Lumpy Forecasts*

Isaac Baley[†]
UPF, CREI, BSE and CEPR

Javier Turén[‡]
Pontificia Universidad Católica de Chile

July 5, 2024

Work in progress, for the latest version: Click here

Abstract

We document that professional forecasters adjust inflation forecasts in a lumpy way—forecasts are changed infrequently, and when adjusted, they are revised by a significant amount. As the forecasting horizon shrinks, the frequency of revisions, the size of revisions, and forecast errors decrease. Using a fixed-event forecasting framework, we assess the role of the consensus forecast and private information in shaping forecast revisions, both at the extensive and the intensive margins. A model of Bayesian belief formation with forecast revision costs and strategic concerns (i) delivers lumpy forecasts consistent with the survey evidence, (ii) rationalizes forecast efficiency tests without introducing behavioral biases, and (iii) generates state-dependent responses to inflation volatility. We discuss implications for the transmission of monetary policy and forecast survey design.

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

Keywords: fixed-event forecasting, consensus, adjustment costs, forecast stability, strategic concerns, Bayesian learning

*We largely benefited from conversations with João Victor Issler (discussant), Philippe Andrade, George-Marios Angeletos, Saleem Bahaj, Benjamin Born, Ricardo Caballero, Giacomo Candian, Luca Gemmi, Paul Hubert, Alex Kohlhas, Francesco Lippi, Kristoffer Nimark, Ricardo Reis, and many seminar participants. 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, Camila Gutiérrez, and Valentina Hernández 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

Expectations and belief formation are central to macroeconomics. Both theoretical and empirical work has questioned the Full Information Rational Expectations (FIRE) paradigm in favor of other rigidities.[1] For instance, Coibion and Gorodnichenko (2015) argues that professional forecasters deviate from the traditional rational expectations full-information hypothesis as they form expectations facing limitations on the rate by which they observe and acquire information. A common assumption at the core of the literature is that forecasts accurately reflect agents' beliefs. We challenge this assumption, ask to what extent forecasts may differ from agents' beliefs, and explore why it is the case.

First, using a relatively understudied survey of professional forecasters and a fixed-event forecasting framework (Nordhaus, 1987; Patton and Timmermann, 2012, 2011), we document that end-of-year inflation forecasts are *lumpy*: Forecasts are characterized by periods of inaction followed by significant revisions. Forecast lumpiness is present despite relevant information becoming available during the forecasting period, which should, in principle, change agents' beliefs. Thus, forecasts are a rather *imperfect measure* of agents' information sets. This evidence casts doubts on how forecasts proxy for individual beliefs and has important implications for monetary and other policy tools that rely on these forecasts.

Second, we build a forecasting model with three features to rationalize the evidence. First, beliefs are formed using past information and noisy signals using Bayes' Law. Second, forecasters prefer forecast stability, which we operationalize through fixed revision costs. Third, forecasters engage in strategic behavior in that the distance to the consensus (the average) forecast matters for their payoffs. The model allows us to back out the beliefs (what agents have in their minds) and separate them from forecasts (what they would report). We estimate the model parameters to match cross-sectional statistics from the survey data. The model delivers lumpy forecasts consistent with the survey evidence.

Third, we use the calibrated model to shed light on the relevance of lumpy forecasts in different dimensions: (i) we assess the role of forecaster ex-ante heterogeneity in fixed costs and strategic concerns, (ii) we rationalize forecast efficiency tests without introducing behavioral biases, and (iii) we generate the observed response to increases in inflation volatility.

Next, we explain in more detail the empirical and theoretical contributions.

Facts on lumpy forecasts The first part of the paper documents new facts on lumpy forecast revisions. Using the Economic Forecasts ECFC survey conducted by Bloomberg for the US for 2010-2020, we document the evolution of monthly inflation forecasts about end-of-year inflation.

[1]For instance, sticky information model (Mankiw and Reis, 2002), rational inattention (Sims, 2003) and (Mackowiak and Wiederholt, 2009), higher-order uncertainty (Morris and Shin, 2002) and (Woodford, 2003), level-k thinking (García-Schmidt and Woodford, 2019), over- and under-reaction to private information (Bordalo, Gennaioli, Ma and Shleifer, 2018), to name a few.

We show that the frequency of revisions, the variance of revisions, and forecast errors *fall* as the date the forecasted variable is released approaches. Since we focus on predictions for the same variable (annual inflation) at different horizons, these patterns are puzzling as they are precisely at shorter horizons when the most significant amount of relevant information is available. We should, in principle, expect more revisions. We also study agents' learning rates, adjustment hazards, and other revision patterns along the forecasting horizon, which we label the "term structure" of forecast revisions and errors.

Decreasing hazard rate is a common feature of Bayesian learning models with fixed adjustment costs in actions (Baley and Veldkamp, 2025). The quantitative model in section 3 builds on these elements to generate a decreasing hazard of forecast revisions. Additionally, we exploit the observed hazard's slope to calibrate some of the model's parameters.

Related to the triggers of revisions, we show that the consensus gap—the distance from an individual forecast to the average forecast (the consensus) triggers revisions. This effect is significant both at the extensive and the intensive margin. The probability of revision increases with the gap, and conditional on a revision, new forecasts are closer to the consensus. The results suggest that strategic complementarities are at play, in which forecasts depend on the forecasts of others, reflecting, on the one hand, correlated information and, on the other hand, reputational costs.

Forecasting model with fixed revisions costs In the second part of the paper, we build a novel forecasting model with fixed revision and strategic concerns. We interpret the fixed revision cost as a stand for preference for forecast stability. It may reflect the implicit costs of disclosing private information, which is particularly important for consulting companies and other financial institutions that provide forecasting services to their clients. Alternatively, it may reflect forecast congruence; forecasters make decisions based on their forecasts (e.g., investment) and thus do not change them to harm their valuations.

Revision costs reduce the frequency and size of revisions. Using a simulated method of moments (SMM), we estimate the model parameters to match average cross-sectional statistics. We emphasize the importance of matching the shape of the hazard rate. The estimation suggests the presence of significant revision costs, strategic complementarity, and idiosyncratic noise. The model delivers horizon-dependent frequency and variance of revisions and forecast errors that match the empirical term structure of these moments. Alternative calibrations that shut down either motive cannot jointly explain the term structure of revisions and forecast errors.

Three applications We use the calibrated model to illuminate the relevance of lumpy forecasts in three ways.

First, we use the model to investigate the strength of strategic concerns and the preference for forecast stability across the four types of forecasters in our data: (i) banks, (ii) financial institutions, (iii) consulting companies, and (iv) universities and research centers. For this purpose, we recalibrate the model to match type-specific moments. The most significant differences in these moments occur between consulting companies and universities. For instance, relative to universities, consulting companies adjust 1.4 times more frequently than universities and do revisions that are 1.3 times more dispersed. Through the model's lens, these moments imply that universities face higher revision costs and stronger concerns for forecast stability. Thus, our results suggest that forecasters' ex-ante heterogeneity is an important dimension when working with this type of survey. These results complement studies focusing on model heterogeneity as a driver behind professional forecasters' reports being different predictions for the same macroeconomic outcome (Giacomini, Skreta and Turen, 2020).

Second, we assess forecast efficiency regression tests in the spirit of Bordalo, Gennaioli, Ma and Shleifer (2020). By regressing forecast errors on forecast revisions, the test investigates the predictability of forecast errors at the individual level, and thus, it is a test for rational expectations. Following Broer and Kohlhas (2022) and Valchev and Gemmi (2023), we extend the baseline test to incorporate the consensus forecast as a public information source. Consistent with the empirical regressions, we find a non-significant bias, a negative and significant coefficient on forecast revisions (interpreted as an over-reaction to private information), and a positive and significant coefficient on the distance to the consensus (interpreted as an under-reaction to public information). Importantly, our results are generated in a model with Bayesian agents without needing behavioral biases (e.g., extrapolating expectations) as in Bordalo *et al.* (2020).

Third, we investigate the forecasts' response to increases in underlying inflation volatility. We motivate this exercise by the observation that during turbulent years—the years of the Great Recessions (2008-2009) and the COVID-19 pandemic (2020-2021)—monthly inflation volatility spiked by 40% relative to normal years (2010-2019). In the same period, forecast revision became more frequent and dispersed, sharply increasing forecast errors. In the model, we implement this change by increasing the volatility of the underlying state to the same amount as in the data and simulating the economy under two scenarios. In the first scenario, we give forecasters the information that volatility has changed (disclosed volatility shock). In the second scenario, we do not give this information and thus keep the policy functions as in the benchmark calibration (undisclosed volatility shock). In both cases, the cross-sectional moments of forecast revisions and errors increase as in the data. However, the forecasts' response under the undisclosed volatility shock better matches the large increase in the empirical moments, especially the forecast errors.

Contributions We contribute to the literature that uses survey data to elicit expectations (see Bachmann, Topa and van der Klaauw (2022) for a comprehensive review).

On the empirical side, our facts on forecast lumpiness complement evidence by Andrade and Le Bihan (2013) showing that expectations from professional forecasters are sticky. More recently, using the well-known Ifo Survey of firms in Germany, Born, Enders, Müller and Niemann (2022) confirms that firms' expectations are also sticky as they are adjusted only infrequently.

On the theoretical side, we contribute to the vast literature on fixed-event forecasting (Nordhaus, 1987; Patton and Timmermann, 2012, 2011; Capistrán and López-Moctezuma, 2014), by developing and estimating a model with fixed revision costs and strategic concerns. Regarding lumpiness, Mankiw and Reis (2002) builds a model where agents update their expectations infrequently because they collect and update information relatively infrequently. In contrast, our model assumes that beliefs are continuously updated with new information, but forecasts are revised infrequently because of adjustment costs. Building on the "sticky information" theory of Mankiw and Reis (2002), Bec, Boucekkine and Jardet (2023) develop a forecasting model applying the price-setting model of Alvarez, Lippi and Paciello (2011); Álvarez, Lippi and Paciello (2018) that features both fixed observation and revision costs.

We also contribute to the literature on heterogeneous agents with aggregate shocks. Solving our model's rational expectations equilibrium is infeasible. Because the whole distribution of forecasts matters in determining current and future consensus, agents must forecast this distribution. Instead, we use an equilibrium concept called *restricted perceptions equilibrium* (RPE) that deviates from rational expectations. This equilibrium concept has been used in signal extraction models like ours, in which agents observe a noisy signal about an underlying state variable, by Evans and Honkapohja (1993), Marcet and Nicolini (2003), and Molavi (2022). footnoteWe borrow the equilibrium term from Evans and Honkapohja (2001).

2 The Anatomy of Inflation Forecasts

This section begins by describing data sources and the fixed-event forecasting framework. Then, we document how forecast revisions and errors evolve over the forecasting horizon.

2.1 Inflation

We construct annual inflation in the United States using the Consumer Price Index (CPI). For any year t, we let cpi_j be the CPI measured j months before the end of year and we let $\overline{cpi}_t = \frac{1}{12}\sum_{j=0}^{11} cpi_j$ be the average CPI in year t. The annual inflation rate π_t in year t is calculated as

$$\pi_t = \log(\overline{cpi}_t) - \log(\overline{cpi}_{t-1}). \tag{1}$$

Following Giacomini, Skreta and Turen (2020), we approximate the annual inflation rate using the sum of year-on-year monthly inflation rates, $x_{m,t}$, as follows:[2]

$$\pi_t \cong \sum_{m=1}^{12} x_{m,t}, \quad \text{with} \quad x_{m,t} = \frac{\log(cpi_{m,t}) - \log(cpi_{m-12,t})}{12}, \quad \forall m = 1, \ldots, 12 \tag{2}$$

[2]See Appendix A.1 for the derivation and conditions under which this approximation holds.

Figure I plots the series of annual inflation π_t (solid black line) and year-on-year monthly inflation $x_{m,t}$ (in red dots) for the sample period. Annual inflation varies from -0.3% in 2010 to 4.7% in 2022.

Figure I – US Annual Inflation and Year-On-Year Monthly Inflation

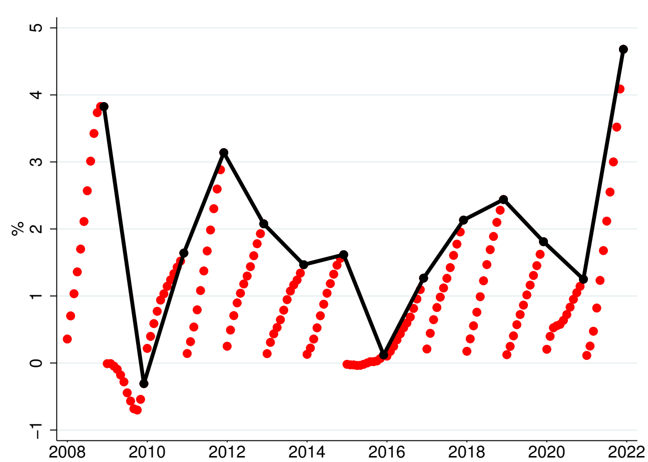

Notes: CPI inflation rates in the US for 2008-2022. Annual inflation rates π_t are shown in a solid black line; year-on-year monthly inflation rates $x_{m,t}$ is in red dots.

2.2 Survey 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 their institutional background. There are four main types of forecasters: banks, financial institutions, consulting firms, and universities and research centers.[3] Besides its similarities, one of the most appealing features of this particular survey is that the most recent forecasts of any other forecaster, the date when each prediction was last updated, and the consensus forecast (the mean forecast) are visible to users of the Bloomberg terminal in *real-time*.

Sample We examine monthly fixed-event forecasts of annual US inflation. Our sample covers the years 2008 to 2022. In the main analysis, we focus on low-volatility years 2010-2019. Section 6 analyses the turbulent years of the Great Recession, 2008-2009, and the COVID-19 pandemic, 2020-21, in which the inflation process was more volatile. For each year, we consider survey participants who forecast inflation for all 12 months before the final figure (end-of-year inflation) is officially published. We remove forecasters who fail to provide at least one annual inflation

[3]See Giacomini, Skreta and Turen (2020) for further details related to the comparison between the Bloomberg survey and other existing surveys.

revision. This criteria leaves approximately 100 forecasters per year. The panel dataset contains the history of forecast updates for all forecasters over a12-month horizon each year.[4]

Incentives Participants in the survey are not anonymous. In fact, through the terminal, it is possible to trace out the entire time series of predictions for any institution across years. As discussed by Croushore (1997), we anticipate that forecasters face reputational concerns given this feature. One typical concern of these financial analysts' surveys is whether the reported forecasts drive the posterior trading behavior of forecasters. Using predictions also collected from Bloomberg surveys, Bahaj, Czech, Ding and Reis (2023) provides empirical evidence that supports this claim. Building on this evidence, we argue that the ultimate investment decisions of these analysts are indeed linked with their reported predictions and, therefore, with their incentives to provide accurate forecasts.

2.3 Fixed-event forecasting

Inflation forecasts Within each year, we denote with f_h^i the inflation forecast in the percentage 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 count 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).

The fixed event is the end-of-year inflation π. All forecasts f_h^i refer to the fixed event. Monthly inflation rates x_h are usually published between the month's second and third week. Figure II illustrates the fixed-event forecasting framework. Each month, participants can provide a prediction f_h^i for the end-of-year inflation while entertaining the possibility of keeping the prediction constant through time.

$$f_h^i = \underbrace{\sum_{j=h+1}^{12} x_j}_{\text{past realizations}} + \underbrace{\mathcal{P}_h^i}_{\text{projection}} + \qquad h = 12, \ldots, 1$$

The forecast consists of two terms: a "sunk" component given by the sum of past realizations $\sum_{j=h+1}^{12} x_j$ and a projection component $\mathcal{P}_h^i$ that reflects the "true" forecasting activity.

[4]Although we have information on the precise dates when a forecast was revised, we analyze a monthly frequency as there are only very few weekly updates. In particular, we use the forecast available on the terminal on the last day of the month to construct our monthly panel data.

Figure II – Fixed-event forecasting

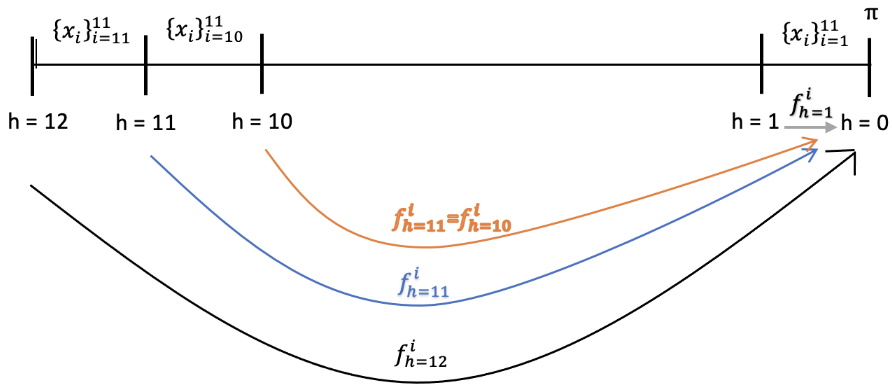

Notes: The figure illustrates how fixed-event forecasts work. The fixed event is the end-of-year inflation π. All forecasts f_h^i refer to the fixed event. Bloomberg allows for multiple revisions within any month, so there is no restriction on the amount of revisions that a participant can do.

2.4 Forecasts revisions and errors

Forecast 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:

$$\Delta f_h^i \equiv f_h^i - f_{h+1}^i. \tag{4}$$

Since forecasts are in percentages, revisions are measured in percentage points. Table I reports summary statistics of forecast revisions averaged across years, forecasters, and horizons. The average revision is close to zero, $\mathbb{E}[\Delta f] = -0.013$, which suggests a symmetric environment in which positive and negative revisions, on average, cancel out. The average revision size (in absolute value) equals $\mathbb{E}[abs(\Delta f)] = 0.110$ and $\mathbb{E}[abs(\Delta f)|\Delta f \neq 0] = 0.247$ depending if we condition on all revisions or we include non-zero revisions only. There are approximately five forecast revisions in a given year, which means forecasts are inactive for 1.6 months on average. The adjustment frequency is 0.427, and downward revisions (0.231) are slightly more likely than upward revisions (0.196). The spike rate, which counts the proportion of significant adjustments (above 20%), is close to 3%.

Forecast errors At any given year, we let e_h^i be the forecast error of individual i at horizon h, defined as the difference between the actual end-of-year inflation π and the reported forecast f_h^i.

$$e_h^i \equiv \pi - f_h^i. \tag{5}$$

Table II provides summary statistics on individual and aggregate forecast errors, averaged across years and horizons. Individuals make small errors on average $\mathbb{E}[e] = -0.055$ and tend to

Table I – Summary Statistics of Forecast Revisions

Average revision	$\mathbb{E}[\Delta f]$	-0.013
Size non-zero revisions	$\mathbb{E}[abs(\Delta f)\|\Delta f \neq 0]$	0.247
Variance non-zero revisions	$\mathbb{V}ar[\Delta f\|\Delta f \neq 0]$	0.110
Avg. number of revisions	$count[\Delta f \neq 0]$	5.059
Months of inaction	$\mathbb{E}[\tau]$	1.594
Adjustment frequency	$\Pr[\Delta f \neq 0]$	0.427
Upward	$\Pr[\Delta f > 0]$	0.196
Downward	$\Pr[\Delta f < 0]$	0.231
Spike rate	$\Pr[abs(\Delta f) > 0.2]$	0.028
Serial correlation (all Δf)	$corr[\Delta f, \Delta f_{-1}]$	-0.043
Serial correlation (non-zero Δf)	$corr[\Delta f, \Delta f_{-1}]$	-0.107
Observations	N	9,256

Notes: Bloomberg data for 2010-2019. Cross-sectional statistics are averaged across years and horizons.

overpredict inflation and thus a negative error e_h^i.

