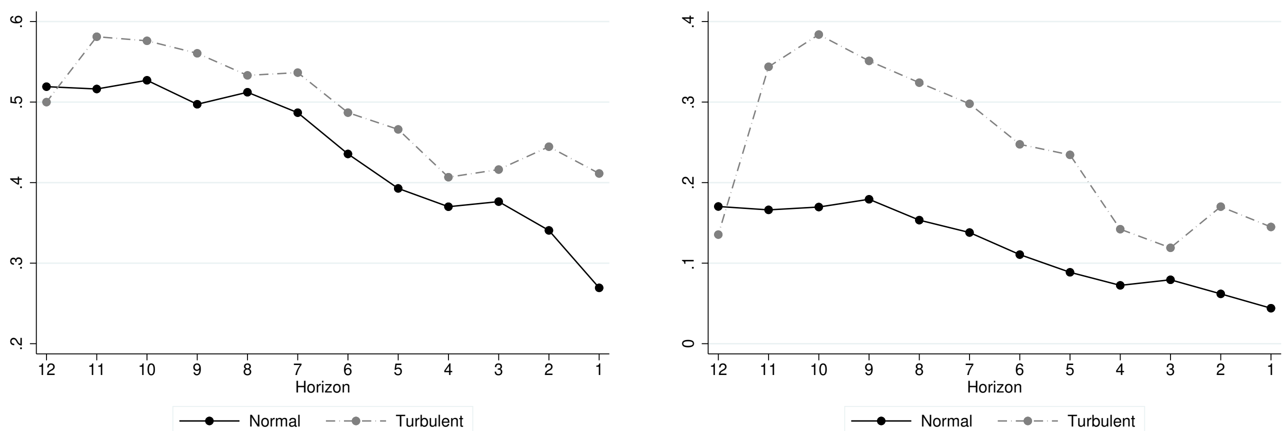

Notes: Bloomberg data. Tranquil = 2010-2020. Turbulent = 2008-09, 2021-23.

2.4 Forecast errors

At any given year, we define the ex-post forecast error e_h^i of individual i at horizon h as the difference between the actual end-of-year inflation π and the forecast f_h^i.

$$e_h^i \equiv \pi - f_h^i \tag{6}$$

Table II summarizes the properties of forecast errors. On average, errors are negative in tranquil times (-0.055), indicating a tendency to overpredict inflation, whereas in turbulent periods, errors are positive (0.237), suggesting underprediction during high-volatility episodes. Forecast errors are larger in turbulent times, with the absolute error size increasing from 0.305 in tranquil periods to 0.789 in turbulent periods, and more volatile, reflected in the higher dispersion of errors in turbulent periods.

Table III – Summary Statistics of Forecast Errors

		All	**Tranquil**	**Turbulent**
Average	$\mathbb{E}(e)$	0.023	-0.055	0.237
Size	$\mathbb{E}(abs(e))$	0.434	0.305	0.789
Positive	$\Pr(e > 0)$	0.414	0.345	0.606
Negative	$\Pr(e < 0)$	0.510	0.572	0.340
Dispersion	$\sigma(e)$	0.663	0.502	1.105
Observations	N	12,619	9,256	3,363

Notes: Bloomberg data. Tranquil = 2010-2020. Turbulent = 2008-09, 2021-23. Averages across years and horizons.

Figure III shows the term structure of forecast errors. Panel (a) presents the Root Mean Square Error (RMSE) across forecast horizons, showing that errors are consistently larger in turbulent periods. The gap between the two regimes is most pronounced at longer horizons, where forecasting is inherently more complex. While the difference narrows at shorter horizons, this pattern is expected given the fixed-horizon structure, as near-term forecasts are naturally more accurate.

Panel (b) displays the hazard rate of forecast revisions, which measures the probability that a forecast is updated as it ages. In both regimes, the hazard rate declines with age, indicating that older forecasts are less likely to be revised. However, revisions occur more frequently in turbulent periods, particularly at short ages, suggesting that forecasters respond more quickly to new information in high-uncertainty environments.

Figure III – Forecast errors across horizons

(a) Root Mean Square Error **(b)** Hazard Rate

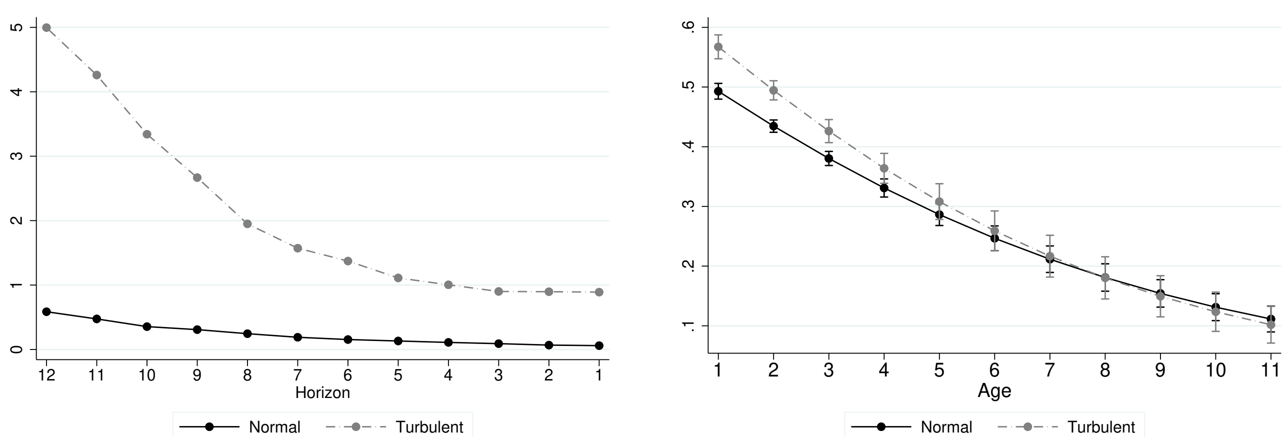

Notes: Bloomberg data. Normal years = 2010-2019. Turbulent years = 2008-2009 and 2020-21.

3 A Structural Model of Lumpy Forecasts

We develop a horizon-dependent fixed-event Bayesian forecasting model with private information, frequent information revelation, fixed revision costs, and strategic concerns. In addition, we introduce an information friction where forecasters may perceive inflation volatility differently from its actual value.

3.1 Forecasting Problem

Many forecasters, indexed by $i \in N$, generate forecasts of end-of-year inflation π. End-of-year inflation inflation π equals the sum of within-year monthly inflations x_h, namely $\pi \equiv \sum_{h=1}^{12} x_h$.

Payoffs At each horizon h, forecaster i chooses a forecast f_h^i based on their information set $\mathcal{I}_h^i$. Changing a forecast entails paying a fixed revision cost $\kappa > 0$ measured in utility units. For a given initial forecast f_{13}^i, forecasts minimize the yearly sum of monthly quadratic losses:

(7)
$$\min_{\{f_h^i\}_{h=12}^1} \mathbb{E}\left[\sum_{h=12}^{1} \underbrace{(f_h^i - \pi)^2}_{\text{accuracy}} + r \underbrace{(f_h^i - F_h)^2}_{\text{strategic}} + \underbrace{\kappa \mathbb{1}_{\{f_h^i \neq f_{h+1}^i\}}}_{\text{stability}} \middle| \mathcal{I}_0^i\right].$$

The first term in the payoff function is the distance between the forecast and the actual end-of-year inflation, reflecting losses from the lack of *accuracy*. The second term is the distance between the forecast and the consensus (the average) $F_h = N^{-1} \sum_{i=1}^{N} f_h^i$, multiplied by the parameter r that measures the strength of *strategic concerns*. If $r > 0$, there is strategic complementarity, as the payoff increases when the forecast is close to the consensus. If $r < 0$, there is strategic substitutability, as the payoff increases when the forecast is far from the consensus. The third term is the fixed cost $\kappa > 0$ paid for any forecast revision, capturing preference for *forecast stability*. We think of κ and r as fundamental frictions that do not change with the aggregate state of the economy.