Table II – Summary Statistics of Forecast Errors

Average error	$\mathbb{E}[e]$	-0.055
Mean squared error	$\mathbb{E}[e^2]$	0.235
Serial correlation	$corr[e, e_{-1}]$	0.877
Observations	N	9,256

Notes: Bloomberg data for 2010-2019. Between-year averages of within-year statistics.

2.5 Term structure of revisions and errors

Next, we examine the "term structure" of forecast revisions and errors—how they evolve along the forecasting horizon h. Figure IIIa shows that the magnitude of revisions becomes smaller as the horizon h shrinks. Figure IIIb shows the term structure of forecast errors across horizons. The average squared forecast errors decrease with the horizon. As expected, as the fixed event (end of the year) approaches, more information is accumulated, making the prediction more precise. Despite the monotonic decrease, the forecast error does not converge to zero, even at $h = 1$. This is a tell-tale sign that forecast accuracy is not the only driving force behind forecasters' activities.

Random walk benchmark We compare the Bloomberg forecasts with a "naive" random walk benchmark to isolate mechanical drivers of the term structure of revisions and errors. In this case, the projection is given by $\mathcal{P}_h = hx_{h+1}$. The random-walk projection implies forecast revisions

Figure III – Term Structure of Forecast Revisions and Errors

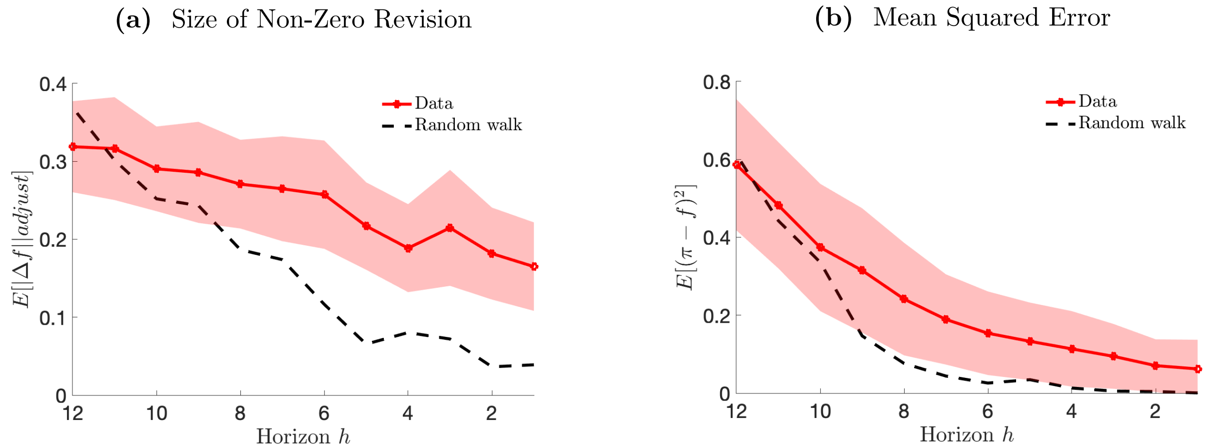

Notes: Bloomberg data for normal years. Panel (a) plots the absolute value of non-zero revisions $\mathbb{E}[|\Delta f|adjust]$. Panel (b) plots the mean squared forecast error $\mathbb{E}[(\pi - f_h^i)^2]$.

Δf_h^{rw} and forecast errors e_h^{rw} the evolve with the horizon according to

$$\Delta f_h^{rw} = (h+1)\Delta x_{h+1}, \tag{6}$$

$$e_h^{rw} = \sum_{j=1}^{h} x_j - h x_{h+1}. \tag{7}$$

The random-walk case is shown as a dashed black line. We see that the size of revisions and forecast errors drop faster under the random walk than the Bloomberg forecasts. This behavior is a tell-tale that forecasts may reflect motives other than accuracy. Next, we explore two potential explanations: forecast lumpiness and strategic concerns.

2.6 Forecasts are lumpy

Figure IVa shows the unconditional probability of updating a forecast across the horizon. Forecasts are updated infrequently. On average, around 38% of forecasters choose to update their predictions throughout the year. The share of updaters also drops as the date of publication of the final inflation figure approaches. The increasing inaction is puzzling as relevant information arrives monthly, which could be used to improve the accuracy of the prediction further.

To see the forecast lumpiness from a different but related angle, we consider the hazard rate of revisions. The hazard rate is a dynamic cross-sectional moment that helps study learning, assess learning speeds, and discriminate across models. It equals the probability of a revision conditional on the forecast's "age", that is, conditional on the time elapsed since the last revision: $h(age) = \Pr[\Delta f \neq 0|age]$. Figure IVb plots the estimated hazard that controls for observed heterogeneity, conditioning on forecaster and year-fixed effects. The hazard is downward slopping: The probability of adjusting a newly set forecast is 0.5; it drops below 0.3 for six-month-old

Figure IV – Term Structure of Forecast Revisions

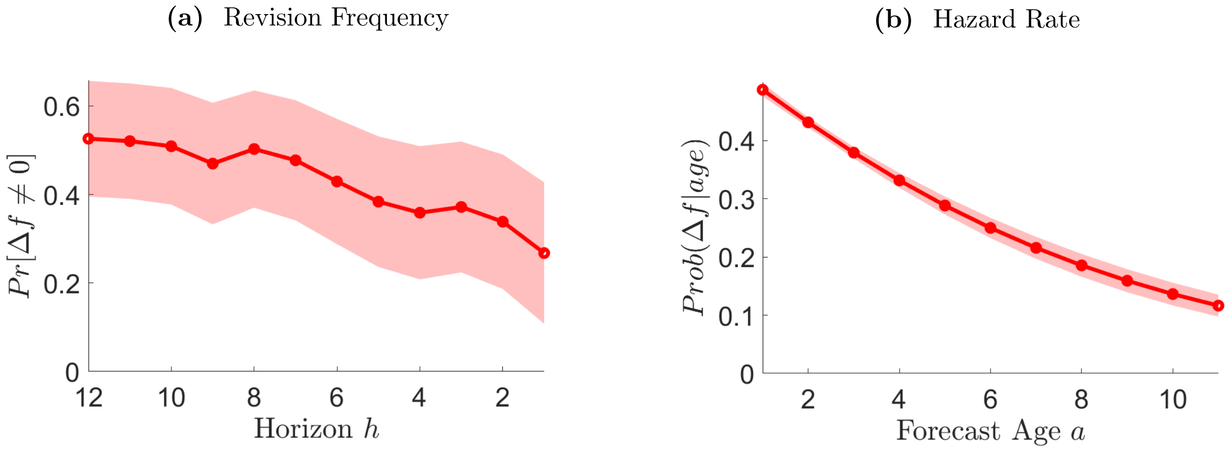

Notes: Bloomberg data for normal years. The left panel shows the frequency of non-zero revisions $\Pr[\Delta f \neq 0]$. The right panel shows the hazard rate of forecast revision $h(age)$.

forecasts and reaches 0.1 for eleven-month-old forecasts.[5]

2.7 Gap to consensus triggers revisions

Next, we explore the role of the consensus forecast in triggering revisions. At each horizon, the consensus forecast is the average forecast across participants:

$$F_h \equiv \frac{1}{N}\sum_{i=1}^{N} f_h^i. \tag{8}$$

Forecasters observe the consensus in real-time through the terminal. Let c_h^i be the gap between individual i's forecast at horizon $h+1$ and the consensus forecast at horizon h:

$$c_h^i \equiv f_{h+1}^i - F_h. \tag{9}$$

We examine how the consensus gap c_h^i affects the probability of updating a prediction, that is, the extensive margin.

Extensive margin We adopt the empirical strategy used in Karadi, Schoenle and Wursten (2021) in the context of firms' price-setting decisions to our setup. First, we divided the gap interval into 15 equally sized bins and computed the frequency of revising a forecast for each bin. The two extreme bins allow the consensus gap to be lower or greater than 1.3%. Second, we run a linear probability model for the probability of a forecaster doing either positive or

[5]Appendix A.4 shows the adjustment hazard conditional on the number of revisions. The age dependence of forecast updating (i.e., the slope of the hazard rate) changes with the number of revisions.

Figure V – Consensus Triggers Revisions

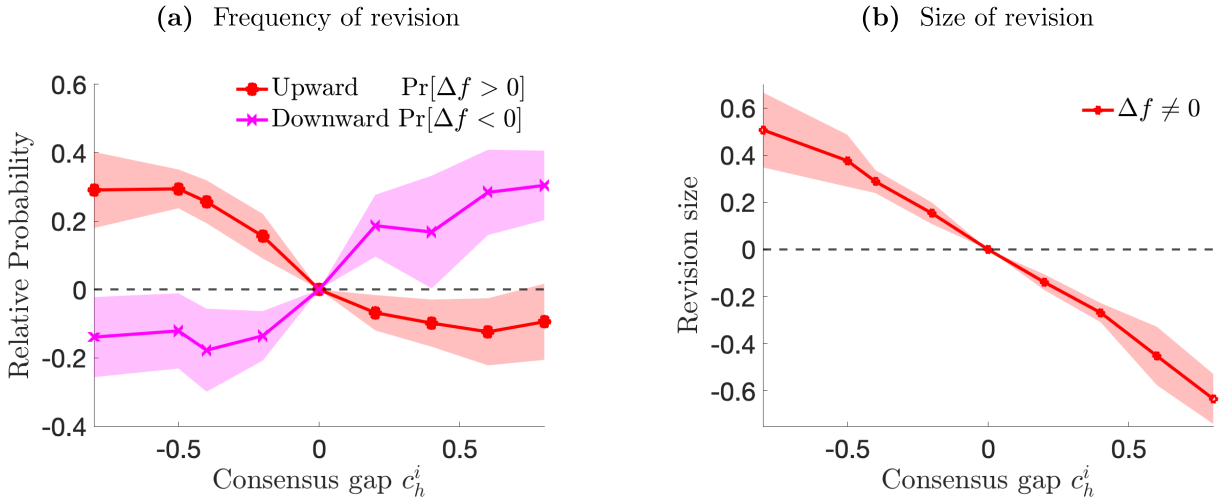

Notes: Bloomberg data.

negative revisions (we run two separate specifications depending on the sign), where we include bin dummies as independent variables. In this case, the consensus gap bin $-0.1\% < c_h^i < 0.1\%$ acts as the reference group. In the specification, we also include forecaster (α_i) and year-horizon ($\alpha_{t,h}$) fixed effects:

$$\Pr[\Delta f_{t,h}^i \neq 0] \;=\; \beta_0 \;+\; \sum_{j=1}^{B} \beta_j \mathbb{1}(c_{t,h}^i(j)) \;+\; \alpha_i \;+\; \alpha_{t,h} \;+\; \epsilon_{t,h}^i \tag{10}$$

Where $\mathbb{1}(c_h^i(j))$ corresponds to the dummies for each of the $B = 15$ bins. As noticed, by estimating the β_j's coefficients, we pin down any possible non-linearities in the relationship between the extensive margin of revisions and c_h^i. Moreover, we run separate regressions for upward and downward forecast revisions to account for potential asymmetries in the dependence on the gap. Figure Va plots the estimated coefficients associated with each dummy, showing the effect of the consensus gap on the probability of positive and negative revisions, respectively, relative to the omitted category.

Two interesting features arise. First, as the relative distance between the forecasts and either gap increases, the probability of a revision increases as well; however, the likelihood of revising upward or downward depends on the sign of the gap. When gaps are above zero, the probability of doing a positive revision ($f_h^i > f_{h+1}^i$) drops while the probability of revising downwards ($f_h^i < f_{h+1}^i$) significantly increases. Likewise, when gaps are negative, the probability of revising upward substantially increases, and the probability of revising downward decreases. Second, the extensive margin reaction appears asymmetric; the updating probability reacts differently depending on whether the forecast is below or above the focal point. Overall, the evidence suggests that distance to the consensus is relevant as it triggers forecast revisions.

Conditioning on agent revisions, we now study the determinants behind the magnitude of revisions as a function of the consensus gap. Figure Vb plots the average revision as a function of the consensus gap c_h^i. As with the extensive margin, we plot residual revisions that control for a battery of fixed effects. All figures show that positive deviations call for negative revisions and the contrary for negative deviations. The strong negative correlation implies that larger deviations call for larger revisions. When either gap is positive (negative), the agent partially revises downwards (upwards) to close this gap.

2.8 Taking stock

To summarize, forecasts are lumpy: they exhibit significant periods of inaction that are followed by large adjustments. The frequency and magnitude of revisions fall with the forecasting horizon. Forecast errors also decrease with the horizon but do not converge to zero. The distance to consensus forecasts (the average predictions of participants) matters for the extensive margin as it triggers forecast revisions.

3 A model of lumpy forecasts

This section builds a fixed-event Bayesian forecasting model with frequent information revelation, strategic concerns, and fixed forecast revision costs.

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:

$$\min_{\{f_h^i\}_{h=12}^1} \mathbb{E}\left[\sum_{h=12}^{1} \underbrace{(f_h^i - \pi)^2}_{\text{accuracy}} + r \underbrace{(f_h^i - F_h)^2}_{\text{strategic}} + \underbrace{\kappa \mathbb{1}_{\{f_h^i \neq f_{h+1}^i\}}}_{\text{stability}} \middle| \mathcal{I}_0^i\right]. \tag{11}$$

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*.

Inflation process Forecasters believe monthly inflation follows an autoregressive process:

$$x_h \;=\; c_x \;+\; \phi_x x_{h+1} \;+\; \varepsilon_h^x, \qquad \varepsilon_h^x \sim \mathcal{N}(0, \sigma_x^2), \tag{12}$$

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. The parameters c_x, ϕ_x and σ_x^2 are common knowledge.

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. The AR(1) assumption for the inflation process implies a public signal about the current *monthly* inflation. Using x_{h+1}, we construct the following common prediction:

$$x_h^{AR} \;\equiv\; \mathbb{E}[x_h | x_{h+1}] \;=\; c_x \;+\; \phi_x x_{h+1}. \tag{13}$$

The variance of the public signal is σ_x^2.

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):

$$\widetilde{x}_h^i \;=\; x_h \;+\; \zeta_h^i, \qquad \zeta_h^i \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_\zeta^2). \tag{14}$$

The idiosyncratic signal noise σ_ζ^2 reflects the heterogeneity in beliefs or models across agents.[6]

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.[7] Therefore, the individual information set $\mathcal{I}_h^i$ at the time of choosing the forecast is

$$\mathcal{I}_h^i \;=\; \widetilde{x}_h^i \;\cup\; \mathcal{I}_h \;=\; \widetilde{x}_h^i \cup \; \{x_{h+1}, x_{h+2}, \ldots, F_{h+1}, F_{h+2}, \ldots\}. \tag{15}$$

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.[8]

[6] We do not explicitly include public (correlated) noise in this signal because the AR(1) signal plays this role. See Valchev and Gemmi (2023) for a model explicitly introducing correlated noise.

[7] 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.

[8] Compare here to the lumpy observation model by Bec, Boucekkine and Jardet (2023).

Figure VI – Timeline of events

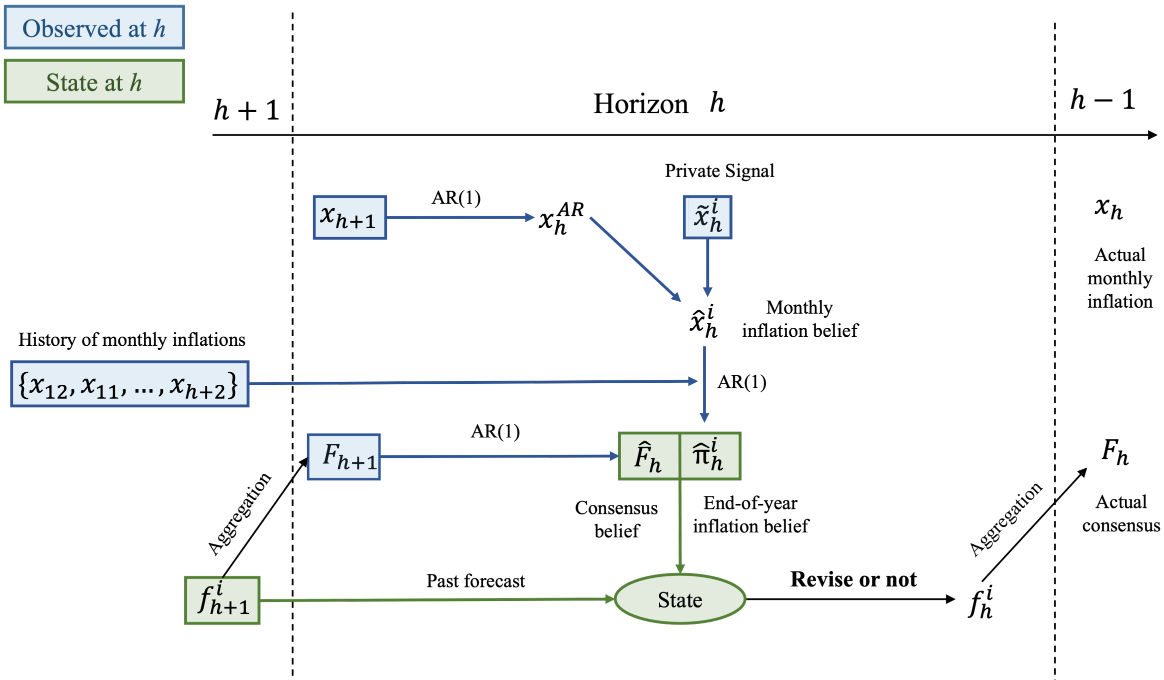

Notes: The figure illustrates the timeline of information revelation, belief formation, and forecast revisions for three contiguous horizons $h+1, h, h-1$.

3.2 Belief formation

Proposition 1 writes the sequential problem in (11) as a function of inflation and consensus beliefs, using the law of iterated expectations and conditioning payoffs on horizon-specific information sets. All proofs are in the Appendix.

Proposition 1. *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:*

$$\min_{\{f_h^i\}_{h=12}^1} \sum_{h=12}^{1} \mathbb{E}\left[\Sigma_h \;+\; \mathbb{E}[(f_h^i - \hat{\pi}_h^i)^2|\mathcal{I}_h^i] \;+\; r\mathbb{E}[(f_h^i - \hat{F}_h)^2|\mathcal{I}_h^i] \;+\; \kappa \mathbb{1}_{\{f_h^i \neq f_{h+1}^i\}} \middle| \mathcal{I}_0^i\right]. \tag{16}$$

where $\Sigma_h \equiv \Sigma_h^\pi + r\Sigma^F$ *is a weighted sum of inflation and consensus uncertainty .*

Next, we characterize individual beliefs about end-of-year inflation $\hat{\pi}_h^i$ and the consensus $\hat{F}_h$. To guide the characterization, Figure VI shows how information becomes available and how these two beliefs are formed.

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:

$$(17) \qquad F_h \;=\; F_{h+1} \;+\; \varepsilon_h^F, \qquad \varepsilon_h^F \sim \mathcal{N}(0, \sigma_F^2),$$

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

Monthly Inflation Beliefs Forecasters combine the public signal x_h^{AR} in (13) and their private signal $\tilde{x}_h^i$ in (14) to construct an individual monthly inflation belief $\hat{x}_h^i$:

$$(18) \qquad \hat{x}_h^i \;\equiv\; \mathbb{E}[x_h|\mathcal{I}_h^i] \;=\; \frac{\sigma_x^{-2} x_h^{AR} + \sigma_\zeta^{-2}\tilde{x}_h^i}{\sigma_x^{-2} + \sigma_\zeta^{-2}} \;=\; (1-\alpha)x_h^{AR} + \alpha\tilde{x}_h^i.$$

where we define the Bayesian weight on the private signal as $\alpha \equiv \sigma_\zeta^{-2}/(\sigma_x^{-2} + \sigma_\zeta^{-2})$. The weight α increases in the precision of the private signal σ_ζ^{-2} and decreases in the precision of inflation σ_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 $\hat{\pi}_h^i$:

$$(19) \qquad \hat{\pi}_h^i \;=\; \underbrace{h\left(\frac{c_x}{1-\phi_x}\right) + \frac{1-\phi_x^h}{1-\phi_x}\left(\hat{x}_h^i - \frac{c_x}{1-\phi_x}\right)}_{\text{AR(1) projection using } h \text{ info}} + \underbrace{\sum_{j=h+1}^{12} x_j}_{\text{realized, } j>h}, \qquad h = 12, \ldots, 1.$$

The first part of the expression (19) 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 $\{\phi_x, \sigma_x^2\}$ and signal noise σ_ζ^2; it decreases with the horizon and is independent of agents' identity:[9]

(20)
$$\Sigma_h^\pi = \; [(1-\alpha)^2\sigma_x^2 + \alpha^2\sigma_\zeta^2]\left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 + \frac{\sigma_x^2}{(1-\phi_x)^2}\left[(h-1) - \frac{2\phi_x(1-\phi_x^{h-1})}{1-\phi_x} + \frac{\phi_x^2(1-\phi_x^{2(h-1)})}{1-\phi_x^2}\right].$$

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

[9] Appendix C.1 presents a detailed derivation of z_h^i and Σ_h^z.

part of (20) reflects the accumulated uncertainty caused by the remaining $(h-1)$ unforecastable shocks that will hit the process until the release date.

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

$$(21) \qquad z_h \;=\; h\left(\frac{c_x}{1-\phi_x}\right)+\frac{\phi_x(1-\phi_x^h)}{1-\phi_x}\left(x_{h+1}-\frac{c_x}{1-\phi_x}\right)\;+\sum_{j=h+1}^{12} x_j, \quad h=12,\ldots,1.$$

It is useful to establish a relationship between individual beliefs $\hat{\pi}_h^i$ under the information set $\mathcal{I}_h^i$ in (15) and public beliefs z_h under the information set $\mathcal{I}_h$ in (21). The following relationship likes individual and common beliefs:

$$(22) \qquad \hat{\pi}_h^i = \; z_h \;+\; \nu_h^i, \quad \text{with} \quad \nu_h^i \sim \mathcal{N}\left(0, \left[\frac{1-\phi_x^h}{1-\phi_x}\right]^2 \alpha^2(\sigma_x^2+\sigma_\zeta^2)\right).$$

where α is the updating weight defined in (18).

Discussion of assumptions Although participants interpret public information differently, the prediction that builds on an AR(1) process is a tractable and accurate proxy for a fixed-event forecast. We provide further discussion about the features and accuracy of the proxy in Section G. See Giacomini, Skreta and Turen (2020). In our setup, given timing assumptions and the fact that past monthly inflation is known, forecasters do not learn from the consensus.

3.3 Optimal forecast revision policy

Proposition 2 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 common 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 2. *The value of a forecaster i at horizon h with state $(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i)$ equals*

$$(23) \qquad \mathcal{V}_h(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i) \;=\; \min\{\underbrace{\mathcal{V}_h^I(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i)}_{inaction},\; \underbrace{\mathcal{V}_h^A(\hat{\pi}_h^i, \hat{F}_h)}_{action}\}$$

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 (19) *and* (20)*, and consensus beliefs in* (17)*.*

Optimal policy The solution entails an *horizon-specific* inaction region

$$\mathcal{R}_h \equiv \{(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i) : \mathcal{V}_h^I(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i) \geq \mathcal{V}_h^A(\hat{\pi}_h^i, \hat{F}_h)\} \tag{24}$$

and a reset forecast $f_h^{i*}(\hat{\pi}_h^i, \hat{F}_h)$, such that, for given the beliefs, the forecast remains unchanged $f_h^i = f_{h+1}^i$, i.e., a zero revision $\Delta f_h = f_h^i - f_{h+1}^i$, whenever $f_{h+1}^i \in \mathcal{R}_h$. The forecast changes to f_h^{i*}, i.e., a revision of size $\Delta f_h = f_h^{i*} - f_{h+1}^i$, whenever $f_{h+1}^i \notin \mathcal{R}_h$.

For a given past forecast f_{h+1}^i, we can express the inaction region as a function of beliefs $\mathcal{R}_h(\hat{\pi}_h^i, \hat{F}_h | f_{h+1}^i)$. It includes the set of beliefs such that it is optimal not to reset the forecast.

3.4 Equilibrium

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)

$$\hat{F}_h = g(\hat{F}_{h+1}, \delta) + \epsilon_h^{\hat{F}}, \quad \epsilon_h^{\hat{F}} \sim \mathcal{N}(0, \sigma_F^2) \tag{25}$$

(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* (19) *and the perceived consensus process* $\{\hat{F}_h\}$ *in* (25)*, forecasts* $\{f_h^i\}$ *are optimal and solve the stopping-time problem* (23) *with optimal values* $\mathcal{V}_h(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i)$*;*
2. *parameters* (δ, σ_F) *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:* $\mathbb{C}ov[\epsilon_h^F, \epsilon_j^F] = 0$ *and* $\mathbb{V}ar[\epsilon_h^F] = \sigma_F^2$*.*

Note that under the RPE, 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 (25). For this reason, Marcet and Nicolini (2003) described this as a learning model that is "internally consistent".

Model solution We compute the decision rules of forecasters using backward induction. Appendix D.2 derives expressions to compute the distributions of expected beliefs needed to compute the value functions.

3.5 Calibration

Externally set parameters We feed the AR(1) parameters estimated directly from the data. By relying on the available information to forecasters in real time, we estimate the AR(1) process parameters using a rolling window over the sample years. For the year-on-year monthly inflation process x_t, we estimate $(c_x, \phi_x, \sigma_x) = (0.013, 0.932, 0.036)$. More details are in Appendix B.

Internally calibrated parameters Using the simulated method of moments, we estimate values for the six remaining parameters by matching the cross-sectional moments in normal years. We calibrate three parameters: the strength of strategic concerns r; the fixed revision cost κ, 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 value of the hazard's slope informs the value of 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.[10]

Internal consistency of consensus beliefs Finally, recall that agents in our model assume a random walk process for the consensus in (17). We thus have an additional parameter to set, the perceived volatility of the consensus process σ_F. Assuming rational expectations, the consensus's perceived and actual probability distributions should coincide. This assumption imposes structure and disciplines the value of σ_F. See Appendix D.4.

Table III 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.[11]

[10] Alvarez, Lippi and Paciello (2011), Baley and Blanco (2019) and Argente and Yeh (2022) apply the idea that signal noise modulates the slope of the hazard rate to calibrate signal noise in price-setting models.

[11] See Appendix D.4 for more details on the consistency of consensus beliefs.

Table III – Internally calibrated parameters

Parameter		Value	Moment	Data	Model
κ	revision cost	0.05	$\Pr[\Delta f \neq 0]$	0.43	0.42
r	strategic concerns	0.41	$\mathbb{E}[abs(\Delta f)\|adjust]$	0.25	0.22
σ_ζ	signal noise	0.04	hazard slope	-0.04	-0.04

3.6 Suggestive evidence

Participants of the Bloomberg survey can log into the system and change their predictions as many times as they want during any month. To keep participation active, Bloomberg sets a regular monthly reminder asking forecasters for their latest predictions.

4 The model in action

This section explores various dimensions of the forecasting model.

4.1 Illustration for one forecaster

To explain the model's workings, Figure VII illustrates how one agent's beliefs, forecasts, revisions, and uncertainty evolve during one year.

The fist panel shows the agent's inflation beliefs $\hat{\pi}_h^i$ (green line) and consensus beliefs $\hat{F}_h^i$ (blue line). Both beliefs change from period to period. However, the agent's forecast f_h^i (gray line) exhibits lumpy behavior, remaining fixed for some periods, followed by considerable revisions that bring the forecasts closer to a linear combination of the two beliefs. We see that, towards the year's end, inflation beliefs are very close to the actual inflation realization, but the forecast is not. Revision costs do not compensate for higher accuracy in forecasting.

The third panel shows the extensive margin of adjustment (gray areas), marking the periods of inaction between horizons 9 and 7 and 3 and 1. It also shows the intensive margin of adjustment given by the revision size Δf (dark line). Lastly, the fourth panel shows belief uncertainty (dashed pink line), which equals the weighted sum of conditional variances of inflation beliefs Σ_h^π and consensus beliefs σ_F^2 and is continuously decreasing; it also shows the mean squared forecast error (solid pink line), which equals the average error computed using forecasts and is non-monotonic. Moreover, uncertainty falls to zero while the mean-squared error remains positive at the end of the year.

4.2 Forecasts vs. beliefs

The left panel plots the realized value for the end-of-year inflation $\pi = 1.85$ (solid red line) and the cumulative sum of monthly inflation realizations $\sum_{h=12}^{1} x_h$ (dashed-dotted red line), which by definition equals π at $h = 1$. We also show the average belief $\hat{\Pi}_h \equiv N^{-1} \sum_{i=1}^{N} \hat{\pi}_h^i$ (black solid

Figure VII – Simulation for one forecaster

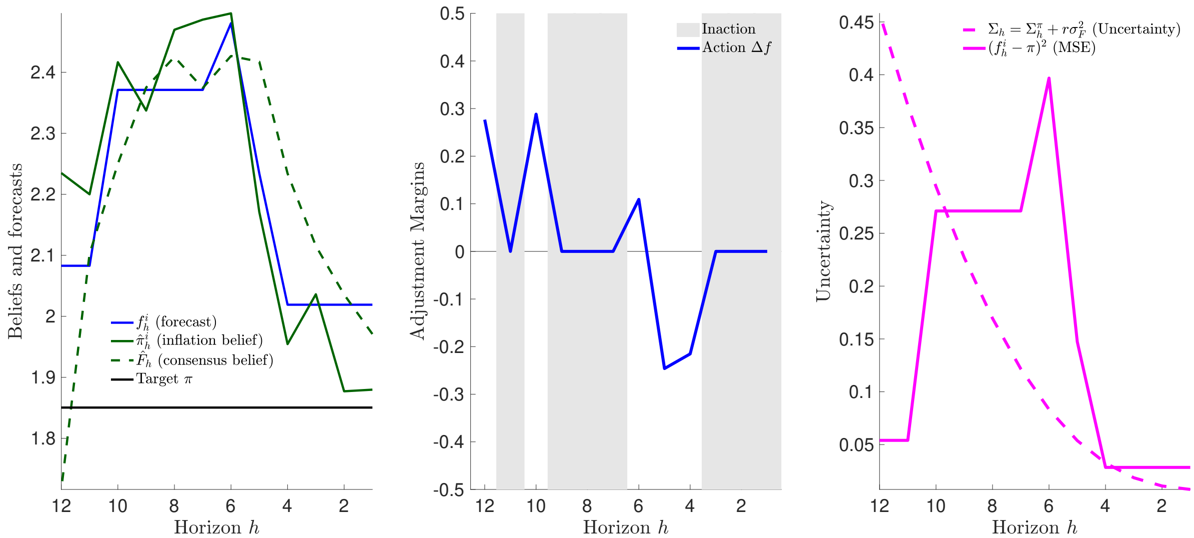

Notes: Illustration of one forecaster for one year.

line) and the average or census forecast $F_h \equiv N^{-1} \sum_{i=1}^{N} \hat{f}_h^i$ (blue dashed line). We observe that these two aggregate series get closer to π as the year goes by, and at the end of the year, average forecasts do not fully close the gap with the target variable. We also see that average beliefs are more volatile than the average forecast.

To highlight the difference between forecasts and beliefs, the right panel in Figure VIII shows the distribution of *non-zero* forecast revisions Δf_h^i in the data and the model, as well as inflation belief revisions $\Delta \hat{\pi}_h^i$ recovered from the model, across all years and horizons. Both distributions are centered around zero. The belief distribution is unimodal, while the forecast distribution is bimodal in both data and model.

4.3 Term structure of cross-sectional statistics

We assess the model's capacity to generate the term structure of cross-sectional statistics across horizons. Figure D.8 shows the term structure of frequency of revisions, variance of revisions, and forecast errors. While we only targeted average values, the model does a great job tracking their patterns along the forecasting period. Both the frequency of revisions (extensive margin, figure IXa) and the size of revisions conditioning on adjustment (intensive margin, figure D.8a). The slope of the hazard is matched throughout all forecast ages.

Volatility vs. option effects The model predicts these two margins would decrease over the horizon in both cases. Hence, although the inaction region is horizon-dependent, at longer horizons, the volatility effect dominates (when there is higher uncertainty about π) relative to the option value effect. At shorter horizons, the option value effects dominate, leading to less updating

Figure VIII – Time-series and cross-sectional distributions

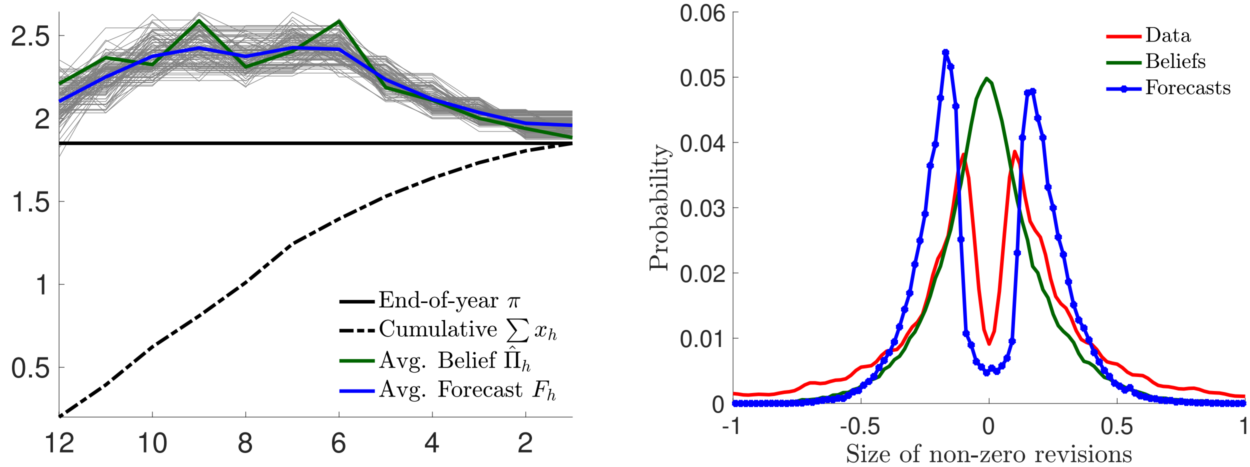

Notes: Model simulation for 100 forecasters with baseline parameterization. Sample pools together forecasters, years and horizons.

frequency. Moreover, it is precisely at shorter horizons when the contribution of each extra piece of information (monthly release) would only marginally affect the inflation belief, making it less likely to hit out of the bands, leading to fewer revisions as predicted by the data. The contribution of each extra bit of information only marginally affects the inflation beliefs, which is also reflected in the decreasing magnitude of forecast revisions. These implications, taken together, imply a decreasing hazard rate. The "volatility" effect dominates the "option value" effect.

Hazard rate The downward-sloping hazard observed in the data is a tell-tale sign of learning with a state-dependent model of forecast revisions. To see this, consider two alternative models. First, consider a model in which forecasters do not face revision costs but instead are "inattentive" and revise forecasts with a constant probability (or degree of attention) at random dates as in Andrade and Le Bihan (2013). In that model, the hazard rate is flat as the likelihood of revision is the same across all forecast ages (akin to the Calvo (1983) price-setting model). Second, consider a model with revision costs but without learning (uncertainty is constant). In that model, the hazard rate is increasing over the forecast's age, as a recently set forecast is at the center of the inaction region, and it takes time for each to reach either border of action. The combination of learning and revision costs delivers a decreasing hazard.

4.4 Untargeted moments

We present additional suggestive evidence of lumpy forecast revisions.

Consensus response We examine how the consensus gap c_h^i shapes the probability of updating a prediction. Figure X plots the average frequency of upward revisions (left column) and the

Figure IX – Cross-sectional statistics across horizons

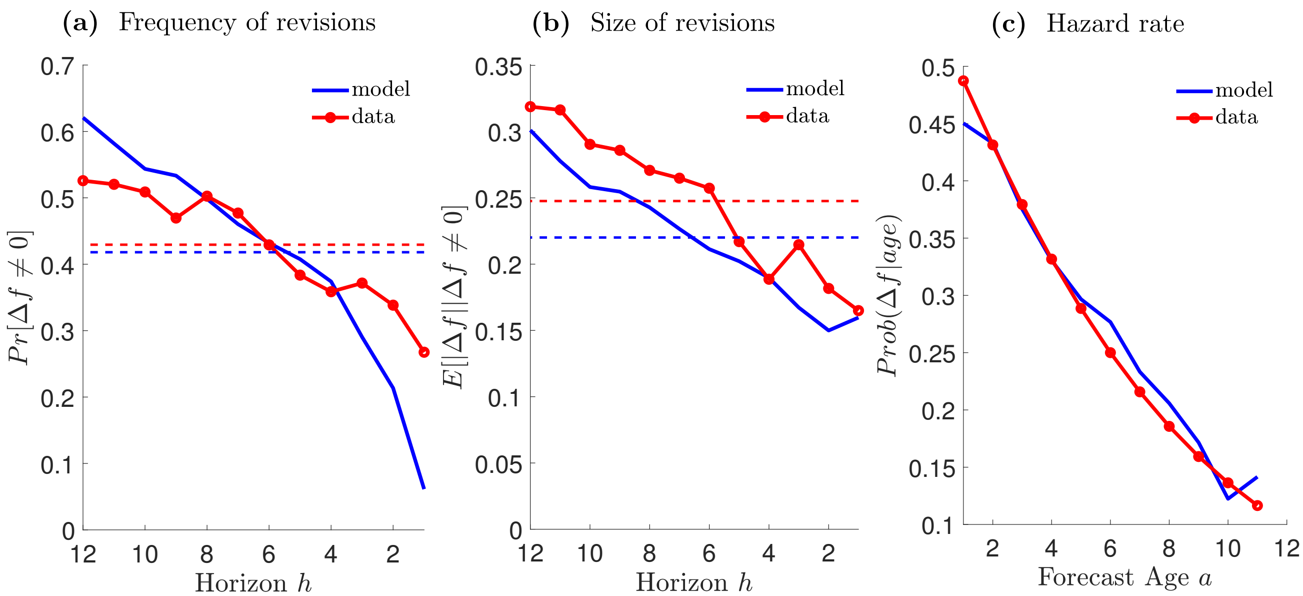

Notes: Model, benchmark calibration.

average frequency of downward revisions (right column). The patterns are qualitatively consistent with the data, although the model is much more responsive.[12]

Mean squared forecast errors Mean squared forecast errors $\mathbb{E}[(\pi - f_h^i)^2]$ (figure XI) decrease with the horizon; this is unsurprising as more information accumulates as the end-of-year event approaches.