Recursive problem Let $\hat{\pi}_h^i \equiv \mathbb{E}[\pi | \mathcal{I}_h^i]$ and $\Sigma_h^\pi \equiv \mathbb{E}[(\hat{\pi}_h^i - \pi)^2 | \mathcal{I}_h^i]$ be the conditional mean and variance of end-of-year inflation beliefs. Let $\hat{F}_h \equiv \mathbb{E}[F_h | \mathcal{I}_h^i]$ and $\Sigma^F \equiv \mathbb{E}[(\hat{F}_h - F_h)^2 | \mathcal{I}_h^i]$ be the conditional mean and variance of consensus beliefs. Then, for given initial forecasts f_{13}^i, forecasters solve the following problem:

(8)
$$\min_{\{f_h^i\}_{h=12}^1} \sum_{h=12}^{1} \Sigma_h + (f_h^i - \hat{\pi}_h^i)^2 + r(f_h^i - \hat{F}_h)^2 + \kappa \mathbb{1}_{\{f_h^i \neq f_{h+1}^i\}},$$

where $\Sigma_h \equiv \Sigma_h^\pi + r\Sigma^F$ is a weighted sum of inflation and consensus uncertainty. Thus, Σ_h accounts for the unforecastable part of the process at each horizon h.

Next, we describe the actual inflation process and belief formation.

3.2 Inflation process and information regimes

Actual monthly inflation follows an autoregressive process of order 1:

$$x_h = c_x + \phi_x x_{h+1} + \varepsilon_h^x, \qquad \varepsilon_h^x \sim \mathcal{N}(0, \sigma_x^2), \tag{9}$$

where c_x is a constant, ϕ_x is the persistence parameter, and ε_h^x is an *iid* normally distributed noise with volatility σ_x^2.

Three regimes We consider three different regimes under which forecasters form their expectations about inflation realizations:

1. *Known Volatility Regime:* In this scenario, forecasters have complete knowledge of the process generating inflation (but still ignore the actual realizations and must forecast them). Since all parameters are common knowledge, agents adjust their expectations immediately and optimally whenever these parameters change. This represents the benchmark case in which there is no uncertainty about the structure of the process governing inflation.

2. *Unknown Volatility Regime:* Here, forecasters know the mean and persistence of the inflation process but do not observe changes in volatility σ_x^2. They assume that the volatility remains constant at a perceived level $\tilde{\sigma}_x^2$, even when the true volatility differs. Because agents do not recognize shifts in volatility, their forecasting policies remain unchanged, and only realized shocks drive the observed forecast dynamics. This regime isolates the impact of surprise volatility shifts on forecast errors and adjustments.

3. *Learning Regime:* In this case, forecasters do not immediately know when volatility changes but attempt to learn about it over time. They update their beliefs about $sigma_x 2$ using observed inflation data and apply a filtering mechanism. Agents in this regime gradually incorporate new information, leading to an evolving forecasting strategy that adapts as they refine their beliefs about the underlying inflation volatility.

These three regimes allow us to assess how different levels of information about the inflation process influence the frequency, size, and accuracy of forecast revisions. By comparing them, we can evaluate how forecasters respond to changing economic conditions when they are fully informed, unaware of shifts, or actively learning about them. From now on, we distinguish perceived $\tilde{\sigma}_x^2$ from actual volatility σ_x^2.

3.3 Signals and information dynamics

Public signal At the beginning of each horizon h, previous monthly inflation x_{h+1} is revealed, reflecting the official release from the statistical agency. Previous inflation and the AR(1) assumption imply a public signal about current *monthly* inflation:

$$x_h^{AR} \equiv \mathbb{E}[x_h|x_{h+1}] = c_x + \phi_x x_{h+1}. \tag{10}$$

The variance of the public signal is $\tilde{\sigma}_x^2 = \mathbb{V}ar[x_h|x_{h+1}] = \mathbb{V}ar[\varepsilon_h^x]$.

Private signal Following Patton and Timmermann (2010), at the beginning of each horizon, each forecaster receives an unbiased private signal $\widetilde{x}_h^i$ about what inflation in that month will be (recall that the actual monthly inflation is only released at the end of the month):

$$\widetilde{x}_h^i = x_h + \zeta_h^i, \qquad \zeta_h^i \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_\zeta^2). \tag{11}$$

The idiosyncratic signal noise σ_ζ^2 reflects the heterogeneity in beliefs, private information, or models across agents.

Information dynamics At the end of the period, and after f_h^i is decided, monthly inflation x_h and the consensus forecast F_h are observed by everyone. These timing assumptions eliminate a fixed point between individual choices and the consensus, as in a beauty contest (Morris and Shin, 2002), greatly simplifying the model solution with revision costs. Therefore, the individual information set $\mathcal{I}_h^i$ at the time of choosing the forecast is

$$\mathcal{I}_h^i = \widetilde{x}_h^i \cup \mathcal{I}_h = \widetilde{x}_h^i \cup \{x_{h+1}, x_{h+2}, \ldots, F_{h+1}, F_{h+2}, \ldots\}. \tag{12}$$

We denote the public information set at horizon h as $\mathcal{I}_h \equiv \{(x_j, F_j) : j \geq h+1\}$, which includes releases of past inflation and past consensus.

3.4 Belief Formation

Consensus Beliefs The consensus is the average forecast $F_h = N^{-1} \sum_{i=1}^{N} f_h^i$. However, since the consensus is observed with delay (e.g., at horizon h, F_{h+1} is observed), forecasters must form expectations about the contemporaneous consensus when choosing their forecasts. Forecasters entertain random walk beliefs:

$$F_h = F_{h+1} + \varepsilon_h^F, \qquad \varepsilon_h^F \sim \mathcal{N}(0, \sigma_F^2), \tag{13}$$

where volatility σ_F^2 is common knowledge. Given this assumption, the common horizon-specific consensus beliefs are $F_h|\mathcal{I}_h^i \sim \mathcal{N}(F_{h+1}, \sigma_F^2)$. Our equilibrium definition below specifies the consistency of these beliefs.

Monthly Inflation Beliefs Forecasters combine the public signal x_h^{AR} in (10) and their private signal $\tilde{x}_h^i$ in (11) to construct an individual monthly inflation belief $\hat{x}_h^i$:

$$(14) \qquad \hat{x}_h^i \equiv \mathbb{E}[x_h|\mathcal{I}_h^i] = \frac{\tilde{\sigma}_x^{-2} x_h^{AR} + \sigma_\zeta^{-2}\tilde{x}_h^i}{\tilde{\sigma}_x^{-2} + \sigma_\zeta^{-2}} = (1-\alpha)x_h^{AR} + \alpha\tilde{x}_h^i,$$

where we define the Bayesian weight on the private signal as $\alpha \equiv \sigma_\zeta^{-2}/(\tilde{\sigma}_x^{-2} + \sigma_\zeta^{-2})$. The weight α increases in the precision of the private signal σ_ζ^{-2} and decreases in the precision of inflation $\tilde{\sigma}_x^{-2}$.