Table IV – Autocorrelations: Model vs. Data

		Data	**Model**
Forecast errors	$corr[e, e_{-1}]$	0.877	0.626
All revisions	$corr[\Delta f, \Delta f_{-1}]$	-0.043	-0.161
Non-zero revisions	$corr[\Delta f, \Delta f_{-1} \vert \Delta f \neq 0]$	-0.107	-0.296

Notes: Data for normal years. Model with benchmark calibration.

Autocorrelations Second, we present autocorrelations of forecast errors and revisions. Table IV presents autocorrelations for forecast errors and revisions (for all and non-zero revisions). The model values are very close to those in the data.

[12] Introducing free adjustment opportunities, in the spirit of the CalvoPlus model of price-setting, which combines state and time-dependent adjustments, may bring the model closer to data.

Figure X – Consensus and Probability of Revision: Data vs. Model

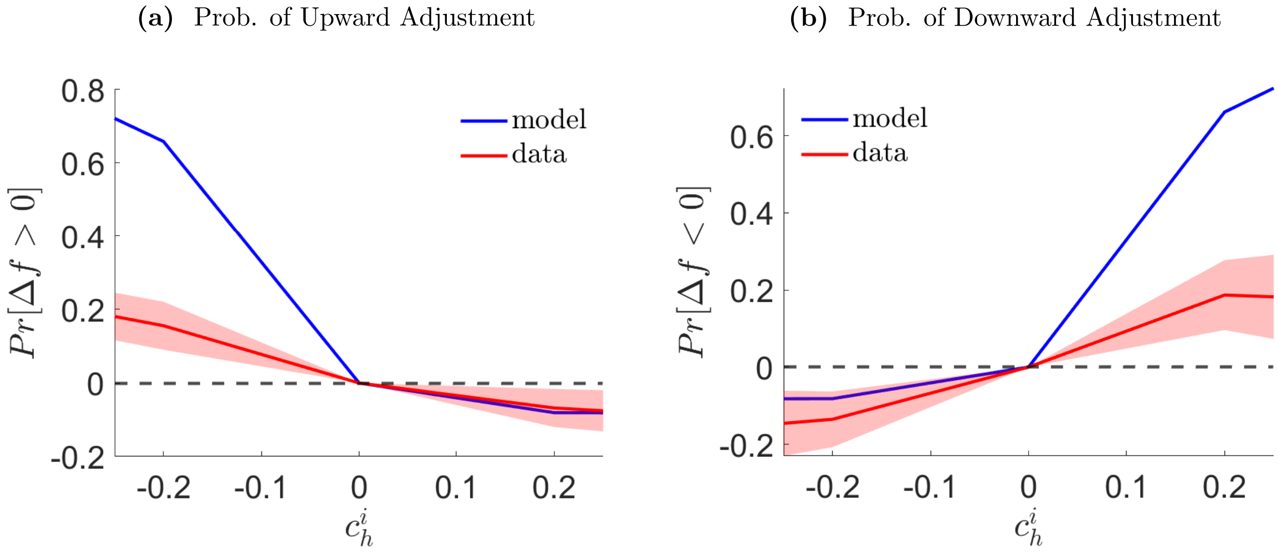

Notes: Data for normal years. The model with benchmark calibration.

Figure XI – Mean squared error

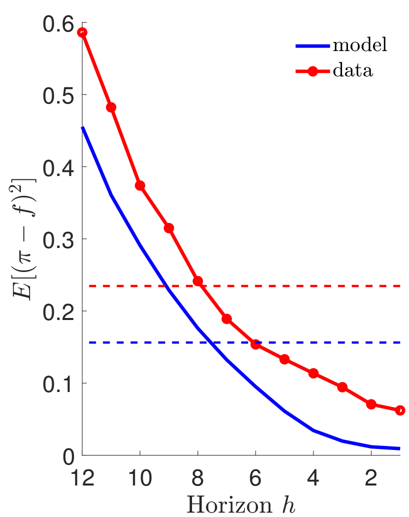

4.5 Forecaster heterogeneity

In our first exercise, we explore heterogeneity across forecaster types. The survey contains four types of forecasters: (i) banks, (ii) financial and investment institutions, (iii) economic consulting companies, and (iv) universities and research centers. Table V shows substantial heterogeneity in average cross-sectional moments in normal years. The most significant differences occur between consulting firms and universities. For instance, relative to universities, consulting firms adjust 30% more than universities.

Table V – Cross-sectional moments by forecaster type

	All		Financial Inst.		Banks		Consulting		Universities	
Moment	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model
$\Pr[\Delta f \neq 0]$	0.43	0.43	0.45	0.46	0.38	0.37	0.47	0.45	0.34	0.34
$\mathbb{V}ar[\Delta f]$	0.05	0.05	0.06	0.06	0.06	0.07	0.08	0.07	0.06	0.07
$\mathbb{E}[(\pi - f)^2]$	0.23	0.21	0.28	0.21	0.22	0.23	0.25	0.23	0.24	0.24
Observations	12,355		5,366		2,567		2,982		1,440	

Notes: Bloomberg data. Data for normal years 2010-2019.

Calibration by forecaster type The differences in the cross-sectional moments of forecast revisions and forecast errors reflect potential differences in strategic concerns (r) and preference for forecast stability (κ). Given the observed heterogeneity, we recalibrate four versions of the model, each matching type-specific moments. We consider heterogeneity in two parameters: fixed costs κ_G and strategic concerns r_G. Results are reported in Table VI. Through the lens of the model, these moments imply that universities face higher revision costs and stronger concerns for forecast stability. Thus, our results suggest that forecast heterogeneity is an important dimension to consider when working with this type of survey.

Table VI – Calibration by forecaster type

Parameter	**All**	**Financial Inst.**	**Banks**	**Consulting**	**Universities**
κ	0.05	0.06	0.07	0.05	0.08
r	0.41	0.70	0.45	0.66	0.14
σ_ζ	0.04	0.04	0.03	0.03	0.05
σ_F	0.11	0.14	0.13	0.11	0.15

Notes: Calibration that targets the group-specific moments reported in Table V. The last line reports the RPE-implied consensus volatility σ_F.

Next, we show the parameters relative to the values obtained for financial institutions.

Table VII – Calibration by forecaster type, relative to financial institutions

Parameter	**Financial Inst.**	**Banks**	**Consulting**	**Universities**
κ	1.00	1.17	0.83	1.33
r	1.00	0.64	0.94	0.20
σ_ζ	1.00	0.75	0.75	1.25

Notes: Parameters relative to financial institutions

5 Forecast rationality tests

In this section, we show how lumpy forecasts affect forecast efficiency tests.

5.1 Rationality tests using forecast errors

The first type of forecast rationality analyzes forecast error predictability at the individual level. The tests build on the work by Bordalo, Gennaioli, Ma and Shleifer (2020), extended by Broer and Kohlhas (2022) and Valchev and Gemmi (2023) to incorporate the consensus as a source of public information.

Test specification and interpretation Let $\pi_t - f^h_{t-h}$ be the individual forecast error at horizon h about annual inflation (known at time t), $f^i_{t-h} - f^i_{t-h+1}$ be the forecast revision between consecutive months for the same variable, and $F_{t,h} - f^i_{t,h+1}$ be the surprise contained in public information. Relying on the panel structure, we consider the following OLS regression with forecaster (α_i) and year (α_t) fixed effects:

$$\pi_t - f^i_{t,h} \; = \; \gamma^h_0 \; + \; \gamma^h_1(f^i_{t,h} - f^i_{t,h+1}) \; + \; \gamma^h_2 \; (F_{t,h} - f^i_{t,h+1}) \; + \; \alpha_i \; + \; \alpha_t \; + \; \epsilon^i_{t,h}. \tag{26}$$

Relative to the literature, which runs this regression at fixed horizon h, we obtain different coefficients for each horizon.

The coefficients are interpreted in the following way. If $\gamma^h_1 > 0$, a positive revision predicts a higher realization of inflation relative to the forecast, meaning that the average forecaster under-reacts to their information. In contrast, if $\gamma^h_1 < 0$ indicates that the average forecasters overreact to his information. Analogously, the sign of γ^h_2 reflects how information in public surprises affects forecast errors.

We run the regression in the survey data (red line) and model-generated data under the baseline calibration (blue line). Results are plotted in Figure XII. Consistent with the empirical regressions, forecasts in our model feature (i) the zero bias (left panel), (ii) over-reaction to private information (middle panel), and (iii) under-reaction to the consensus (right panel). These patterns hold across all horizons. Standard errors are two-way clustered at the forecaster and year level.

Next, we repeat the regressions using inflation beliefs $\hat{\pi}_{i,t}$ recovered from the model (green line). The results can be interpreted as those that would arise in a frictionless environment without a preference for stability $\kappa = 0$ and without strategic concerns $r = 0$. In this case, the consensus becomes the average inflation belief, denoted by $\hat{\Pi}_{t,h} \equiv N^{-1} \sum_{i=1}^N \hat{\pi}_{i,t}$. The specification is as follows:

$$\underbrace{\pi_t - \hat{\pi}^i_{t,h}}_{\text{error}} \; = \; \underbrace{\gamma^h_0}_{\text{bias}} \; + \; \gamma^h_1 \underbrace{(\hat{\pi}^i_{t,h} - \hat{\pi}^i_{t,h+1})}_{\text{revision}} \; + \gamma^h_2 \; \underbrace{(\hat{\Pi}_{t,h} - \hat{\pi}^i_{t,h+1})}_{\text{public info}} \; + \epsilon^i_h. \tag{27}$$

Proposition (3) shows that model-implied beliefs are rational ($\gamma^h_0 = \gamma^h_1 = 0$), even under information frictions.

Figure XII – Beliefs rationality tests

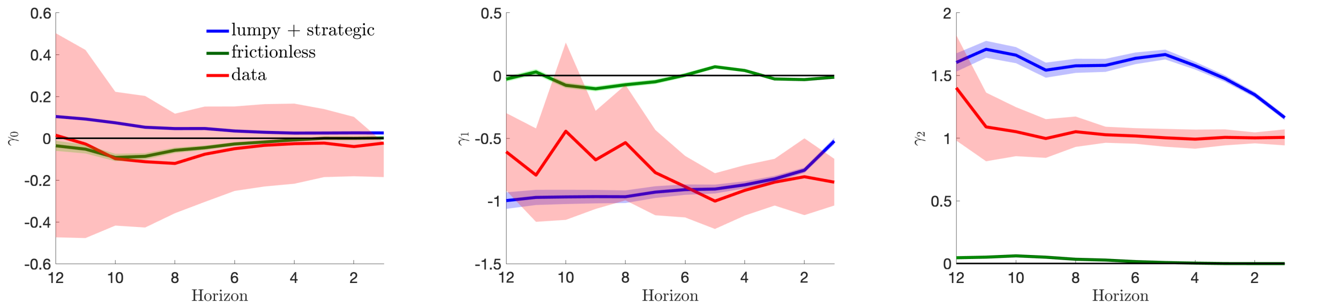

Notes: Data for normal years. Model under benchmark calibration.

Proposition 3. *Assume there is no preference for forecast stability ($\kappa = 0$) and no strategic concerns ($r = 0$). Then:*

(i) individual forecasts are equal to individual beliefs: $f_h^i = \widehat{\pi}_h^i$

(ii) efficiency tests confirm rationality (forecast errors are unpredictable): $\gamma_0^h = \gamma_1^h = \gamma_2^h = 0$.

5.2 Rationality tests using multi-horizons predictions

The previous tests build on the potential predictability of forecast errors at different horizons. While informative, the muti-horizon framework we exploit in this paper allows us to test rationality in a much richer context where we can rely on all the available information across the term structure of the forecasts. We are interested in extending the previous results to stress whether the evidence deviates from the presence of *joint* rationality throughout all horizons h. Thus, we rely on tests for forecast rationality that exploit the information across multi-horizons.

We follow Patton and Timmermann (2012)'s rationality test, which proposed and estimated a battery of different tests aiming to uncover the internal consistency of the set of forecasts. Specifically, the paper introduces ten different implications of rational behavior that should be fulfilled under the assumption that the forecaster's loss function is quadratic and the predicted process is covariance stationary. In each case, an empirical test for these implications can be derived. Seven conditions imply methods for handling multivariate inequality tests. On the other hand, the three are traditional coefficient tests after the MZ regressions. Notably, for all cases, Patton and Timmermann (2012) proposed tests constructed under the null hypothesis of rationality. Hence, situations when we fail to reject the null support the presence of internal consistency of predictions in line with rational behavior. Given all the tests, the authors rely on Bonferroni bounds to obtain a joint test.

Table VIII shows the p-values associated with each rationality test. The last three rows of the Table report the Bonferroni-based combination tests. We conduct the tests for the Bloomberg data (Column I), the simulated forecasts from the calibrated model (Column II), and simulated beliefs from the calibrated model. According to the results for both the data and the baseline model, we rejected the presence of rationality across the multi-horizons. Moreover, only for the "Beliefs" scenario are the results consistent with the internal consistency of forecasters.

Table VIII – Multi-horizon fixed-event rationality tests

Test	(I) Data	(II) Forecasts	(III) Beliefs
Inequality-based			
1. Increasing MSE	0.879	0.966	0.985
2. Increasing MSR^r	0.851	0.931	0.972
3. Decreasing MSF^r	**0.038**	0.94	0.944
4. Decreasing covariance	**0.071**	0.937	0.966
5. Decreasing covariance with proxy^r	**0.067**	0.907	0.959
6. Variance bound	0.16	0.774	0.771
7. Variance bound with proxy^r	**0.092**	0.711	0.825
MZ Regresion-based			
8. Univar opt revision	NaN	1.000	1.000
9. Univar opt revision with proxy	1.000	1.000	1.000
10. Univar MZ short h	**0.000**	**0.000**	0.194
Joint tests			
Bonferroni I (1+4 +6 +8 +10)	**0.000**	**0.000**	0.968
Bonferroni II (2+ 3+ 5+ 7+9)	0.192	1.000	1.000
Bonferroni All (1-10)	**0.000**	**0.000**	1.000

Notes: r = test does not rely on target variable.

6 Response to inflation volatility

Vavra, Lippi

In our third and last exercise, we investigate the forecasts' response to changes in underlying inflation volatility. In four years in our sample, economic conditions were far from ordinary. In particular, at the two ends of the sample, we have the Great Recession between 2008 and 2009 and the COVID-19 pandemic between 2020 and 2021. The volatility of monthly inflation differed over these two episodes: in turbulent years, inflation volatility is 10% higher than in normal years. We split the sample between normal years, 2010-2019, and turbulent years, 2008-2009 and 2020-21.

Figure XIII below shows the term structure of (a) frequency of revisions, (b) variance of revisions, and (c) forecast errors squared. In turbulent years, forecast revisions are more frequent and dispersed, and forecast errors sharply increase.

We explore the model's ability to deliver these changes in cross-sectional moments. We implement increasing monthly inflation volatility as in the data, $1.1 \times \sigma_x^2$. We also consider two different unconditional means. Low c_x for 2008-2009 and high c_x for 2020-2021.

Figure XIII – Term Structure of Forecast Revisions: Normal vs. Turbulent Tears

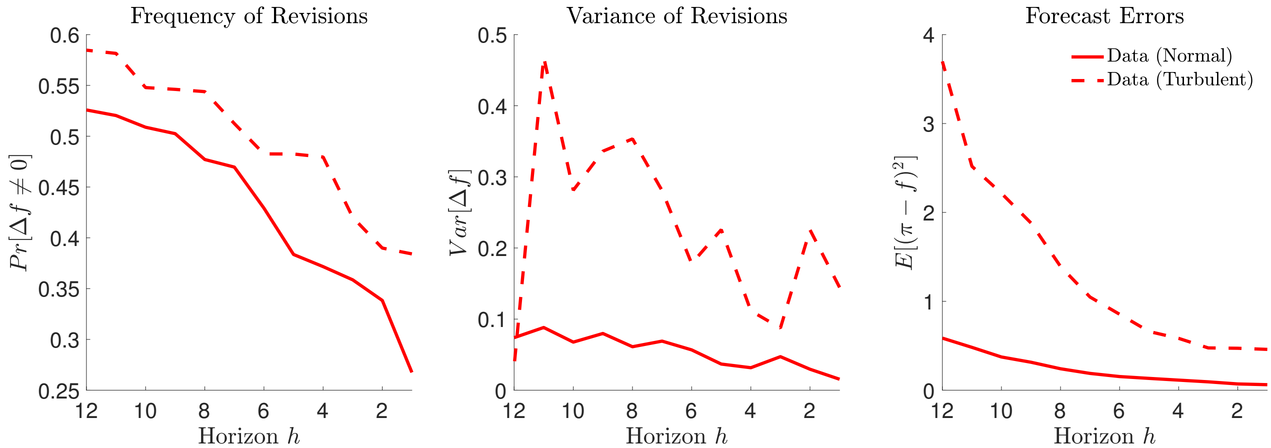

Notes: Bloomberg data. Notes: Normal years = 2010-2019. Turbulent years = 2008-2009 and 2020-21.

We give forecasters the information that volatility and mean have changed. Figure XIV shows the results. Relative to the baseline calibration (solid line that matches moments in normal years), higher volatility increases all moments, regardless of the change in the mean.

Figure XIV – Response to changes in inflation volatility

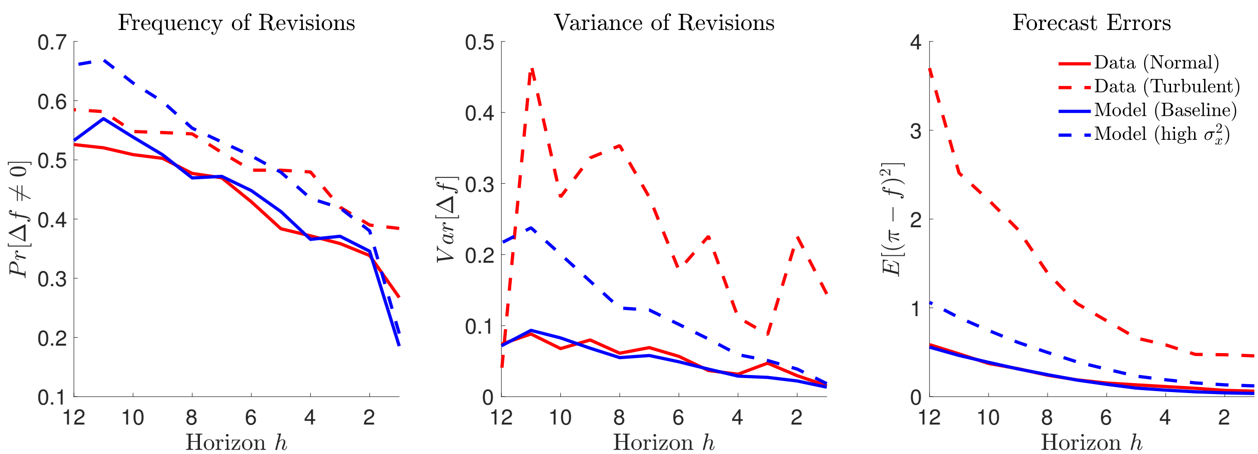

Notes: Data for normal years. Model with benchmark calibration.

7 Conclusion

TBC. In parallel work, we study how lumpy forecasts affect the transmission of monetary policy shocks to expectations. Implications for survey design.

References

Álvarez, F., Lippi, F. and Paciello, L. (2018). Monetary shocks in models with observation and menu costs. *Journal of the European Economic Association*, **16** (2), 353–382.