End-of-Year Inflation Beliefs At each horizon, forecasters form end-of-year inflation beliefs $\pi|\mathcal{I}_h^i \sim \mathcal{N}(\hat{\pi}_h^i, \Sigma_h^\pi)$ by projecting their monthly beliefs using the AR(1) structure. These beliefs are normal. Forecasters combine past "official" releases $\{x_j\}_{j>h}$ with their individual monthly beliefs $\hat{x}_h^i$ to obtain the conditional mean $\hat{\pi}_h^i$:

$$(15) \qquad \hat{\pi}_h^i = \underbrace{h\left(\frac{c_x}{1-\phi_x}\right) + \frac{1-\phi_x^h}{1-\phi_x}\left(\hat{x}_h^i - \frac{c_x}{1-\phi_x}\right)}_{\text{AR(1) projection using } h \text{ info}} + \underbrace{\sum_{j=h+1}^{12} x_j}_{\text{realized, } j>h}, \quad h = 12, \ldots, 1.$$

The first part of the expression (15) uses the AR(1) statistical model to project the monthly belief $\hat{x}_h^i$ into the future. The second part equals the sum of the true monthly inflation values released to date. The conditional variance $\Sigma_h^\pi \equiv \mathbb{E}[(\pi - \hat{\pi}_h^i)^2]$ is a function of the AR(1) parameters $\{\phi_x, \tilde{\sigma}_x^2\}$ and signal noise σ_ζ^2; it decreases with the horizon and is independent of agents' identity:

$$(16) \qquad \Sigma_h^\pi = [(1-\alpha)^2\tilde{\sigma}_x^2 + \alpha^2\sigma_\zeta^2]\left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 + \frac{\tilde{\sigma}_x^2}{(1-\phi_x)^2}\left[(h-1) - \frac{2\phi_x(1-\phi_x^{h-1})}{1-\phi_x} + \frac{\phi_x^2(1-\phi_x^{2(h-1)})}{1-\phi_x^2}\right].$$

The first term of Σ_h^π corresponds to the uncertainty driven by the AR(1) projection and the noisy signal (weighted by α) for the current release of monthly inflation. Likewise, the second part of (16) reflects the accumulated uncertainty caused by the remaining $(h-1)$ unforecastable shocks that will hit the process until the release date.

Average Beliefs Given the public releases of monthly past values, the AR(1) assumption implies a public signal z_h about *yearly* inflation, given by:

$$(17) \qquad z_h = h\left(\frac{c_x}{1-\phi_x}\right) + \frac{\phi_x(1-\phi_x^h)}{1-\phi_x}\left(x_{h+1} - \frac{c_x}{1-\phi_x}\right) + \sum_{j=h+1}^{12} x_j, \quad h = 12, \ldots, 1.$$

It is useful to establish a relationship between individual beliefs $\hat{\pi}_h^i$ in (15) under the information set $\mathcal{I}_h^i$ and public beliefs z_h in (17) under the information set $\mathcal{I}_h$. The following relationship links individual and common beliefs:

$$\hat{\pi}_h^i = z_h + \nu_h^i, \quad \text{with} \quad \nu_h^i \sim \mathcal{N}\left(0, \left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 \alpha^2(\tilde{\sigma}_x^2 + \sigma_\zeta^2)\right). \tag{18}$$

where α is the updating weight defined in (14).

3.5 Equilibrium

We now define our notion of equilibrium. We focus on a restricted perceptions equilibrium (RPE), representing a slight deviation from rational expectations (Evans and Honkapohja, 1993). We posit that forecasters believe the consensus follows a random walk, and ex-post, they cannot distinguish the actual consensus process from a random walk. Forecasters are *internally rational* (Marcet and Nicolini, 2003), as they use an "internally consistent" learning model. This equilibrium concept delivers enormous tractability by eliminating the fixed point between the consensus and the aggregation of individual forecasts.

Definition 1. *A restricted perceptions equilibrium (RPE) consists of:*

(i) perceived consensus process $\{\hat{F}_h\}$ *given by a function* g *parametrized by* (δ, σ_F)

$$\hat{F}_h = g(\hat{F}_{h+1}, \delta) + \epsilon_h^{\hat{F}}, \quad \epsilon_h^{\hat{F}} \sim \mathcal{N}(0, \sigma_F^2) \tag{19}$$

(ii) inflation beliefs $\{\hat{\pi}_h^i\}$ *and forecasts* $\{f_h^i\}$ *for all agents* i *and horizons* h

such that:

1. *given inflation beliefs* $\{\hat{\pi}_h^i\}$ *in* (15) *and the perceived consensus process* $\{\hat{F}_h\}$ *in* (19)*, forecasts* $\{f_h^i\}$ *are optimal and solve the forecasting problem* (8)*;*

2. *parameters* (δ, σ_F^2) *are such that the forecast errors arising from predicting the actual consensus using the perceived law of motion, i.e.,* $\epsilon_h^F \equiv F_h - g(F_{h+1}, \delta)$*, satisfy:* $\mathbb{C}ov[\epsilon_h^F, \epsilon_j^F] = 0$ $\forall h \neq j$ *and* $\mathbb{V}ar[\epsilon_h^F] = \sigma_F^2$*.*

In the restricted perceptions equilibrium, the actual consensus process given by the aggregation of individual forecasts, $F_h = N^{-1}\sum_{i=1}^N f_h^i$, differs from the prediction. However, in this equilibrium concept, agents are assumed to use the δ, σ_F that best predicts future prices given (19).

3.6 Optimal Forecasting Policy

Proposition 1 writes the problem in recursive form as a stopping-time problem using the principle of optimality. The individual state includes the past forecast, the mean and variance of inflation beliefs, and the mean and variance of consensus beliefs. It is equivalent to working with posterior beliefs instead of the signals. The aggregate state includes past realizations of monthly inflation and consensus. Because total uncertainty evolves deterministically and is shared across agents, we include it in the aggregate state. We thus index value function with the horizon h to account for the aggregate state.

Proposition 1. *The value of a forecaster i at horizon h with state $(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i)$ equals*

$$\mathcal{V}_h(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i) \;=\; \min\{\underbrace{\mathcal{V}_h^I(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i)}_{\text{inaction}},\; \underbrace{\mathcal{V}_h^A(\hat{\pi}_h^i, \hat{F}_h)}_{\text{action}}\} \tag{20}$$

where the value of inaction $\mathcal{V}_h^I$ and the value of action $\mathcal{V}_h^A$ are, respectively,

$$\begin{aligned}
\mathcal{V}_h^I(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i) &= \Sigma_h + (f_{h+1}^i - \hat{\pi}_h^i)^2 + r(f_{h+1}^i - \hat{F}_h)^2 + \mathbb{E}[\mathcal{V}_{h-1}(\hat{\pi}_{h-1}^i, \hat{F}_{h-1}, f_{h+1}^i)|\mathcal{I}_h^i] \\
\mathcal{V}_h^A(\hat{\pi}_h^i, \hat{F}_h) &= \kappa + \Sigma_h + \min_{f_h^i}\left\{(f_h^i - \hat{\pi}_h^i)^2 + r(f_h^i - \hat{F}_h)^2 + \mathbb{E}[\mathcal{V}_{h-1}(\hat{\pi}_{h-1}^i, \hat{F}_{h-1}, f_h^i)|\mathcal{I}_h^i]\right\}
\end{aligned}$$

subject to the evolution of inflation beliefs in (15) and (16), and consensus beliefs in (19).

Inaction region and reset forecast The optimal policy consists of a *horizon-specific* 3-dimensional inaction region $\mathcal{R}_h$ given by the set of states for which the value of inaction (keeping the current forecast) is greater or equal to the value of action (revising the forecast)

$$\mathcal{R}_h \equiv \{(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i) : \mathcal{V}_h^I(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i) \geq \mathcal{V}_h^A(\hat{\pi}_h^i, \hat{F}_h)\}, \tag{21}$$

and a reset forecast $f_h^{i*}(\hat{\pi}_h^i, \hat{F}_h)$ where the forecast is set when revising. Thus, given the current forecast, it remains unchanged if beliefs lie inside the inaction region and resets at any horizon h when those beliefs fall outside it. Revisions are then given by

$$\Delta f_h = \begin{cases} 0 & \text{if} \quad f_{h+1}^i \in \mathcal{R}_h \\ f_h^{i*} - f_{h+1}^i & \text{if} \quad f_{h+1}^i \notin \mathcal{R}_h. \end{cases} \tag{22}$$

Panel (a) in Figure IV shows the forecast revision policy at horizon $h = 6$. We plot it in two dimensions by fixing the current forecast at $f_6^i = 2$ and varying inflation beliefs $\hat{\pi}_h^i$ in the x-axis and consensus beliefs $\hat{F}_h$ in the y-axis. We use the parametrization in Table V. The negative slope in the inaction region arises because the two beliefs are "substitutes" in that a smaller distance to the consensus belief may compensate for a greater distance from the inflation belief, or vice versa.

Figure IV – Forecast Revision Policy: Inaction and Reset

(a) Inaction Region and Reset Forecast $f_h^{i\,*}$

(b) Width of Inaction Region

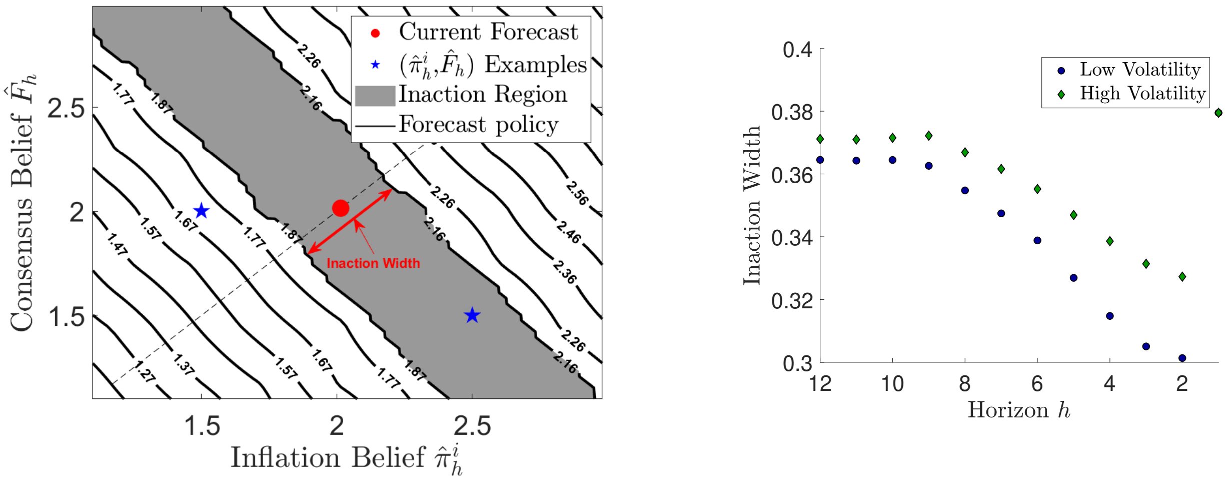

Notes: Panel (a) illustrates the reset forecast policy $f_h^{i\,*}$ and inaction region $\mathcal{R}_h$ for $h = 6$ given a current forecast $f_h^i = 2$. We also show two examples of beliefs, one inside and one outside the inaction region. Panel (b) plots the inaction region width (the segment on the 45-degree line) for different horizons and two levels of inflation volatility σ_x, low and high.

How volatility affects inaction Panel (b) plots the width of the inaction region measured on the 45-degree line, against the forecasting horizon. The width of the inaction region shrinks with the horizon.[1] At long horizons, belief uncertainty is at its highest level; forecasters anticipate that their belief would hit the band very often and thus optimally widen it to minimize adjustment cost payments an *option effect*. As belief uncertainty falls, the option effect is smaller, and the band shrinks. A shrinking inaction region implies that adjustment size falls with the horizon. We plot for high and low volatility. We see clearly the option effect. The impact on the adjustment frequency is nuanced because frequency depends on the option effect and the *volatility effect*.

3.7 Calibration and Solution

Externally set parameters Frequency is monthly. We feed the AR(1) parameters estimated directly from the data. We estimate the AR(1) process parameters using a rolling window for tranquil and turbulent years over the sample years.

For the year-on-year monthly inflation process we estimate $(c_x, \phi_x, \sigma_x) = (0.013, 0.932, 0.036)$ in tranquil times. These values imply an unconditional annual inflation of $\mu_\pi = 12c_x/(1-\phi_x) = 2.23$ with annual volatility $\sigma_\pi^2 = \sigma_x^2 \sum_{h=1}^{12}(1-\phi_x^h)^2/(1-\phi_x^2) = 0.49$. Table IV summarizes this information, describing the parametrization at the monthly frequency and the implied values for the yearly inflation process[2]

[1] We see a widening of the inaction region at $h = 1$ arising from the finite-horizon nature of the problem.

[2] The inflation process estimation details appear in Appendix B.

Table IV – Parameters of Two Inflation Regimes

		Monthly (x)			Yearly (π)	
		Tranquil	**Turbulent**		**Tranquil**	**Turbulent**
Mean	c_x	0.013		c_π	2.23	
Persistence	ρ_x	0.932		ρ_π	0.43	
Volatility	σ_x	0.036		σ_π	0.49	

Table V – Internally calibrated parameters

Parameter		Value	Moment	Data	Model
κ	Revision cost	0.06	$\Pr[\Delta f \neq 0]$	0.43	0.40
r	Strategic concerns	0.73	$\mathbb{E}[abs(\Delta f)\vert adjust]$	0.25	0.19
σ_ζ	Signal noise	0.03	Hazard Slope	-0.04	-0.04
σ_F	Consensus volatility	0.13	Internal Consistency	—	—

Internally calibrated parameters Using the simulated method of moments (SMM), we estimate values for the three remaining parameters by matching the cross-sectional moments in tranquil times. We take these years as a steady-state, which allow us to back out the structural frictions: the fixed revision cost κ, the strength of strategic concerns r, and the private noise σ_ζ.