Alvarez, F. E., Lippi, F. and Paciello, L. (2011). Optimal price setting with observation and menu costs. *The Quarterly journal of economics*, **126** (4), 1909–1960.

Andrade, P. and Le Bihan, H. (2013). Inattentive professional forecasters. *Journal of Monetary Economics*, **60** (8), 967–982.

Argente, D. and Yeh, C. (2022). Product life cycle, learning, and nominal shocks. *Review of Economic Studies*, **Forthcoming**.

Bachmann, R., Topa, G. and van der Klaauw, W. (2022). *Handbook of Economic Expectations*. Elsevier.

Bahaj, S., Czech, R., Ding, S. and Reis, R. (2023). *The market for inflation risk*. Tech. rep., Bank of England.

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

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

Bec, F., Boucekkine, R. and Jardet, C. (2023). Why are inflation forecasts sticky? theory and application to france and germany. *International Journal of Central Banking*, **19** (4), 215–249.

Bordalo, P., Gennaioli, N., Ma, Y. and Shleifer, A. (2018). Over-reaction in macroeconomic expectations, working Paper.

—, —, — and — (2020). Overreaction in macroeconomic expectations. *American Economic Review*, **110** (9), 2748–82.

Born, B., Enders, Z., Müller, G. J. and Niemann, K. (2022). *Firm expectations about production and prices: Facts, determinants, and effects*. Tech. rep., mimeo.

Broer, T. and Kohlhas, A. N. (2022). Forecaster (mis-) behavior. *Review of Economics and Statistics*, pp. 1–45.

Calvo, G. (1983). Satggered prices in a utility maximizing framework. *Journal of Monetary Economics*, **12**, 383–398.

Capistrán, C. and López-Moctezuma, G. (2014). Forecast revisions of mexican inflation and gdp growth. *International Journal of Forecasting*, **30** (2), 177–191.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, **105** (8), 2644–78.

Croushore, D. D. (1997). The livingston survey: Still useful after all these years. *Business Review-Federal Reserve Bank of Philadelphia*, **2**, 1.

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

García-Schmidt, M. and Woodford, M. (2019). Are low interest rates deflationary? a paradox of perfect-foresight analysis. *American Economic Review*, **109** (1), 86–120.

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

Karadi, P., Schoenle, R. and Wursten, J. (2021). Measuring price selection in microdata: it's not there.

Kozlowski, J., Veldkamp, L. and Venkateswaran, V. (2020a). *Scarring body and mind: the long-term belief-scarring effects of Covid-19*. Tech. rep., National Bureau of Economic Research.

—, — and — (2020b). The tail that wags the economy: Beliefs and persistent stagnation. *Journal of Political Economy*, **128** (8), 2839–2879.

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

Mankiw, G. and Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *Quarterly Journal of Economics*, **117**, 1295–1328.

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

Mincer, J. A. and Zarnowitz, V. (1969). The evaluation of economic forecasts. In *Economic forecasts and expectations: Analysis of forecasting behavior and performance*, NBER, pp. 3–46.

Molavi, P. (2022). Simple models and biased forecasts. *arXiv preprint arXiv:2202.06921*.

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

NORDHAUS, W. D. (1987). Forecasting efficiency: concepts and applications. *The Review of Economics and Statistics*, pp. 667–674.

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.

— and — (2011). Predictability of output growth and inflation: A multi-horizon survey approach. *Journal of Business & Economic Statistics*, **29** (3), 397–410.

— and — (2012). Forecast rationality tests based on multi-horizon bounds. *Journal of Business & Economic Statistics*, **30** (1), 1–17.

SIMS, C. (2003). Implications of rational inattention. *Journal of Monetary Economics*, **50(3)**, 665–90.

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

WOODFORD, M. (2003). Imperfect common knowledge and the effects of monetary policy. *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, p. 25.

Lumpy Forecasts

Isaac Baley and Javier Turén

Online Appendix

Contents

A Data

A.1 Inflation definitions

We show how to approximate yearly inflation with year-on-year monthly inflation. Let $\overline{cpi}_t = \frac{1}{12}\sum_{h=1}^{12} cpi_{t,h}$ be the average cpi in year t. Then, the annual inflation equals:

$$
\begin{aligned}
\text{(A.1)} \qquad \pi_t &= \log(\overline{cpi}_t) - \log(\overline{cpi}_{t-1}) \\
&= \log\left(\frac{1}{12}\sum_{h=1}^{12} cpi_{t,h}\right) - \log\left(\frac{1}{12}\sum_{h=1}^{12} cpi_{t-1,h}\right) \\
&\approx^{Jensen} \frac{1}{12}\sum_{h=1}^{12}\log\left(cpi_{t,h}\right) - \frac{1}{12}\sum_{h=1}^{12}\log\left(cpi_{t-1,h}\right) \\
&= \frac{1}{12}\sum_{h=1}^{12}\left(\log\left(cpi_{t,h}\right) - \log\left(cpi_{t-1,h}\right)\right) \\
&= \frac{1}{12}\sum_{h=1}^{12}\left(\log\left(cpi_{h}\right) - \log\left(cpi_{h+12}\right)\right) \\
&= \sum_{h=1}^{12}\underbrace{\frac{1}{12}\left(\log\left(cpi_{h}\right) - \log\left(cpi_{h+12}\right)\right)}_{x_h} \\
&= \sum_{h=1}^{12} x_h
\end{aligned}
$$

When is year-on-year monthly inflation a good approximation of annual inflation? To answer, consider a second-order Taylor approximation of $\log(p)$ around $\mathbb{E}[p]$, which yields:

$$\text{(A.2)} \qquad \log(p) \approx \log(\mathbb{E}[p]) + \frac{1}{\bar{p}}(p - \mathbb{E}[p]) - \frac{1}{2\mathbb{E}[p]^2}(p - \mathbb{E}[p])^2$$

Take expectations on both sides (note that $\mathbb{E}[p]$ is a constant):

$$\text{(A.3)} \qquad \mathbb{E}[\log(p)] \approx \log(\mathbb{E}[p]) - \frac{\mathbb{V}ar[p]}{2\mathbb{E}[p]^2} = \log(\mathbb{E}[p]) - \frac{\mathbb{CV}^2[p]}{2}$$

Applying the decomposition to annual inflation, letting p, p' be the CPI in consecutive years, we obtain:

$$\text{(A.4)} \qquad \pi = \log(\mathbb{E}[p]) - \log(\mathbb{E}[p']) = \underbrace{\mathbb{E}[\log(p) - \log(p')]}_{\text{average year-on-year inflation } \mathbb{E}[x]} + \underbrace{\frac{\mathbb{CV}^2[p] - \mathbb{CV}^2[p']}{2}}_{\text{differences in within-year dispersion}}$$

Therefore, for similar within-year price dispersion ($\mathbb{CV}^2[p] \approx \mathbb{CV}^2[p']$), then $\pi \approx \mathbb{E}[x]$.

A.2 Inflation summary statistics

Table A.1 shows the mean and the variance of yearly inflation for all years in our sample, separated between normal (2010-2019), Turbulent (2008-2009 & 2020-2021) and Pandemic years (2020-2021). years. We focus our analysis on normal and pandemic years, mostly since the observed inflation dynamics were very different between the Great Recession and the COVID-19 pandemic. While in the former episode, the US experienced a deflation (in fact, inflation was -0.3% in 2009), in the latter, during the COVID-19 pandemic, inflation spiked up to 4.7% during 2021. Given this fact and since these two episodes are also wide apart, we focus only on these two groups of years.

Table A.1 – Summary Statistics of Inflation

		All	**Normal**	**Turbulent**	**Pandemic**
Average	$\mathbb{E}[\pi]$	1.896	1.795	2.175	2.95
Volatility	$\mathbb{V}ar[\pi]$	1.621	0.622	4.267	2.478
Years		14	10	4	2

Notes: CPI Index.
Normal years = 2010-2019. Turbulent years = 2008-09 and 2020-21.
Pandemic years = 2020-2021.

A.3 Cross-sectional statistics by year

Figure A.1 – Adjustment Frequency and Size by Horizon and Year

(a) Frequency **(b)** Size **(c)** Forecast errors

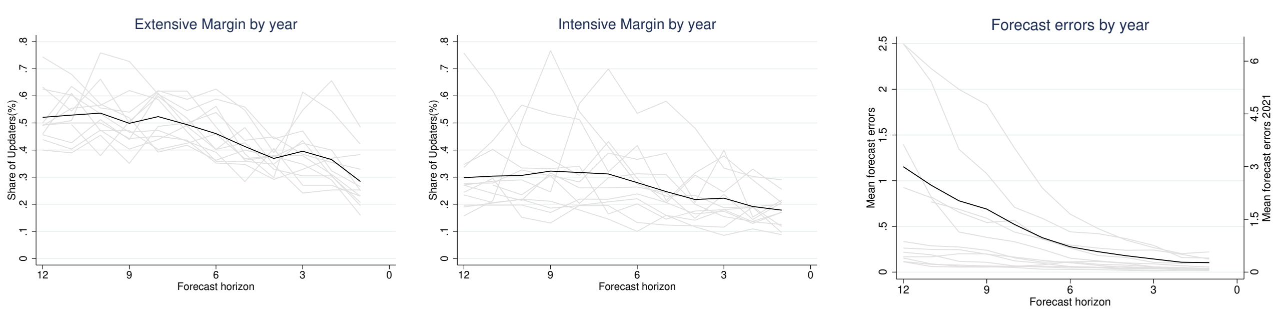

Notes: Bloomberg data.

A.4 Hazard rates by number of revisions

We compute the hazard rates based on the number of revisions the forecasters have done in the past. In this sense, we explore whether the age-dependence of updating probabilities changes as a function of the revision being the first, second, third, and so forth. This is shown in Figure A.2.

Figure A.2 – Hazard Rates by revisions

(a) Hazard Rate: 1st, 2nd and 3rd Rev. **(b)** Hazard Rate: 4th, 5th and 6th Rev.

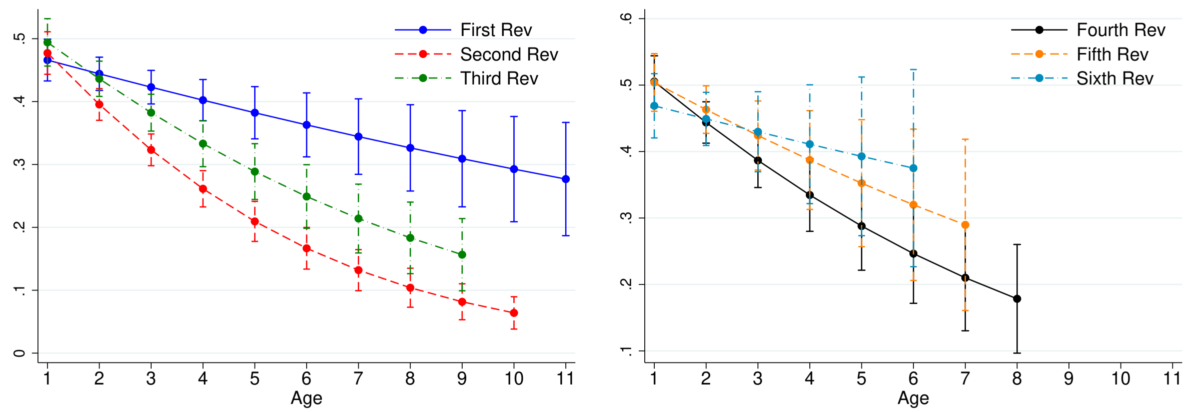

Notes: Bloomberg data.

Independently of the revision, the decaying pattern of the hazard rates remains across specifications. While the chance of an immediate revision right at age one is roughly the same across the number of revisions (between 45% and 50%), the likelihood drops as more revisions accumulate throughout the year. Although the relations are not monotonic, in most cases, the age probabilities are not statistically different across groups. We interpret the decaying probability as a function of revisions as an implication of the fixed-event scheme. In contrast, with age growths, we get closer to the final release date, and the fact that, as time passes, we accumulate more relevant information to predict annual inflation.

A.5 Cross-sectional statistics at long horizons

Notice that from 18 to 13 months ahead, there is information about future end-of-year inflation, but only for horizons 12 to 1 months ahead, inflation is realized.

A.6 Cross-sectional statistics by forecaster type

Figure A.4 shows the term structure of adjustment frequency, size, and hazard rate for each of the four groups. These term structures are broadly consistent with the general patterns observed for the average moments, with universities being the group that adjusts less often but for more significant amounts across horizons while consulting firms do the opposite. The hazard rates for forecasters belonging to either "Financial & Investment" or "Economic Consulting" are the steepest relative to the other two groups. Hence, although they decrease, the updating probability is less sensible to the age of both Banks and Universities.

(a) Frequency of revisions

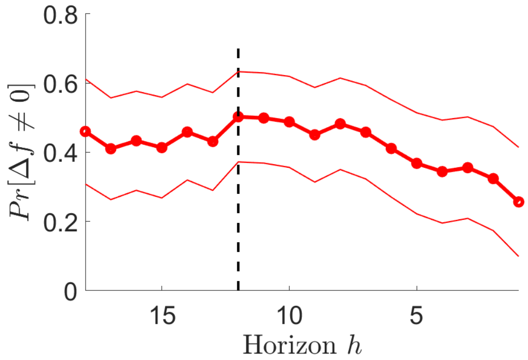

(b) Size of non-zero revisions

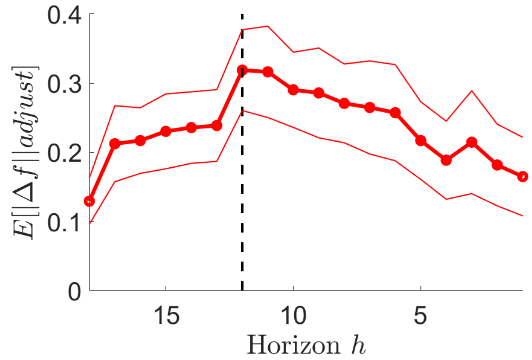

(c) Variance of revisions

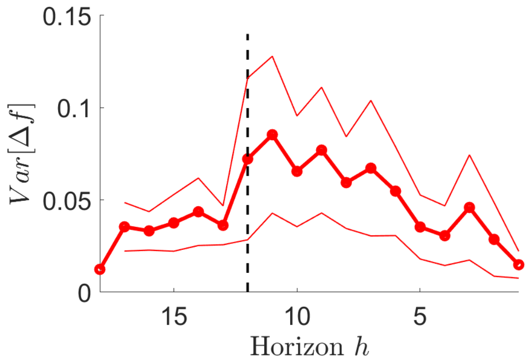

(d) Mean squared error

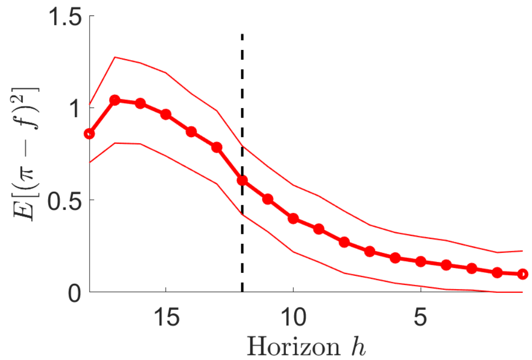

Figure A.4 – Term Structure of Revisions and Errors: By Forecaster Type

(a) Frequency

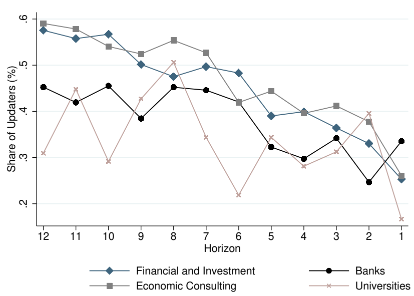

(b) Size

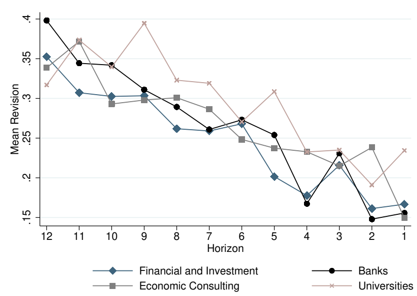

(c) Hazard

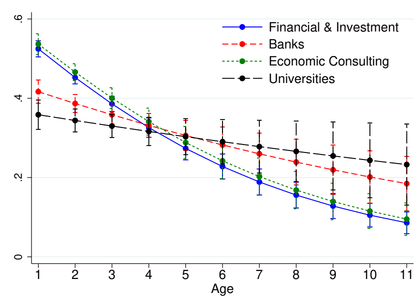

Notes: Bloomberg data.

A.7 Cross-sectional statistics in turbulent times

Figure A.5 – Forecast Errors by Groups

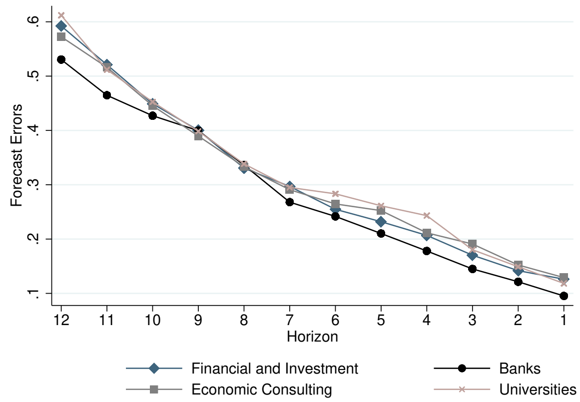

Notes: Bloomberg data.

Table A.2 – Summary Statistics of Forecast Revisions

		All	**Turbulent**	**Normal**
Average	$\mathbb{E}(\Delta f)$	-0.002	0.028	-0.013
Size	$\mathbb{E}(abs(\Delta f)\|\Delta f \neq 0)$	0.307	0.453	0.247
Dispersion	$\mathbb{V}ar(\Delta f)$	0.104	0.227	0.055
Number of revisions	$count(\Delta f \neq 0)$	5.204	5.602	5.059
Duration (months)	$\mathbb{E}(\tau)$	1.497	1.231	1.594
Inaction rate	$\Pr(\Delta f = 0)$	0.523	0.392	0.569
Frequency	$\Pr(\Delta f \neq 0)$	0.444	0.492	0.427
Upward	$\Pr(\Delta f) > 0$	0.228	0.318	0.196
Downward	$\Pr(\Delta f) < 0$	0.216	0.173	0.231
Spike rate	$abs(\Delta f/f) > 0.2$	0.081	0.231	0.028
Positive spikes	$\Delta f/f > 0.2$	0.076	0.227	0.023
Negative spikes	$\Delta f/f < -0.2$	0.005	0.004	0.005
Serial correlation (all)	$corr(\Delta f, \Delta f_{-1})$	-0.035	-0.035	-0.043
Serial correlation (non-zero)	$corr(\Delta f, \Delta f_{-1})$	-0.085	-0.078	-0.107
Annual Inflation	π	1.896	2.175	1.795
Observations	N	12,619	3,363	9,256

Sources: Bloomberg data. Notes: Normal years = 2010-2019. Turbulent years = 2008-2009 and 2020-21. Numbers are averages across sample years and 12 horizons.

Table A.3 – Summary Statistics of Forecast Errors

		All	**Turbulent**	**Normal**
Average	$\mathbb{E}(e)$	0.023	0.237	-0.055
Size	$\mathbb{E}(abs(e))$	0.434	0.789	0.305
Positive	$\Pr(e > 0)$	0.414	0.606	0.345
Negative	$\Pr(e < 0)$	0.510	0.340	0.572
Dispersion	$\sigma(e)$	0.663	1.105	0.502
Serial correlation	$corr(e, e_{-1})$	0.882	0.878	0.877
Observations	N	12,619	3,363	9,256

Sources: Bloomberg data. Notes: Normal years = 2010-2019. Turbulent years = 2008-2009 and 2020-21. Numbers are between-year averages of within-year statistics.

Figure A.6 – Time series

(a) Revision frequency **(b)** Revision size

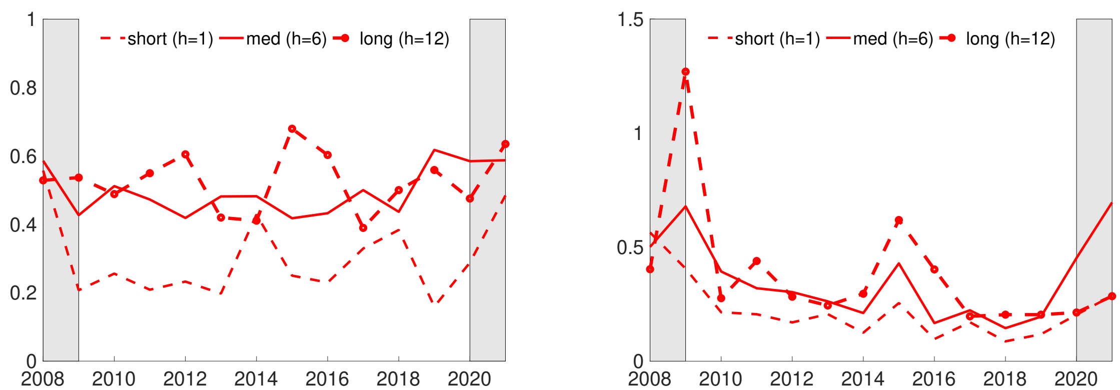

Notes: Data for normal years. Standard errors are clustered two ways at the forecaster and year level. Model with benchmark calibration.

B Estimate of inflation process

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

$$x_h \;=\; c_x + \phi_x x_{h+1} + \varepsilon_h^x, \qquad \varepsilon_h^x \sim \mathcal{N}(0, \sigma_x^2), \tag{B.5}$$

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. We estimate the three parameters $(c_x, \phi_x, \sigma_x^2) = (0.013, 0.932, 0.036)$ using the monthly inflation rate from the CPI for 2010-2019. Figure B.7 plots the resulting estimates and 95% confidence intervals.

Figure B.7 – Rolling Estimates for Inflation Parameters

(a) Constant **(b)** Persistence **(c)** Volatility

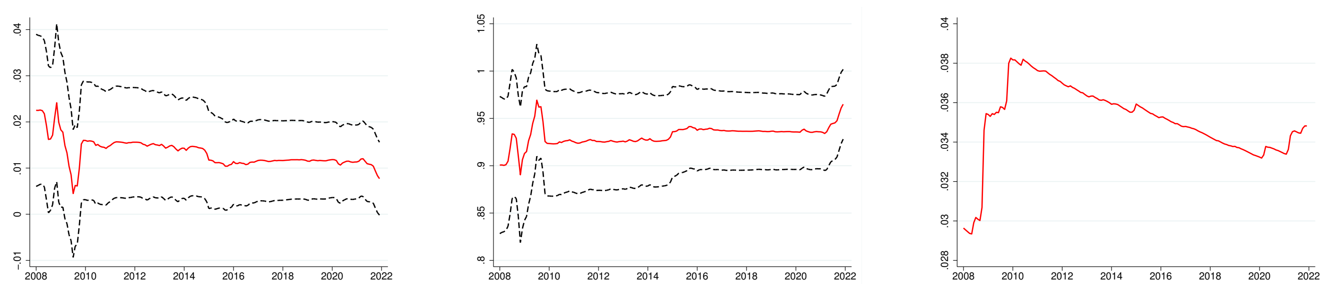

Notes: Bloomberg data.

C Proofs

C.1 Inflation process

Demeaned monthly inflation We begin with the assumption of an AR(1) process for monthly inflation:

$$x_h \;=\; c_x \;+\; \phi_x x_{h+1} \;+\; \varepsilon^x_h. \tag{C.1}$$

This process has an unconditional mean of $\frac{c_x}{1-\phi_x}$ and an unconditional variance of $\frac{\sigma^2_x}{1-\phi^2_x}$. For any h, we can rewrite (C.1) as deviations from the unconditional mean:

$$x_h - \frac{c_x}{1-\phi_x} \;=\; \phi_x\left(x_{h+1} - \frac{c_x}{1-\phi_x}\right) \;+\; \varepsilon^x_h. \tag{C.2}$$

Annual inflation Annual inflation π is approximately equal to the sum of the twelve realizations of monthly inflation x_h within each target year $\pi = \sum_{h=1}^{12} x_h$. See appendix A.1. Without loss of generality, we can derive π as a function of the initial value of monthly inflation x_{12}:

$$\begin{aligned}
x_1 &= \frac{c_x}{1-\phi_x} + \phi_x^{11}\left(x_{12} - \frac{c_x}{1-\phi_x}\right) + \sum_{j=0}^{10}\phi_x^j\varepsilon^x_{j+1} \\
&\dots \\
x_{10} &= \frac{c_x}{1-\phi_x} + \phi_x^{2}\left(x_{12} - \frac{c_x}{1-\phi_x}\right) + \phi_x\varepsilon^x_{11} + \varepsilon^x_{10} \\
x_{11} &= \frac{c_x}{1-\phi_x} + \phi_x\left(x_{12} - \frac{c_x}{1-\phi_x}\right) + \varepsilon^x_{11},
\end{aligned}$$

Summing up the monthly values $x_1, x_2, \dots, x_{12}$ we get an expression for annual inflation at horizon $h = 12$:

$$\pi \;=\; 12\left(\frac{c_x}{1-\phi_x}\right) \;+\; \frac{1-\phi_x^{12}}{1-\phi_x}\left(x_{12} - \frac{c_x}{1-\phi_x}\right) \;+\; \sum_{j=1}^{11}\frac{1-\phi_x^j}{1-\phi_x}\varepsilon^x_j. \tag{C.3}$$

Similarly, for any h within the year, we can derive an expression for π. Importantly, as h shrinks (as we get closer to the release date), we start summing the actual lagged values of inflation starting at $h = 12$ until h while we project the remaining months of the year using the last piece of available information x_h. In particular, annual inflation at any given horizon $h = 12, 11, \dots, 1$ can be written as follows:

$$\pi \;=\; h\left(\frac{c_x}{1-\phi_x}\right) \;+\; \frac{(1-\phi_x^h)}{1-\phi_x}\left(x_h - \frac{c_x}{1-\phi_x}\right) \;+\; \sum_{i=h+1}^{12} x_j \;+\; \sum_{j=1}^{h-1}\frac{1-\phi_x^j}{1-\phi_x}\varepsilon^x_j, \tag{C.4}$$

where $\sum_{i=h+1}^{12} x_j = 0$ for $i = 12$. If $h = 1$ then $\pi = \sum_{h=1}^{12} x_h$. The unconditional mean and variance of end-of-year inflation are:

$$\mathbb{E}[\pi] = \frac{12c_x}{1-\phi_x} \tag{C.5}$$

$$\mathbb{V}ar[\pi] = \sigma_x^2 \sum_{j=1}^{h-1} \left(\frac{1-\phi_x^j}{1-\phi_x}\right)^2. \tag{C.6}$$

To compute annual inflation from the perspective of $h = 13$, we use the fact that

$$x_{12} - \frac{c_x}{1-\phi_x} = \phi_x \left(x_{13} - \frac{c_x}{1-\phi_x}\right) + \varepsilon_{12}^x. \tag{C.7}$$

Thus, when summing up the monthly values $x_1, x_2, \ldots, x_{12}$, we get

$$\pi = 12\left(\frac{c_x}{1-\phi_x}\right) + \phi_x \frac{1-\phi_x^{12}}{1-\phi_x}\left(x_{13} - \frac{c_x}{1-\phi_x}\right) + \sum_{j=1}^{12} \frac{1-\phi_x^j}{1-\phi_x}\varepsilon_j^x. \tag{C.8}$$

C.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. In turn, the monthly beliefs are constructed using the AR(1) one-period ahead prediction and the private signal $\widetilde{x}_h^i = x_h + \zeta_{ih}$. In addition, the historical values of lagged monthly inflation are observed without noise. Thus, the forecasters information set at each horizon $\mathcal{I}_h^i = \{\widetilde{x}_h^i, x_{h+1}, x_{h+2}, \ldots\}$.

Conditional mean Taking the conditional expectation of equation (C.4), given information up to horizon h, delivers the conditional mean $\hat{\pi}_h^i \equiv \mathbb{E}[\pi|\mathcal{I}_t^i]$:

$$\hat{\pi}_h^i = h\left(\frac{c_x}{1-\phi_x}\right) + \frac{1-\phi_x^h}{1-\phi_x}\left(\hat{x}_h^i - \frac{c_x}{1-\phi_x}\right) + \sum_{i=h+1}^{12} x_j \qquad \text{for } h = 12, \ldots, 1 \tag{C.9}$$

which corresponds to equation (19) in the text.

Conditional variance To compute the conditional variance, we first define forecast errors as the difference between end-of-year inflation π in (C.4) and the conditional mean $\varepsilon_h^i \equiv \pi - \hat{\pi}_h^i$ in (C.9):

$$\varepsilon_h^i = \pi - \hat{\pi}_h^i = \frac{1-\phi_x^h}{1-\phi_x}((1-\alpha)\varepsilon_h^x + \alpha\zeta_{ih}) + \sum_{j=1}^{h-1} \frac{1-\phi_x^j}{1-\phi_x}\varepsilon_j^x \qquad \forall h \tag{C.10}$$

Where $\alpha \equiv \frac{\sigma_\zeta^{-2}}{\sigma_x^{-2}+\sigma_\zeta^{-2}}$ as discussed in the main text. Squaring and taking expectations, we obtain the variance of the forecast error $\Sigma_h^\pi \equiv \mathbb{E}[(\varepsilon_h^i)^2]$ at each horizon h:

$$\text{(C.11)} \qquad \Sigma_h^\pi \;=\; \left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 ((1-\alpha)^2\sigma_x^2 + \alpha^2\sigma_\zeta^2) \;+\; \frac{\sigma_x^2}{(1-\phi_x)^2}\sum_{j=1}^{h-1}(1-\phi_x^j)^2 \qquad \forall h$$

where we used that shocks are i.i.d $\varepsilon_h^x \overset{\text{iid}}{\sim} \mathcal{N}(0,\sigma_x^2)$, $\zeta_h^i \overset{\text{iid}}{\sim} \mathcal{N}(0,\sigma_\zeta^2)$, $\eta_h \overset{\text{iid}}{\sim} \mathcal{N}(0,\sigma_\eta^2)$ and uncorrelated $\mathbb{E}[\zeta_h^i, \eta_h] = 0$. We simplify the last term with the sum as follows:

$$\begin{aligned}
\sum_{j=1}^{h-1}(1-\phi_x^j)^2 \;&=\; (1-\phi_x)^2 \;+\; (1-\phi_x^2)^2 \;+\; \dots \;+\; (1-\phi^{h-1})^2 \\
&=\; 1 \;-\; 2\phi_x \;+\; \phi_x^2 \;+\; 1 \;-\; 2\phi_x^2 \;+\; \phi_x^4 \;+\; \dots \;+\; 1 \;-\; 2\phi^{h-1} \;+\; \phi^{2(h-1)} \\
&=\; (h-1) - 2(\phi_x + \phi_x^2 + \dots + \phi_x^{h-1}) + (\phi_x^2 + \phi_x^4 + \dots + \phi^{2(h-1)}) \\
&=\; (h-1) - \frac{2\phi_x(1-\phi_x^{h-1})}{1-\phi_x} + \frac{\phi_x^2(1-\phi_x^{2(h-1)})}{1-\phi_x^2}
\end{aligned}$$

Substituting back into (C.11), we obtain the expression for the signal variance in

$$\text{(C.12)} \qquad \Sigma_h^z \;=\; [(1-\alpha)^2\sigma_x^2 + \alpha^2\sigma_\zeta^2]\left(\frac{1-\phi^h}{1-\phi}\right)^2 + \frac{\sigma_x^2}{(1-\phi)^2}\left[(h-1) - \frac{2\phi(1-\phi^{h-1})}{1-\phi} + \frac{\phi^2(1-\phi^{2(h-1)})}{1-\phi^2}\right].$$

The conditional variance is common across forecasters; thus, we denote it as $\Sigma_{z,h}$.

C.3 Relationship between individual vs. aggregate beliefs

The public signal z_h in (21) projects the past release x_{h+1} to obtain the yearly forecast:

$$\text{(C.13)} \qquad z_h \;=\; h\left(\frac{c_x}{1-\phi_x}\right) \;+\; \frac{\phi_x(1-\phi_x^h)}{1-\phi_x}\left(x_{h+1} - \frac{c_x}{1-\phi_x}\right) \;+\; \sum_{j=h+1}^{12} x_j.$$

In contrast, the private signal $\hat{\pi}_h^i$ in (**??**) projects the noisy observation $\widetilde{x}_h^i$ to obtain the yearly forecast (note the extra ϕ in the second term of the expression above, reflecting the timing of the information):

$$\text{(C.14)} \qquad \hat{\pi}_h^i \;=\; h\left(\frac{c_x}{1-\phi_x}\right) \;+\; \frac{1-\phi_x^h}{1-\phi_x}\left(\hat{x}_h^i - \frac{c_x}{1-\phi_x}\right) \;+\; \sum_{j=h+1}^{12} x_j.$$

Next, we establish a relationship between the public and the private signals about yearly inflation. In the first line, we substitute the expression for the noisy signal $\widetilde{x}_h^i = x_h + \zeta_h^i$. In the second line, we substitute the expression for the (demeaned) actual monthly inflation $x_h - c_x/(1-\phi_x) = \phi_x(x_{h+1} - c_x/(1-\phi_x)) + \varepsilon_h^x$. Lastly, in the third line, we define the noise term

$\nu_h^i \equiv \frac{1-\phi_x^h}{1-\phi_x}\left(\varepsilon_h^x \;+\; \eta_h\right)$ and compute its noise.

$$
\begin{aligned}
\hat{\pi}_h^i &= h\left(\frac{c_x}{1-\phi_x}\right) + \frac{1-\phi_x^h}{1-\phi_x}\left(x_h - \frac{c_x}{1-\phi_x} + \zeta_h^i\right) + \sum_{j=h+1}^{12} x_j \\
&= h\left(\frac{c_x}{1-\phi_x}\right) + \frac{1-\phi_x^h}{1-\phi_x}\left(\phi_x\left(x_{h+1} - \frac{c_x}{1-\phi_x}\right) + \varepsilon_h^x + \zeta_h^i\right) + \sum_{j=h+1}^{12} x_j \\
&= z_h + \nu_h^i, \qquad \nu_\eta^i \sim \mathcal{N}\left(0, \left[\frac{1-\phi_x^h}{1-\phi_x}\right]^2 \alpha^2(\sigma_v^2 + \sigma_\zeta^2)\right).
\end{aligned}
\tag{C.15}
$$

C.4 Martingale property of beliefs

Beliefs follow a martingale: expected of future belief equal current beliefs, i.e., $\mathbb{E}[z_{h-1}^i|\mathcal{I}_h^i] = \hat{\pi}^i$. First, we use the relationship between public and private beliefs in (C.15) to set the expectation of individual noise ν to zero.

$$
\mathbb{E}[z_{h-1}^i|\mathcal{I}_h^i] = \mathbb{E}[z_{h-1} + \nu_{h-1}^i|\mathcal{I}_h^i] = \mathbb{E}[z_{h-1}|\mathcal{I}_h^i]. \tag{C.16}
$$

Next, we show that the expected public belief equals current public belief. Substituting in the expression for z_{h-1} in (21) and applying the expectation conditional on $\mathcal{I}_h^i$, we get:

$$
\mathbb{E}[z_{h-1}|\mathcal{I}_h^i] = (h-1)\frac{c_x}{1-\phi_x} + \frac{\phi_x(1-\phi_x^{h-1})}{1-\phi_x}\left(\widehat{x}_h^i - \frac{c_x}{1-\phi_x}\right) + \widehat{x}_h^i + \sum_{j=h+1}^{12} x_j.
$$

In the last sum, we separate $\widehat{x}_h^i \equiv \mathbb{E}[x_h|\mathcal{I}_h^i]$ that is not yet released from the rest of known values for $h = 12, ..., h+1$. Finally, we rearrange the expression to recover the expression for individual beliefs $\hat{\pi}^i$ plus three summands that cancel out:

$$
\mathbb{E}[z_{h-1}|\mathcal{I}_h^i] = \underbrace{h\frac{c_x}{1-\phi_x} + \frac{1-\phi_x^h}{1-\phi_x}\left(\widehat{x}_h^i - \frac{c_x}{1-\phi_x}\right) + \sum_{j=h+1}^{12} x_j}_{=\hat{\pi}^i} \; \underbrace{-\frac{c_x}{1-\phi_x} - \frac{1-\phi_x}{1-\phi_x}\left(\widehat{x}_h^i - \frac{c_x}{1-\phi_x}\right) + \widehat{x}_h^i}_{=0}.
$$

We conclude that $\mathbb{E}[z_{h-1}|\mathcal{I}_h^i] = \hat{\pi}^i$. As data on monthly inflation arrives, forecasters add the new observations to their dataset and update their estimates. Belief changes tend to be very persistent, even if the shocks that caused the beliefs to change are transitory. As a result, any changes in beliefs induced by new information are approximately permanent (Kozlowski, Veldkamp and Venkateswaran, 2020a,b).