We target three moments: the frequency of revisions $\Pr[\Delta f \neq 0] = 0.43$, the average absolute value of revisions $\mathbb{E}[abs(\Delta f)||adjust] = 0.25$ and the slope of the hazard rate between horizons 12 and 6 equal to -0.04. The hazard's slope informs idiosyncratic signal noise. Learning is slow when signals are very noisy, and the hazard rate declines slowly. In contrast, learning is faster when signals are less noisy, and the hazard rate declines faster.

Internal consistency of consensus beliefs Forecasters in our model assume a random walk process for the consensus in (13). This assumption imposes structure and disciplines the value of σ_F. Starting with a guess for the volatility of the consensus process σ_F^2, we compute individual decision rules for each horizon h using backward induction. We then simulate the model, calculate the volatility of the realized consensus, and iterate on σ_F^2 to ensure belief consistency.[3]

Estimated Parameters Table V shows the baseline parameterization, the moments in the data, and the model fit. The calibrated parameters are as follows. First, the fixed adjustment cost of $\kappa = 0.05$ implies a preference for forecast stability. Second, the positive value for $r = 0.41$ signals strategic complementarities. Lastly, the private noise $\sigma_\zeta = 0.04$ is as significant as the volatility of the inflation process, $\sigma_x = 0.036$. Given their relative precision, the weight on private signals equals $\alpha = 0.56$. Finally, setting $\sigma_F = 0.11$ delivers consistent consensus beliefs.[4]

[3] Appendix **??** explains the solution algorithm and other computational details.

[4] Appendix **??** presents the details on the consistency of consensus beliefs.

3.8 Untargeted Term Structures

Figure V shows the term structure of the revisions frequency, size, and hazard rate. While we only targeted average values (the dashed lines), the model matches the empirical patterns along the forecasting period. The model can quantitatively match the downward slopping patterns of the frequency of revisions (extensive margin, Figure VIIa) and the size of non-zero revisions (intensive margin, Figure VIIb). In addition, the model accurately matches the hazard's level.

Figure V – Cross-sectional statistics across horizons

(a) Frequency of revisions **(b)** Size of revisions **(c)** Hazard rate

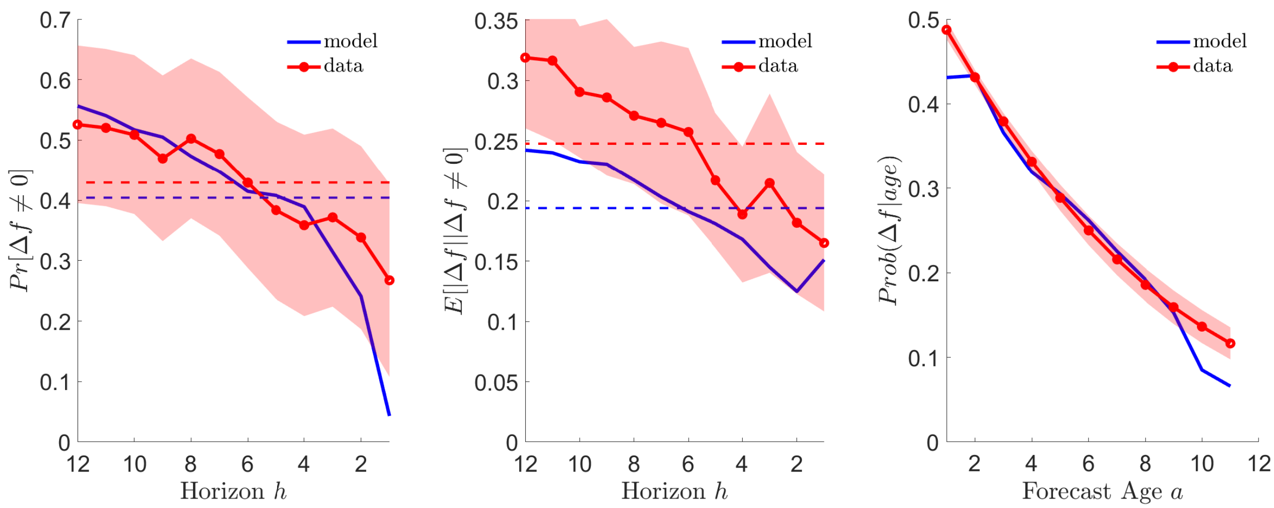

Notes: Cross-sectional moments obtained from the model's simulation under the benchmark calibration.

4 State-dependent Forecasting

This section studies how forecast behavior responds to different macroeconomic conditions, emphasizing the state-dependent nature of expectation formation. We explore how forecasts and beliefs adjust when the volatility of inflation changes, considering different informational environments where forecasters may or may not be aware of these shifts. We then extend the analysis to study changes in the mean of inflation, distinguishing between inflationary and deflationary turbulent periods, and finally, assess how large inflation shocks influence the dynamics of forecast revisions. By systematically varying the state of the economy, we aim to disentangle the role of different frictions in expectation formation and understand how forecasters adapt to changing macroeconomic regimes.

4.1 Changes in inflation volatility

We begin by examining how forecasts and beliefs respond to changes in the underlying volatility of inflation. Inflation volatility has exhibited substantial variation across different periods, with turbulent years displaying nearly twice the level of volatility observed in tranquil times. This increase in volatility can alter forecast behavior through multiple channels, affecting both the frequency and magnitude of revisions. To isolate these effects, we study three cases: one in which forecasters know the volatility regime, one in which the volatility shift is unknown, and one in which agents must learn about the change over time. Each case provides insight into how forecasters process information under different uncertainty conditions and whether their responsiveness to inflation shocks varies with their perception of volatility.

4.1.1 Known volatility change

Figure VI examines the case where forecasters know the volatility regime, allowing them to adjust their expectations optimally when inflation uncertainty rises. The comparison between tranquil (solid red) and turbulent (dashed red) periods in the data shows that both the frequency and size of revisions increase in volatile times. The model (solid blue) successfully captures the baseline revision patterns. When volatility rises by 10 percent (dashed blue), it aligns more closely with the turbulent data, though it underpredicts the rise in revision size. Forecast errors also rise in turbulent periods but more modestly than the data. Since the model with known volatility only partially replicates this increase, it suggests that additional uncertainty beyond an increase in volatility shifts plays a role.

Figure VI – Known changes in volatility

(a) Frequency of Revisions **(b)** Size of Revisions **(c)** Mean Forecast Errors

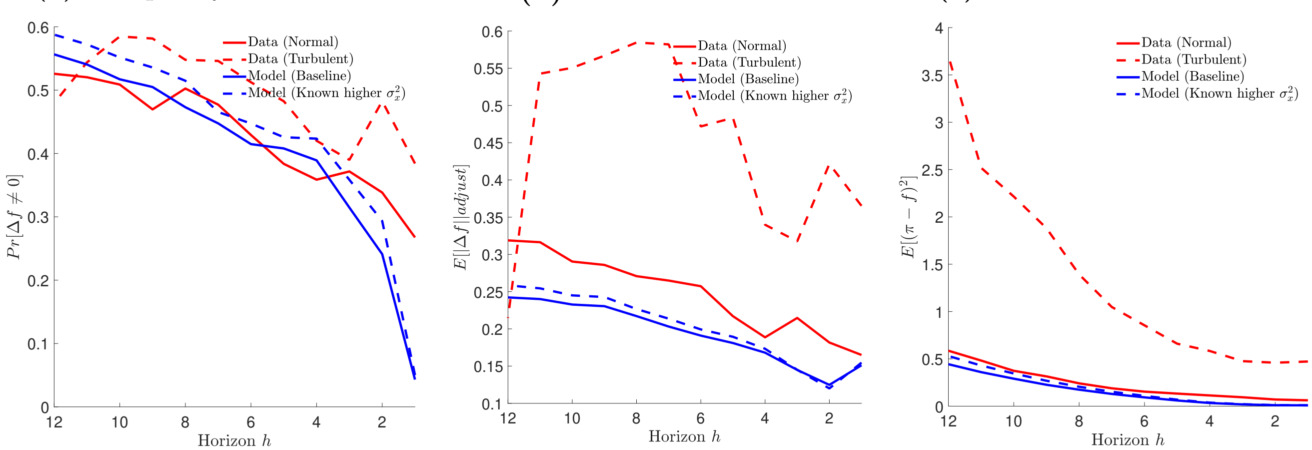

Notes: Red lines are from Bloomberg data. Tranquil = 2010-20. Turbulent = 2008-2009 and 2021-23. Blue lines use model-simulated data.