C.5 Proof of Proposition 1

First, using the law of iterated expectations, we condition payoffs on the horizon-specific information sets:

$$
\mathbb{E}\left[\sum_{h=12}^{1} \mathbb{E}[(f_h^i - \pi)^2|\mathcal{I}_h^i] + r\,\mathbb{E}[(f_h^i - F_h)^2|\mathcal{I}_h^i] + \kappa\mathbb{1}_{\{f_h^i \neq f_{h+1}^i\}}\Big|\mathcal{I}_0^i\right]
$$

Second, we add and subtract beliefs $\hat{\pi}_h^i \equiv \mathbb{E}[\pi|\mathcal{I}_h^i]$ and $\hat{F}_h^i \equiv \mathbb{E}[F_h|\mathcal{I}_h^i]$ and open the squares:

$$\begin{aligned}
& \mathbb{E}\left[\sum_{h=12}^{1} \mathbb{E}[(f_h^i - \hat{\pi}_h^i + \hat{\pi}_h^i - \pi)^2|\mathcal{I}_h^i] \;+ r\; \mathbb{E}[(f_h^i - \hat{F}_h^i + \hat{F}_h^i - F_h)^2|\mathcal{I}_h^i] \;+\; \kappa \mathbb{1}_{\{f_h^i \neq f_{h+1}^i\}} \Big| \mathcal{I}_0^i\right] \\
= \;& \mathbb{E}\left[\sum_{h=12}^{1} \mathbb{E}[(f_h^i - \hat{\pi}_h^i)^2|\mathcal{I}_h^i] + \mathbb{E}[(\hat{\pi}_h^i - \pi)^2|\mathcal{I}_h^i] + 2\mathbb{E}[(f_h^i - \hat{\pi}_h^i)(\hat{\pi}_h^i - \pi)|\mathcal{I}_h^i] \Big| \mathcal{I}_0^i\right] \\
+ \;& r\mathbb{E}\left[\sum_{h=12}^{1} \mathbb{E}[(f_h^i - \hat{F}_h^i)^2|\mathcal{I}_h^i] + \mathbb{E}[(\hat{F}_h^i - F_h)^2|\mathcal{I}_h^i] \;+\; 2\mathbb{E}[(f_h^i - \hat{F}_h^i)(\hat{F}_h^i - F_h)|\mathcal{I}_h^i] \Big| \mathcal{I}_0^i\right] \\
+ \;& \kappa\mathbb{E}\left[\sum_{h=12}^{1} \mathbb{1}_{\{f_h^i \neq f_{h+1}^i\}} \Big| \mathcal{I}_0^i\right]
\end{aligned}$$

Third, we rewrite using conditional variances $\Sigma_h^\pi \equiv \mathbb{E}[(\hat{\pi}_h^i - \pi)^2|\mathcal{I}_h^i]$ and $\Sigma_h^F \equiv \mathbb{E}[(\hat{F}_h^i - F_h)^2|\mathcal{I}_h^i]$ and the fact that beliefs are unbiased $\mathbb{E}[(\hat{\pi}_h^i - \pi)|\mathcal{I}_h^i] = \mathbb{E}[(\hat{F}_h^i - F_h)|\mathcal{I}_h^i] = 0$:

$$\sum_{h=12}^{1} \mathbb{E}\left[\Sigma_h^\pi \;+\; r\Sigma_h^F \;+\; \mathbb{E}[(f_h^i - \hat{\pi}_h^i)^2|\mathcal{I}_h^i] \;+\; r\mathbb{E}[(f_h^i - \hat{F}_h^i)^2|\mathcal{I}_h^i] \;+\; \kappa \mathbb{1}_{\{f_h^i \neq f_{h+1}^i\}} \Big| \mathcal{I}_0^i\right].$$

C.6 Proof of Proposition 2

Given the stationarity of the problem and the stochastic processes, we apply the Principle of Optimality to the sequential problem and express it as a sequence of stopping-time problems Let τ be the stopping data associated with the optimal decision given the state $(\hat{\pi}_h^i, \hat{F}_h^i)$. The stopping time problem is given by:

C.7 Benchmark without fixed costs ($\kappa = 0$)

As a benchmark, we assume away the adjustment cost and set $\kappa = 0$. The problem becomes static, and the optimal forecast at horizon h minimizes per-period losses:

$$\min_{\{f_h^i\}} \; (f_h^i - \hat{\pi}_h^i)^2 \;+\; r(f_h^i - \hat{F}_h)^2 \tag{C.17}$$

The FOC yields:

$$f_h^i \;=\; \frac{1}{1+r}\left(\hat{\pi}_h^i + r\hat{F}_h\right). \tag{C.18}$$

Substituting the Bayesian beliefs yields:

$$f_h^i \;=\; z_h \;+\; r\left(c_F + \phi_F F_{h+1}\right) \;+\; \nu_h^i, \quad \text{with} \quad \nu_h^i \sim \mathcal{N}\left(0, \left[\frac{1-\phi_x^h}{1-\phi_x}\right]^2 \alpha^2(\sigma_x^2 + \sigma_\zeta^2)\right). \tag{C.19}$$

where α is the weight on private signals: $\alpha \equiv \frac{\sigma_\zeta^{-2}}{\sigma_x^{-2}) + \sigma_\zeta^{-2}}$.

MSE without fixed costs

$$\begin{aligned}
\mathbb{E}[e_h^i] = \mathbb{E}[(\pi - f_h^i)^2] &= \mathbb{E}[(\pi - \frac{1}{1+r}\hat{\pi}_h^i - \frac{r}{1+r}\hat{F}_h)^2] \\
&= \mathbb{E}\left[\left(\frac{1}{1+r}(\pi - \hat{\pi}_h^i) - \frac{r}{1+r}(\hat{F}_h - \pi)\right)^2\right] \\
&= \left(\frac{1}{1+r}\right)^2 \Sigma_h^\pi + \left(\frac{r}{1+r}\right)^2 \mathbb{E}\left[(\hat{F}_h - \pi)^2\right] + 2r\left(\frac{1}{1+r}\right)^2 \mathbb{E}\left[(\pi - \hat{\pi}_h^i)(\pi - \hat{F}_h)\right] \\
&= \left(\frac{1}{1+r}\right)^2 \Sigma_h^\pi + \left(\frac{r}{1+r}\right)^2 \mathbb{E}\left[(F_{h+1} - F_h)^2 + (F_h - \pi)^2\right] \\
&= \left(\frac{1}{1+r}\right)^2 \Sigma_h^\pi + \left(\frac{r}{1+r}\right)^2 \sigma_F^2 + \left(\frac{r}{1+r}\right)^2 \mathbb{E}\left[(F_h - \pi)^2\right] \\
&= \left(\frac{1}{1+r}\right)^2 \Sigma_h^\pi + \left(\frac{r}{1+r}\right)^2 \sigma_F^2 + \left(\frac{r}{1+r}\right)^2 \mathbb{E}\left[E_h^2\right]
\end{aligned}$$

D Computational strategy

Solving the problem requires computing expectations of future beliefs. Since all random variables are normal, this amounts to knowing the first two moments of these distributions. Next, we characterize these moments. Afterward, we use these moments to compute expectations.

D.1 Initial forecast

At the beginning of each year, we assume initial forecasts equal the 13-months ahead belief, which is optimal without frictions ($\kappa = r = 0$):

$$f^i_{13} \;=\; \hat{\pi}^i_{13} \;=\; z_{13} \;+\; \nu^i_{13}, \qquad \nu^i_{13} \sim \mathcal{N}(0, \sigma^2_{13}) \tag{D.20}$$

where z_{13} is constructed using the projection formula in (21)

$$z_{13} \;=\; 12\left(\frac{c_x}{1-\phi_x}\right) \;+\; \phi_x \frac{1-\phi_x^{12}}{1-\phi_x}\left(\hat{x}_{13} - \frac{c_x}{1-\phi_x}\right) \tag{D.21}$$

and the monthly belief equals $\hat{x}^i_{13} \;=\; \alpha[c_x + \phi_x x_{14}] + (1-\alpha)\tilde{x}^i_{13}$.

D.2 Distributions of expected beliefs

The law of motion of individual states implies the following values at $h-1$:

$$\hat{\pi}^i_{h-1} \;=\; \left(\frac{\Sigma_{z,h-1}}{\sigma^2_o + \Sigma_{z,h-1}}\right)\mu_o \;+\; \left(1 - \frac{\Sigma_{z,h-1}}{\sigma^2_o + \Sigma_{z,h-1}}\right)\hat{\pi}^i\hat{\pi}^i_{h-1} \tag{D.22}$$

$$\hat{F}_{h-1} \;=\; c_F \;+\; \phi_F F_h \tag{D.23}$$

Expected consensus beliefs The mean and variance of the distribution of expected consensus beliefs at $h-1$, from the perspective of horizon h (with knowledge up to F_{h+1}), are:

$$\mathbb{E}[\hat{F}_{h-1}|\mathcal{I}^i_h] \;=\; c_F + \phi_F \mathbb{E}[F_h|\mathcal{I}^i_h] \;=\; c_F(1+\phi_F) + \phi^2_F F_{h+1} \tag{D.24}$$

$$\mathbb{V}ar[\hat{F}^i_{h-1}|\mathcal{I}^i_h] \;=\; \phi^2_F \mathbb{V}ar[F_h|\mathcal{I}^i_h] \;=\; \phi^2_F \sigma^2_F \tag{D.25}$$

Expected inflation beliefs The mean and variance of the distribution of expected inflation beliefs at $h-1$, from the perspective of horizon h, are:

$$\mathbb{E}[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] \;=\; \left(\frac{\Sigma_{z,h-1}}{\sigma^2_o + \Sigma_{z,h-1}}\right)\mu_o \;+\; \left(1 - \frac{\Sigma_{z,h-1}}{\sigma^2_o + \Sigma_{z,h-1}}\right)\mathbb{E}[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] \tag{D.26}$$

$$\mathbb{V}ar[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] \;=\; \left(\frac{\sigma^2_o \Sigma_{z,h-1}}{\sigma^2_o + \Sigma_{z,h-1}}\right)^2 \mathbb{V}ar[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] \tag{D.27}$$

Now we compute the mean $\mathbb{E}[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h]$ and variance $\mathbb{V}ar[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h]$ of the idiosyncratic signal from the perspective of horizon h—inputs into the formulas above.

Expected signals We evaluate the formula for $\hat{\pi}^i_h$ in (??) at $h-1$, and separate the observation x_h from the sum yields:

$$\hat{\pi}^i_{h-1} \;=\; (h-1)\left(\frac{c_x}{1-\phi_x}\right) \;+\; \frac{1-\phi_x^{h-1}}{1-\phi_x}\left(\widetilde{x}^i_{h-1} - \frac{c_x}{1-\phi_x}\right) \;+\; x_h + \sum_{j=h+1}^{12} x_j. \tag{D.28}$$

Then, we take the expectation conditional on $\mathcal{I}^i_h$:

$$\mathbb{E}[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] = (h-1)\left(\frac{c_x}{1-\phi_x}\right) + \frac{1-\phi_x^{h-1}}{1-\phi_x}\left(\mathbb{E}[\widetilde{x}^i_{h-1}|\mathcal{I}^i_h] - \frac{c_x}{1-\phi_x}\right) + \mathbb{E}[x_h|\mathcal{I}^i_h] + \sum_{j=h+1}^{12} x_j \tag{D.29}$$

Next, we use the fact that $\mathbb{E}[\widetilde{x}^i_{h-1}|\mathcal{I}^i_h] = \mathbb{E}[x_{h-1}|\mathcal{I}^i_h]$ (because public and private noise have zero mean) and $\mathbb{E}[x_{h-1}|\mathcal{I}^i_h] = c_x + \phi_x \mathbb{E}[x_h|\mathcal{I}^i_h]$ (by the AR(1) assumption). Substituting into the previous expression:

$$\mathbb{E}[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] = (h-1)\left(\frac{c_x}{1-\phi_x}\right) + \frac{1-\phi_x^{h-1}}{1-\phi_x}\left(c_x + \phi_x\mathbb{E}[x_h|\mathcal{I}^i_h] - \frac{c_x}{1-\phi_x}\right) + \mathbb{E}[x_h|\mathcal{I}^i_h] + \sum_{j=h+1}^{12} x_j$$

Rearranging, we obtain:

$$\mathbb{E}[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] \;=\; h\left(\frac{c_x}{1-\phi_x}\right) \;+\; \phi_x\frac{1-\phi_x^{h-1}}{1-\phi_x}\left(\frac{\mathbb{E}[x_h|\mathcal{I}^i_h] - c_x}{1-\phi_x}\right) \;+\; \mathbb{E}[x_h|\mathcal{I}^i_h] - \frac{c_x}{1-\phi_x} + \sum_{j=h+1}^{12} x_j$$

Lastly, we substitute the AR(1) assumption $\mathbb{E}[x_h|\mathcal{I}^i_h] = c_x + \phi_x x_{h+1}$:

$$\mathbb{E}[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] \;=\; h\left(\frac{c_x}{1-\phi_x}\right) \;+\; \phi_x^2\frac{1-\phi_x^{h-1}}{(1-\phi_x)^2}x_{h+1} \;+\; \phi_x\left(x_{h+1} - \frac{c_x}{1-\phi_x}\right) + \sum_{j=h+1}^{12} x_j. \tag{D.30}$$

For the variance, we apply the variance operator to (D.28) and note that the terms in the sum disappear because they are known at h. Thus we are left with two terms.

$$\begin{aligned}
\mathbb{V}ar[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] &= \left(\frac{1-\phi_x^{h-1}}{1-\phi_x}\right)^2 \mathbb{V}ar[\widetilde{x}^i_{h-1}|\mathcal{I}^i_h] \;+\; \mathbb{V}ar[x_h|\mathcal{I}^i_h] \\
&= \left(\frac{1-\phi_x^{h-1}}{1-\phi_x}\right)^2 \left(\phi_x^2\mathbb{V}ar[x_h|\mathcal{I}^i_h] + \sigma_x^2 + \sigma_\zeta^2 + \sigma_\eta^2\right) \;+\; \mathbb{V}ar[x_h|\mathcal{I}^i_h] \\
&= \left(\frac{1-\phi_x^{h-1}}{1-\phi_x}\right)^2 \left(\phi_x^2\sigma_x^2 + \sigma_x^2 + \sigma_\zeta^2 + \sigma_\eta^2\right) \;+\; \sigma_x^2
\end{aligned}$$

where we use $\mathbb{V}ar[x_h|\mathcal{I}^i_h] = \sigma_x^2$ and the structure of the signal and the AR(1) assumption to write

$$\widetilde{x}^i_{h-1} \;=\; x^i_{h-1} + \zeta^i_{h-1} + \eta_{h-1} \;=\; c_x + \phi_x x_h + \varepsilon^x_{h-1} + \zeta^i_{h-1} + \eta_{h-1}. \tag{D.31}$$

D.3 Computing expectations

We approximate the expected continuation value of the value of action and inaction derived in Proposition 2 as follows

$$\mathbb{E}[\mathcal{V}_{h-1}(\hat{\pi}^i_{h-1}, \hat{F}_{h-1}, f^i_{h+1})|\mathcal{I}^i_h] = \sum_{\hat{\pi}^i_{h-1}} \sum_{\hat{F}_{h-1}} \mathcal{V}_{h-1}(\hat{\pi}^i_{h-1}, \hat{F}_{h-1}, f^i_{h+1})\omega(\hat{\pi}^i)\omega(\hat{F}) \tag{D.32}$$

where weights $\{\omega(\hat{\pi}^i), \omega(\hat{F})\}$ are constructed with Gaussian quadrature over grids for $\hat{\pi}^i$ and $\hat{F}$.

Integration weights $\omega_{\hat{F}}$ are such that $\hat{F}_{h-1}|\mathcal{I}^i_h \sim \mathcal{N}(\mathbb{E}[\hat{F}_{h-1}|\mathcal{I}^i_h], \mathbb{V}ar[\hat{F}^i_{h-1}|\mathcal{I}^i_h])$ with

$$\begin{aligned}
\mathbb{E}[\hat{F}_{h-1}|\mathcal{I}^i_h] &= c_F(1+\phi_F) + \phi_F^2 F_{h+1} \\
\mathbb{V}ar[\hat{F}^i_{h-1}|\mathcal{I}^i_h] &= \phi_F^2 \sigma_F^2
\end{aligned}$$

Integration weights $\omega_{\hat{\pi}^i}$ are such that $\hat{\pi}^i_{h-1}|\mathcal{I}^i_h \sim \mathcal{N}\left(\mathbb{E}[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h], \mathbb{V}ar[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h]\right)$, with

$$\begin{aligned}
\mathbb{E}[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] &= \frac{\Sigma_{z,h-1}}{\sigma_o^2 + \Sigma_{z,h-1}}\mu_o + \frac{\sigma_o^2}{\sigma_o^2 + \Sigma_{z,h-1}}\left[h\left(\frac{c_x}{1-\phi_x}\right) + \phi_x^2 \frac{1-\phi_x^{h-1}}{(1-\phi_x)^2} x_{h+1} + \phi_x\left(x_{h+1} - \frac{c_x}{1-\phi_x}\right) + \sum_{j=h}^{1} \right. \\
\mathbb{V}ar[\hat{\pi}^i_{h-1}|\mathcal{I}^i_h] &= \left(\frac{\sigma_o^2 \Sigma_{z,h-1}}{\sigma_o^2 + \Sigma_{z,h-1}}\right)^2 \left[\sigma_x^2 + \left(\frac{1-\phi_x^{h-1}}{1-\phi_x}\right)^2 \left(\phi_x^2\sigma_x^2 + \sigma_x^2 + \sigma_\zeta^2 + \sigma_\eta^2\right)\right]
\end{aligned}$$

D.4 Consistency of consensus process

The perceived law of motion for consensus is $F_h = F_{h+1} + \varepsilon^F_h$ with $\varepsilon^F_h \sim \mathcal{N}(0, 0.11^2)$. The perceived process is

$$\hat{F}_t = \hat{F}_{t-1} + \varepsilon^{\hat{F}}_t, \qquad \varepsilon^F_h \sim \mathcal{N}(0, 0.11^2) \tag{D.33}$$

The actual law of motion is

$$F_h = -0.03 + 1.01 F_{h+1} + \varepsilon^F_h, \qquad \varepsilon^F_h \sim \mathcal{N}(0, 0.11^2). \tag{D.34}$$

Figure D.8 – Consensus belief consistency

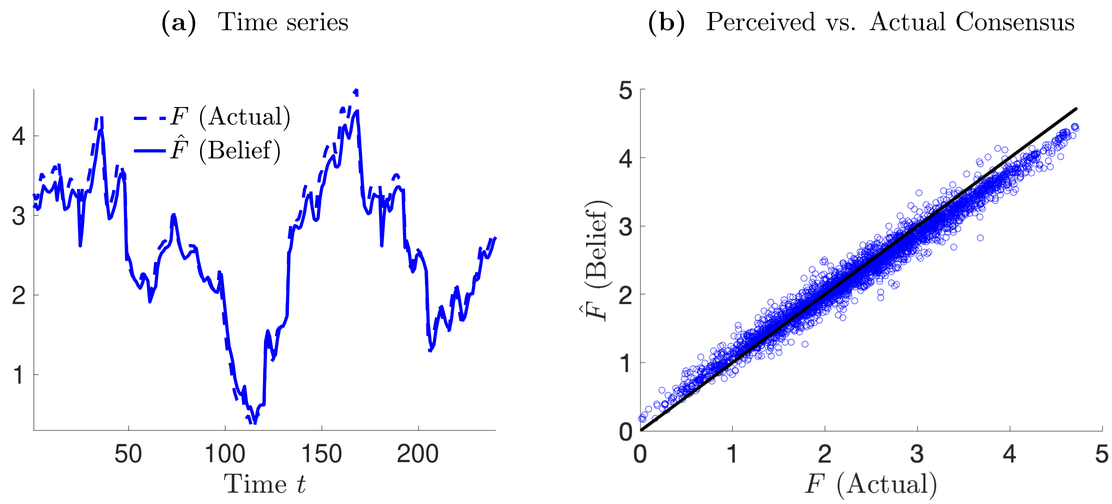

Notes: Model, benchmark calibration.

E Rationality tests

E.1 A history of rationality tests

Underlying the Rationality tests is the hypothesis that forecast errors can not be predicted using the information available when each forecast is made. Building on these implications, Coibion and Gorodnichenko (2015) proposed to test if forecast errors are predictable using forecast revisions. The underlying assumption is that forecast revisions $f_h^i - f_{h+1}^i$ about the same targeted variable (as in our setup) constitute an accurate reflection of the reaction to recently available news. The original test uses the consensus forecast made at horizon h about yearly inflation. In terms of our notation and the definition of the consensus forecast, the rationality test

$$\pi - F_h = \gamma_0 + \gamma_1(F_h - F_{h+1}) + \eta_h.$$

Under FIRE, forecasters entail rational expectations (RE) and share the same full information set (FI). Therefore, under FIRE, $\gamma_1 = 0$, i.e., forecast errors are not predictable from revisions. Bordalo *et al.* (2020) challenges this interpretation and argues that while $\gamma_1 \neq 0$ rejects FIRE, it does not necessarily challenges RE. In fact, at the individual level, although a forecaster observes noisy signals about inflation, she still can rationally use them to update their forecasts. Bordalo *et al.* (2020) argues that a rationality test where unpredictable forecast errors must be performed at the individual level. They proposed the following test:

$$\pi - f_h^i \ = \ \gamma_0 \ + \ \gamma_1(f_h^i - f_{h+1}^i) \ + \ \eta_h^i, \qquad \mathbb{E}[\eta_h^i] = 0. \tag{E.35}$$

With the individual information set assumed in our model $\mathcal{I}_h^i = \widetilde{x}_h^i \cup \mathcal{I}_h = \widetilde{x}_h^i \cup \{x_{h+1}, x_{h+2}, \ldots, F_{h+1}, F_{h+2}, \ldots\}$. Thus, the new information available between consecutive periods h and $h+1$ is $(\widetilde{x}_h^i, x_{h+1}, F_{h+1})$. Combining these two last bits of information and using the assumed known structure for the variables, each forecaster can back out $\epsilon_{h+1}^x, \zeta_{h+1}^i, \epsilon_{h+1}^F$. Moreover, certainty about the lag shock for monthly inflation and noise realization for the signal allows the forecasters to back out ν^i_{h+1}.