4.1.2 Unknown regime change

Figure VII introduces an alternative scenario in which forecasters do not anticipate the increase in volatility but instead continue acting as if volatility remains unchanged. The only difference relative to the baseline is that the actual shocks exhibit greater volatility, while agents' policies remain fixed. Unlike Figure VI where agents adjust their revision behavior in response to expected higher volatility, here they follow the same decision rules as in tranquil times. The effects on revisions are similar, but the key difference emerges in forecast errors. When volatility increases unexpectedly, errors rise significantly more than in the case where forecasters knew the volatility shift in advance. This highlights that part of the increased forecast uncertainty in turbulent times arises purely from volatility effects, independently of any option effects or changes in revision policies.

Figure VII – Unknown changes in volatility

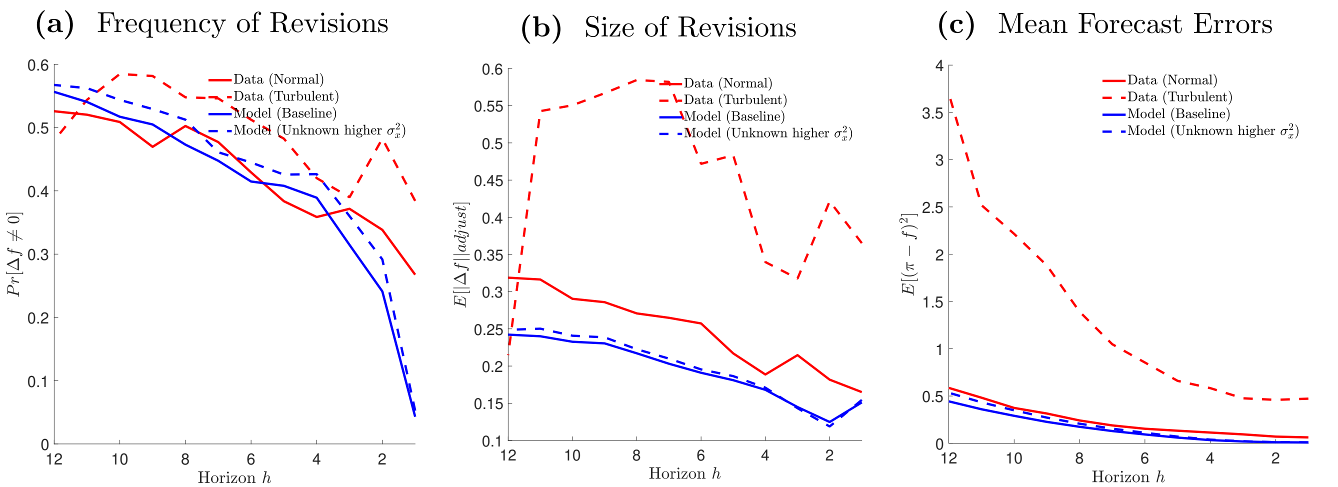

Notes: Red lines are from Bloomberg data. Tranquil = 2010-20. Turbulent = 2008-2009 and 2021-23. Blue lines use model-simulated data.

4.1.3 Learning regime

The third exercise will introduce a learning mechanism where forecasters do not immediately know when inflation volatility changes but must infer it over time. Instead of reacting instantly to shifts in volatility, agents will update their beliefs about σ_x^2 using observed inflation data, applying a Bayesian filtering mechanism (Baley and Veldkamp, 2025). This process introduces gradual adjustments in forecasting behavior as agents refine their volatility estimates with each new observation. As a result, forecast revisions and errors will evolve dynamically, initially exhibiting inertia due to uncertainty about the actual volatility regime before converging as agents gain confidence in their updated beliefs. This framework will allow us to examine how learning frictions influence the responsiveness of expectations to inflation shocks and how long it takes forecasters to recognize and adjust to a new inflation environment.

4.2 Changes in the mean

In the third exercise, we study how changes in the unconditional mean of inflation influence forecast behavior, distinguishing between deflationary and inflationary turbulent periods. As seen in Section 2, forecast errors rise significantly in high-inflation environments, suggesting that uncertainty stems not only from volatility but also from shifts in the long-run level of inflation. This analysis connects to the broader debate on whether inflation shocks are transitory or persistent, a key question for policymakers and forecasters. To examine this, we separate the analysis into Turbulent I (lower inflation) and Turbulent II (higher inflation) and extend the model to incorporate different unconditional means. We consider two scenarios: one in which agents know whether the change is permanent or temporary and another in which they must infer its persistence over time. By comparing these cases, we assess whether the rise in forecast errors during turbulent periods is primarily due to volatility or uncertainty about structural shifts in inflation dynamics.

4.2.1 Known mean change

In the known mean change scenario, forecasters immediately recognize whether the shift in inflation is transitory or permanent and adjust their expectations accordingly. This allows us to test how well the model captures differences in forecast behavior when agents react optimally to structural inflation shifts.

4.2.2 Unknown mean change

In the unknown mean change scenario, forecasters do not initially know whether the change is temporary or permanent and must infer it over time. By observing how forecast errors and revisions evolve, we assess how learning frictions influence expectation formation and whether uncertainty about inflation persistence amplifies forecast inaccuracy.

4.3 Large shocks travel fast?

In this exercise, we analyze the effects of large inflation shocks on forecast revisions, drawing inspiration from inaction models in price-setting under large shocks (Cavallo, Lippi and Miyahara, 2024). The idea is that, beyond changes in volatility, large shocks can push forecasters outside their inaction regions, triggering immediate and widespread revisions. To test this, we will feed the model with actual inflation realizations and examine how well it replicates observed forecast behavior.

Suppose the model correctly captures the mechanisms of lumpy adjustments. In that case, we expect large shocks to induce a rapid response, similar to what is observed in price-setting models where substantial cost shocks lead to immediate price adjustments. By comparing individual forecast and consensus forecasts, this exercise will allow us to assess whether expectation formation exhibits the same non-linearity as price-setting, with large shocks traveling quickly through the forecasting process.

5 Frictions in turbulent times

In this section, we relax the assumption that the frictions governing forecast behavior—fixed costs of revision, strategic concerns, and private signal noise—remain constant at their steady-state values. Instead, we allow these frictions to evolve over time, particularly in response to shifts in the volatility regime. While our ability to precisely estimate these changes is constrained by the limited number of observations in turbulent periods, we nonetheless re-estimate the structural parameters using our SMM procedure, targeting empirical moments specific to high-volatility environments. This allows us to assess how the costs of revision, the strength of strategic complementarities, and the precision of private signals vary across different macroeconomic conditions.

By doing so, we contribute to the literature on expectation formation in volatile periods, which suggests that turbulent times are associated with more dispersed signals and greater uncertainty. If information dispersion increases, private signals may become noisier, leading to greater heterogeneity in expectations. Furthermore, strategic interactions among forecasters may shift: complementarities that drive coordination in tranquil times could weaken, or even turn into substitutability, as agents react more idiosyncratically to new information when uncertainty is high. Our analysis will provide new insights into how expectation frictions adapt in response to economic turbulence and whether the observed changes in forecast behavior can be explained by shifts in these underlying mechanisms.

TBC...

References

Álvarez, F. and Lippi, F. (2014). Price setting with menu cost for multi-product firms. *Econometrica*, **82** (1), 89–135.

Álvarez, F., Lippi, F. and Paciello, L. (2011). Optimal price setting with observation and menu costs. *The Quarterly Journal of Economics*, **126** (4), 1909–1960.

Argente, D. and Yeh, C. (2022). Product life cycle, learning, and nominal shocks. *The Review of Economic Studies*, **89** (6), 2992–3054.

Baley, I. and Blanco, A. (2019). Firm uncertainty cycles and the propagation of nominal shocks. *American Economic Journal: Macroeconomics*, **11** (1), 276–337.

—, Figueiredo, A. and Ulbricht, R. (2022). Mismatch cycles. *Journal of Political Economy*, **130**, 2943–2984.

— and Turen, J. (2024). *Lumpy Forecasts*. CEPR Discussion Paper DP19824, Centre for Economic Policy Research.

— and Veldkamp, L. (2025). *The Data Economy: Tools and Applications*. Princeton University Press, forthcoming.

Borovicková, K. (2016). *Job flows, worker flows and labor market policies*. Tech. rep., New York University.

Cavallo, A., Lippi, F. and Miyahara, K. (2024). Large shocks travel fast. *American Economic Review: Insights*, **6** (2), 123–138.

Evans, G. W. and Honkapohja, S. (1993). Adaptive forecasts, hysteresis, and endogenous fluctuations. *Economic Review-Federal Reserve Bank of San Francisco*, (1), 3.

Giacomini, R., Skreta, V. and Turen, J. (2020). Heterogeneity, inattention, and bayesian updates. *American Economic Journal: Macroeconomics*, **12** (1), 282–309.

Lasry, J.-M. and Lions, P.-L. (2007). Mean field games. *Japanese journal of mathematics*, **2** (1), 229–260.

Marcet, A. and Nicolini, J. P. (2003). Recurrent hyperinflations and learning. *American Economic Review*, **93** (5), 1476–1498.

Midrigan, V. (2011). Menu cost, multiproduct firms, and aggregate fluctuations. *Econometrica*, **79** (4), 1139–1180.

Molavi, P. *et al.* (2019). Macroeconomics with learning and misspecification: A general theory and applications. *Unpublished manuscript*.

Morris, S. and Shin, H. S. (2002). Social value of public information. *American Economic Review*, **92** (5), 1521–1534.

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.

Pfäuti, O. (2024). Inflation—who cares? monetary policy in times of low attention. *Journal of Money, Credit and Banking*, **56** (2), 123–145.

Valchev, R. and Gemmi, L. (2023). *Biased Surveys*. Tech. rep., National Bureau of Economic Research.

Vavra, J. (2014). Inflation dynamics and time-varying volatility. *The Quarterly Journal of Economics*, **129** (1), 215–258.

A Proxy for Forecasters' Beliefs

In this section we derive the expression for the belief $\hat{\pi}_h \equiv \mathbb{E}[\pi|\mathcal{I}_h]$ and its precision.

Optimal Forecast We derive the conditional mean of annual inflation at different horizons h. As discussed we assume that $x_h = c + \phi x_{h+1} + v_h$. The forecasted variable π is approximately equal to the sum of the last twelve realizations of monthly inflation x_h:

$$\pi = \sum_{h=0}^{11} x_h$$

From the AR(1) assumption implies that the unconditional mean of x_h is $\widetilde{\mu} \equiv \frac{c}{1-\phi}$, hence we can rewrite the demeaned process for monthly inflation at each relevant month within the target-year as:

$$\begin{aligned}
x_0 &= \widetilde{\mu} + \phi^{12}(x_{12} - \widetilde{\mu}) + \sum_{j=0}^{11} \phi^j v_j \\
x_1 &= \widetilde{\mu} + \phi^{11}(x_{12} - \widetilde{\mu}) + \sum_{j=0}^{10} \phi^j v_{j+1} \\
\dots \\
x_{11} &= \widetilde{\mu} + \phi(x_{12} - \widetilde{\mu}_i) + v_{11},
\end{aligned}$$

which implies

$$\pi = 12\widetilde{\mu} + \frac{\phi(1-\phi^{12})}{1-\phi}(x_{12} - \widetilde{\mu}) + \sum_{j=0}^{11} \frac{1-\phi^{j+1}}{1-\phi} v_j.$$

At horizon $h \geq 12$ we have:

$$\pi = 12\widetilde{\mu} + \frac{\phi^{h-11}(1-\phi^{12})}{1-\phi}(x_h - \widetilde{\mu}) + \sum_{j=0}^{11} \frac{1-\phi^{j+1}}{1-\phi} v_j + \sum_{j=12}^{h-1} \frac{\phi^{j-11}(1-\phi^{12})}{1-\phi} v_j \tag{23}$$

Within the year, $11 - h$ realizations of x_h are observed. Hence, for $h \leq 11$:

$$\pi = h\widetilde{\mu} + \frac{\phi(1-\phi^h)}{1-\phi}(x_h - \widetilde{\mu}) + \sum_{j=h}^{11} x_j + \sum_{j=0}^{h+1} \frac{1-\phi^{j+1}}{1-\phi} v_j \tag{24}$$

Public signal The signal $z_{i,h}$ is equal to the conditional expectation of equations (23) and (24), respectively, given information up to horizon $h+1$:[5]

$$z_h = 12\widetilde{\mu} + \frac{\phi^{h-10}(1-\phi^{12})}{1-\phi}(x_{h+1} - \widetilde{\mu}_i) \text{ for } h = 18, \ldots, 11 \tag{25}$$

$$= (h+1)\widetilde{\mu}_i + \frac{\phi(1-\phi^{h+1})}{1-\phi}(x_{h+1} - \widetilde{\mu}) + \sum_{j=h+1}^{11} x_j \text{ for } h = 10, \ldots, 1, \tag{26}$$

which correspond to equation (**??**) in the main text.

Signal Precision Forecasters evaluate the precision of their signals through the AR(1) model. Given the signal $z_h = \pi + \varepsilon_h$ and the expressions for π in (23) and z_h in (25), the variance of the forecast error at horizon $h \geq 12$ is derived as follows:

$$\begin{aligned}
\varepsilon_{i,h} &= \pi - z_{i,h} \\
&= \sum_{j=0}^{11} \frac{1-\phi^{j+1}}{1-\phi} v_j + \sum_{j=12}^{h} \frac{\phi^{j-11}(1-\phi^{12})}{1-\phi} v_j \\
\mathbb{E}[\varepsilon_{i,h}^2] &= \sigma_v^2 \sum_{j=0}^{11} \frac{(1-\phi^{j+1})^2}{(1-\phi)^2} + \sigma_v^2 \sum_{j=12}^{h} \frac{\phi^{2j-22}(1-\phi^{12})^2}{(1-\phi)^2},
\end{aligned}$$

where we rely on the assumption v_h is i.i.d and drawn from a Gaussian distribution with common variance, $v_h \sim N(0, \sigma_v^2)$ and the timing of the public release of information. Hence, the conditional variance is constant across forecasters. The first expression on the right hand side of the previous equation is:

$$\begin{aligned}
\sigma_v^2 \sum_{j=0}^{11} \frac{(1-\phi^{j+1})^2}{(1-\phi)^2} &= \frac{\sigma_v^2}{(1-\phi)^2} \sum_{j=0}^{11} (1-\phi^{j+1})^2 \\
&= \frac{\sigma_v^2}{(1-\phi)^2} [(1-\phi)^2 + (1-\phi^2)^2 + \cdots + (1-\phi^{12})^2] \\
&= \frac{\sigma_v^2}{(1-\phi)^2} [12 - 2(\phi + \phi^2 + \cdots + \phi^{12}) + (\phi^2 + \phi^4 + \cdots + \phi^{24})] \\
&= \frac{\sigma_v^2}{(1-\phi)^2} \left[12 - \frac{2\phi(1-\phi^{12})}{1-\phi} + \frac{\phi^2(1-\phi^{24})}{1-\phi^2} \right]
\end{aligned}$$

[5] Although the horizon for the prediction is h we assume the information of monthly inflation is available only until $h+1$. This is assume to account for the monthly delay in the publication of official statistics.

The second expression is:

$$
\begin{aligned}
\sigma_v^2 \sum_{j=12}^{h} \frac{\phi^{2j-22}(1-\phi^{12})^2}{(1-\phi)^2} &= \sigma_v^2 \frac{(1-\phi^{12})^2}{(1-\phi)^2} \sum_{j=12}^{h} \phi^{2j-22} \\
&= \sigma_v^2 \frac{(1-\phi^{12})^2}{(1-\phi)^2} (\phi^2 + \phi^4 + \cdots + \phi^{2h-22}) \\
&= \sigma_v^2 \frac{(1-\phi^{12})^2}{(1-\phi)^2} \frac{\phi^2(1-\phi^{2h-22})}{1-\phi^2} \\
&= \sigma_v^2 \frac{\phi^2(1-\phi^{12})^2(1-\phi^{2h-22})}{(1-\phi)^3(1+\phi)}
\end{aligned}
$$

Summing the two expressions we obtain the expression for b_h^{-1} for $h \geq 11$, while relying on the same derivation as in the first equation, we obtain the expression for $h \leq 10$.

B Rolling Estimates AR(1) model

Let the monthly inflation rate x_h follow an AR(1) process:

$$
x_h = c + \phi x_{h+1} + v_h, \qquad v_h \sim \mathcal{N}(0, \sigma_v^2), \tag{27}
$$

where c is a constant, ϕ is the persistence parameter, and v_h is an *iid* normally distributed noise with volatility σ_v^2. We estimate the three parameters (c, ϕ, σ_v) using the monthly inflation rate from the CPI. Figure VIII plots the resulting estimates and 95% confidence intervals.

Figure VIII – Rolling Estimates AR(1) parameters

(a) **(b)** **(c)**

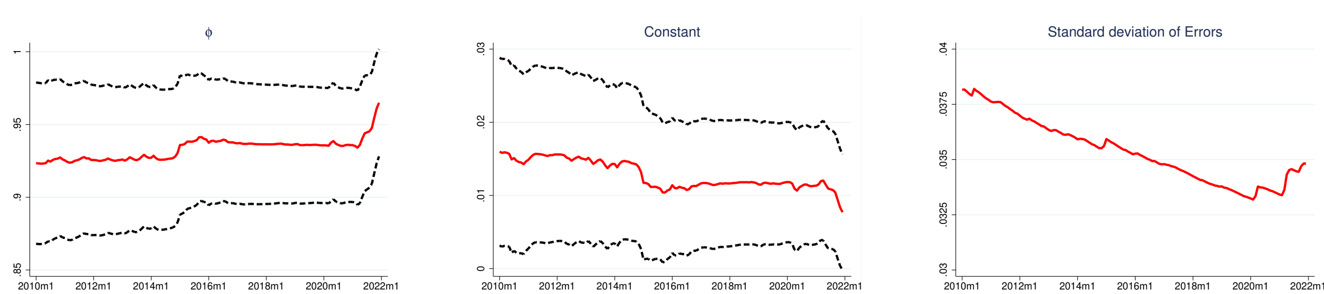

Notes: Bloomberg data.

C Cross-sectional Statistics for All Years

Figure IX – Frequency and Size of Adjustment, by Horizon

(a) Frequency

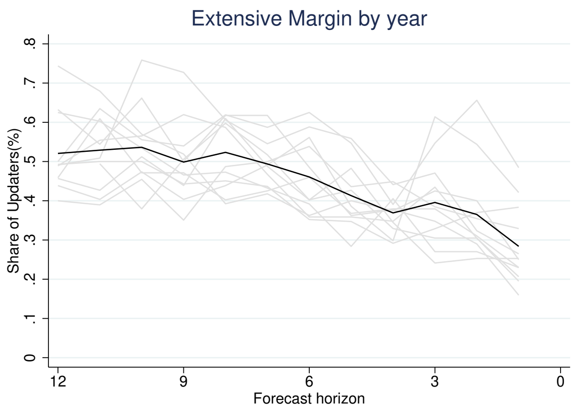

(b) Size

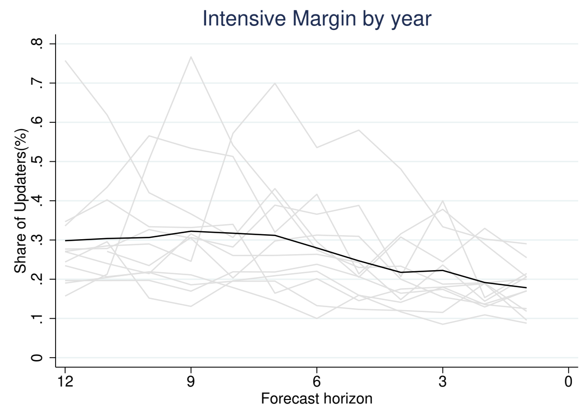

(c) Forecast errors

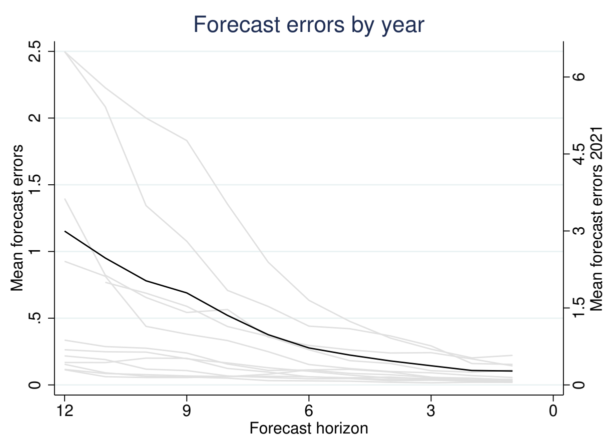

Notes: Bloomberg data.