E.2 Rationality tests in the literature

Table E.4 – Rationality Tests in the Literature

	Coibion-Gorodnichencko (2015)			Bordalo et.al (2020)			Broer and Kohlhas (2022)			
Horizon	Constant	Revision	F-test	Constant	Revision	F-test	Constant	Revision	Consensus	F-test
1	-0.0341	5.7198	0.0720	-0.0657	-0.5379	0.0008	-0.0651	-0.2359	0.0026	0.0008
	(.0418)	(2.7197)		(.0003)	(.1092)		(.0220)	(.1101)	(.0133)	
2	-0.1001	7.1433	0.0004	-0.0665	-0.3284	0.0080	-0.0648	-0.1006	0.0011	0.0184
	(.0233)	(1.0654)		(.0008)	(.0969)		(.0227)	(.2296)	(.0136)	
3	-0.0318	5.9955	0.0007	-0.0657	-0.5048	0.0001	-0.0555	-0.2895	0.0001	0.0078
	(.0426)	(1.0723)		(.0002)	(.0729)		(.0236)	(.1396)	(.01415)	
4	-0.0172	7.9209	0.2830	-0.0707	-0.3527	0.0022	-0.0678	-0.1225	0.0076	0.0112
	(.0775)	(5.3897)		(.00001)	(.0833)		(.02543)	(.1012)	(.0153)	
5	-0.1020	-0.5235	0.3723	-0.0820	-0.3501	0.0030	-0.1044	-0.2546	0.0221	0.0000
	(.1352)	(2.1798)		(.0001)	(.0867)		(.0286)	(.0674)	(.0169)	
6	-0.0812	1.4816	0.0458	-0.1054	-0.4286	0.0012	-0.1676	-0.2687	0.0440	0.0000
	(.1394)	(2.2645)		(.00112)	(.0916)		(.0285)	(.1175)	(.01667)	
7	-0.0441	2.9080	0.0023	-0.1363	-0.5388	0.0013	-0.2014	-0.3205	0.0476	0.0000
	(.1261)	(1.2488)		(.0037)	(.1172)		(.0273)	(.1273)	(.0159)	
8	-0.1902	3.4327	0.0050	-0.1471	-0.5098	0.0020	-0.2788	0.0075	0.0783	0.0000
	(.0622)	(1.2121)		(.0006)	(.1186)		(.0282	(.0912)	(.0161)	
9	-0.2616	5.8216	0.0019	-0.1326	-0.3888	0.0396	-0.3447	0.0146	0.1217	0.0000
	(.0736)	(1.0557)		(.0042)	(.1616)		(.0328)	(.12729	(.01879)	
10	-0.2470	4.3845	0.0059	-0.1031	-0.4776	0.0011	-0.4016	0.2655	0.1671	0.0000
	(.0695)	(1.0055)		(.0034)	(.1016)		(.0433)	(.0805)	(.0247)	
11	-0.0247	2.2792	0.0236	-0.1018	-0.5368	0.0000	-0.5377	-0.1216	0.2778	0.0000
	(.1140)	(1.0906)		(.00007)	(.0619)		(.0507)	(.1007)	(.0293)	
12	0.1126	2.5178	0.0010	-0.1590	-0.8157	0.0016	-0.4858	0.1883	0.2381	0.0000
	(.2179)	(.9034)		(.0142)	(.1737)		(.0717)	(.1163)	(.0363)	

Notes: The first columns of the table reports the estimates of running Coibion and Gorodnichenko (2015)'s aggregate Rationality Test. Standard are estimated using a HAC matrix. The middle columns present the results of Bordalo *et al.* (2020)'s individual test. Finally, the last columns shows the estimates of Broer and Kohlhas (2022). The panel estimations include individual fixed-effects. Standard errors are robust and clustered by forecaster and time. All the tests are estimated using the "Normal years" sample (2010-2019).

E.3 Rationality tests for beliefs ($\kappa = r = 0$)

Consider the case $\kappa = r = 0$. In this case, forecasts equal agent's beliefs $f_h^i = \widehat{\pi}_h^i$. Denote the average belief in horizon h as $\hat{\Pi}_h = \sum_{i=1}^N \hat{\pi}_h^i$.

For beliefs, the rationality regression reads:

$$\pi - \hat{\pi}_{h+1}^i \;=\; \gamma_0 \;+\; \gamma_1(\hat{\pi}_h^i - \hat{\pi}_{h+1}^i) \;+\; \gamma_2(\hat{\Pi}_h^i - \hat{\pi}_{h+1}^i) \;+\; \eta_h^i, \qquad \mathbb{E}[\eta_h^i] = 0. \tag{E.36}$$

First, we note that the last term equals:

$$\hat{\Pi}_h - \hat{\pi}_{h+1}^i \quad = \quad \hat{\Pi}_h - \hat{\pi}_h^i + \hat{\pi}_h^i - \hat{\pi}_{h+1}^i \;=\; \nu_h^i + \hat{\pi}_h^i - \hat{\pi}_{h+1}^i$$

Thus, we can rewrite it as:

$$\pi - \hat{\pi}_{h+1}^i \;=\; \gamma_0 \;+\; \tilde{\gamma}(\hat{\pi}_h^i - \hat{\pi}_{h+1}^i) \;+\; \tilde{\eta}_h^i, \qquad \mathbb{E}[\tilde{\eta}_h^i] = 0. \tag{E.37}$$

where $\tilde{\gamma} \equiv (\gamma_1 + \gamma_2)$ and $\tilde{\eta}_h^i \equiv \gamma_2 \nu_h^i + \eta_h^i$. This equation will pin down the expected value for $\tilde{\gamma}$. As a preliminary, note that the difference (between consecutive horizons) in the z_h's equals to:

$$\begin{aligned}
& z_h - z_{h+1} \\
= \;& (h-(h+1))\left(\frac{c_x}{1-\phi_x}\right) + \frac{\phi_x(1-\phi_x^h)}{1-\phi_x}\left(x_{h+1} - \frac{c_x}{1-\phi_x}\right) - \frac{\phi_x(1-\phi_x^{h+1})}{1-\phi_x}\left(x_{h+2} - \frac{c_x}{1-\phi_x}\right) + \sum_{j=h+1}^{12} x_j - \\
= \;& -\left(\frac{c_x}{1-\phi_x}\right) + \frac{\phi_x(1-\phi_x^h)}{1-\phi_x}\left(x_{h+1} - \frac{c_x}{1-\phi_x}\right) - \frac{\phi_x(1-\phi_x^{h+1})}{1-\phi_x}\left(x_{h+2} - \frac{c_x}{1-\phi_x}\right) + \frac{1-\phi_x}{1-\phi_x} x_{h+1} \\
= \;& -(1-\phi^{h+1})\left(\frac{c_x}{1-\phi_x}\right) \;+\; \frac{(1-\phi_x^{h+1})}{1-\phi_x}\left(x_{h+1} - \phi_x x_{h+2}\right) \\
= \;& -\frac{1-\phi^{h+1}}{1-\phi_x} c_x \;+\; \frac{1-\phi_x^{h+1}}{1-\phi_x}(c_x + \epsilon_{h+1}^x) \\
= \;& \frac{1-\phi_x^{h+1}}{1-\phi_x} \epsilon_{h+1}^x
\end{aligned}$$

Let us provide an expression for belief revisions in terms of shocks:

$$\begin{aligned}
\hat{\pi}_h^i - \hat{\pi}_{h+1}^i \quad &= \quad \left[\, (z_h - z_{h+1}) + \; (\nu_h^i - \nu_{h+1}^i)\right] \\
&= \quad \frac{1-\phi_x^{h+1}}{1-\phi_x}\epsilon_{h+1}^x + \left[\frac{1-\phi^h}{1-\phi}\alpha(\epsilon_h^x + \zeta_h^i) - \frac{1-\phi^{h+1}}{1-\phi}\alpha(\epsilon_{h+1}^x + \zeta_{h+1}^i)\right] \\
&= \quad \underbrace{\frac{1-\phi_x^{h+1}}{1-\phi_x}[(1-\alpha)\epsilon_{h+1}^x + \alpha\zeta_{h+1}^i]}_{\text{A}} + \underbrace{\frac{1-\phi_x^h}{1-\phi_x}\alpha(\epsilon_h^x + \zeta_h^i)}_{\text{B}}
\end{aligned}$$

For the first term, we used the previous result. Moreover, the second component in the second row comes from the definition of ν_h^i at any horizon.

On the other hand, the forecast error $\pi - \widehat{\pi}_h^i$ can be written in terms of shocks:

$$\pi - \widehat{\pi}_h^i = \underbrace{\frac{1-\phi_x^h}{1-\phi_x}((1-\alpha)\varepsilon_h^x - \alpha\zeta_h^i)}_{\text{C}} + \underbrace{\sum_{j=1}^{h-1}\frac{1-\phi_x^j}{1-\phi_x}\varepsilon_j^x}_{\text{D}}$$

The first term (C) accounts for the forecast error due to the uncertainty about the actual value of the current monthly inflation x_h. We interpreted this part as the weighted error between the AR(1) and the private signals, ultimately projected towards the end of the year. The second component (D) accounts for the unforecastable part of the process over the remaining horizons $h-1, \ldots, 1$.

The estimated coefficient $\tilde{\gamma}$ equals:

$$\begin{aligned}
\tilde{\gamma} &= \frac{\mathbb{C}ov((f_h^i - f_{h+1}^i), \pi - \widehat{\pi}_h^i)}{Var(f_h^i - f_{h+1}^i)} \\
&= \frac{\mathbb{C}ov(A,C) + \mathbb{C}ov(A,D) + \mathbb{C}ov(B,C) + \mathbb{C}ov(B,D)}{Var[\hat{\pi}_h^i - \hat{\pi}_{h+1}^i]} \\
&= 0
\end{aligned}$$

We can show that in this case, all the covariance terms in the numerator are equal to zero. In fact, $\mathbb{C}ov(A,C) = \mathbb{C}ov(A,D) = \mathbb{C}ov(B,D) = 0$. All these covariance terms depend on i.i.d. shocks realized at different times. Moreover, the term $\mathbb{C}ov(B,C)$ is equal to:

$$\begin{aligned}
\mathbb{C}ov(B,C) &= \mathbb{C}ov\left(\frac{1-\phi_x^h}{1-\phi_x}\alpha(\epsilon_h^x + \zeta_h^i), \frac{1-\phi_x^h}{1-\phi_x}((1-\alpha)\varepsilon_h^x - \alpha\zeta_h^i)\right) \\
&= \left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 \alpha(1-\alpha)\mathbb{C}ov(\varepsilon_h^x + \zeta_h^i) - \left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 \alpha^2\mathbb{C}ov(\varepsilon_h^x + \zeta_h^i, \zeta_h^i) \\
&= \left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 \alpha(1-\alpha)\sigma_x^2 - \left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 \alpha^2\sigma_\zeta^2 \\
&= \left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 \alpha[(1-\alpha)\sigma_x^2 - \alpha\sigma_\zeta^2] \\
&= \left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 \alpha\left[\left(\frac{\sigma_x^{-2}}{\sigma_x^{-2}+\sigma_\zeta^{-2}}\right)\sigma_x^2 - \left(\frac{\sigma_\zeta^{-2}}{\sigma_x^{-2}+\sigma_\zeta^{-2}}\right)\sigma_\zeta^2\right] \\
&= 0
\end{aligned}$$

Therefore, in the Belief case, we expect $\tilde{\gamma} = \gamma_1 + \gamma_2 = 0$, which is consistent with the idea that agents have rational expectations. Given the process for the error, $\gamma_2 = 0$. Finally, $\gamma_0 = \frac{\sum_{i=0}^{N}\pi - f_h^i}{N}$, for each horizon. As forecasters are indeed rational, their forecast error must remain unpredictable, i.e., $\gamma_0 = 0$

E.4 Rationality tests without fixed costs ($\kappa = 0$)

From the FOC in the frictionless case:

$$f_h^i = \frac{1}{1+r}(z_h + r\hat{F}_h + \nu_h^i), \quad \text{with} \quad \nu_h^i \sim \mathcal{N}\left(0, \left[\frac{1-\phi_x^h}{1-\phi_x}\right]^2 \alpha^2(\sigma_x^2 + \sigma_\zeta^2)\right). \tag{E.38}$$

Individual forecast revisions across consecutive horizons can be unpacked into three different news terms:

$$f_h^i - f_{t+1}^i = \frac{1}{1+r}\left[\underbrace{(z_h - z_{h+1})}_{\text{I}} + r\underbrace{\left(\hat{F}_h - \hat{F}_{h+1}\right)}_{\text{II}} + \underbrace{(\nu_h^i - \nu_{h+1}^i)}_{\text{III}}\right] \tag{E.39}$$

- News I, the difference in the z_h's, equals to:

$$z_h - z_{h+1} = \frac{1-\phi_x^{h+1}}{1-\phi_x}\epsilon_{h+1}^x$$

- News II, the difference in the $\hat{F}_h$'s, equals to:

$$\begin{aligned} \hat{F}_h - \hat{F}_{h+1} &= c_F + \phi_F F_{h+1} - c_F + \phi_F F_{h+2} \\ &= \phi_F(F_{h+1} - F_{h+2}) \end{aligned}$$

- News III, the difference in ν_h^i, equals to:

$$\nu_h^i - \nu_{h+1}^i = \widehat{\pi}_h^i - \widehat{\pi}_{h+1}^i - (z_h - z_{h+1})$$

Similarly, in this case the forecast error $\pi - f_h^i$, can be written as:

$$\pi - f_h^i = \frac{1}{1+r}\left[\underbrace{(\pi - \widehat{\pi}_h^i)}_{\text{I*}} + \underbrace{r\pi}_{\text{II*}} - \underbrace{r\widehat{F}_h}_{\text{III*}}\right] \tag{E.40}$$

- Where I*, the difference $\pi - \widehat{\pi}_h^i$, equals to:

$$\pi - \widehat{\pi}_h^i = \frac{1-\phi_x^h}{1-\phi_x}((1-\alpha)\varepsilon_h^x - \alpha\zeta_{ih}) + \sum_{j=1}^{h-1}\frac{1-\phi_x^j}{1-\phi_x}\varepsilon_j^x$$

  As shown in equation (C.10).

- II*, equals to:

$$r\pi = r\sum_{h=1}^{12} x_h$$

- III*, equals to:

$$r\widehat{F}_h = r(c_f + \phi_F F_{h+1})$$

Hence, in this case, the regression coefficients are:

$$\gamma_1 = \frac{\mathbb{C}ov(f_h^i - f_{h+1}^i, \pi - f_h^i)}{\mathbb{V}ar(f_h^i - f_{h+1}^i)} \tag{E.41}$$

where the numerator $\mathbb{C}ov(f_h^i - f_{h+1}^i, \pi - f_h^i) = \left(\frac{1}{1+r}\right)^2 \mathbb{C}ov(\mathrm{I} + \mathrm{II} + \mathrm{III}, \mathrm{I}^* + \mathrm{II}^* + \mathrm{III}^*)$,

E.5 Rationality tests using forecast revisions

Notation In what follows, we incorporate the subindices $s = short$, $m = medium$, and $l = long$ to describe the forecasts horizons when the forecasts were made. Then $s < m < l$. All the tests below rely on average moments across the cross-section of forecasters. All testable implications exploit the variability across targeted years t, conditioning on the forecast horizons. Let the consensus forecast be denoted by $F_{t,h} = N^{-1} \sum_{i=1}^{N} f_{t,h}^{i}$. To fix the notation, let the consensus error be $E_{t,h} = \pi_t - F_{t,h}$; the mean-squared consensus error be $\text{MSE}_h \equiv \mathbb{E}[(\pi_t - F_{t,h})^2]$, the average consensus revision between consecutive periods as $\Delta F_{t,h|h+1} \equiv F_{t,h} - F_{t,h+1}$, and the mean-squared consensus revision be $\text{MSR}_h \equiv \mathbb{E}[(F_{t,1} - F_{t,h})^2]$

E.5.1 Multivariable inequality tests of rationality

The first seven conditions imply methods for handling multivariate inequality tests.

1. **Weakly increasing mean-squared consensus errors**

   (E.42) $$\text{MSE}_s \ \leq \ \text{MSE}_m \ \leq \ \text{MSE}_l, \qquad \forall s < m < l$$

   Forecast accuracy increases as the horizon shrinks.

2. **Weakly increasing variance of consensus revisions**

   (E.43) $$\mathbb{V}ar[\Delta F_{t,1|2}] \leq \mathbb{V}ar[\Delta F_{t,1|3}] \leq \cdots \leq \mathbb{V}ar[\Delta F_{t,1|H}].$$

   The variance of mean squared revisions must increase in the forecast horizon.

   Proof. By definition $\Delta F_{t,1|H} = F_{t,1} - F_{t,H} = \sum_{j=1}^{H-1} \Delta F_{t,j|j+1}$ for $H > 1$. Take variance on both sides $\mathbb{V}ar[\Delta F_{t,1|H}] = \sum_{j=1}^{H-1} \mathbb{V}ar[\Delta F_{t,j|j+1}]$, then we get the result.

3. **Weakly decreasing consensus variance**

   (E.44) $$\mathbb{V}ar[F_{t,s}] \ \geq \ \mathbb{V}ar[F_{t,m}] \geq \ \mathbb{V}ar[F_{t,l}], \qquad \forall s < m < l$$

   Proof. By definition of consensus error: $\mathbb{V}ar(\pi_t) = \mathbb{V}ar(F_{t,h}) + \mathbb{V}ar(E_{t,h}) + 2\mathbb{C}ov(F_{t,h}, E_{t,h})$. By rationality, the covariance term is zero and thus $\mathbb{V}ar(F_{t,h}) = \mathbb{V}ar(\pi_t) - \mathbb{V}ar(E_{t,h})$. Since mean squared errors must be weakly increasing, then forecasts must be weakly decreasing as well. Note that unlike tests using forecast errors, this test does not rely on data for the targeted variable.

4. **Weakly decreasing covariance between consensus and target**

   (E.45) $$\mathbb{C}ov(F_{t,s}, \pi_t) \ \geq \ \mathbb{C}ov(F_{t,l}, \pi_t).$$

   Proof. The covariance between the average prediction and the target is $\mathbb{C}ov(F_{t,h}, \pi_t) = \mathbb{C}ov(F_{t,h}, F_{t,h} + E_{t,h})$. By rationality, the covariance term is zero and thus $\mathbb{C}ov(F_{t,h}, \pi_t) = \mathbb{V}ar(F_{t,h})$. Since the variance of predictions decreases as the forecast horizon increases, the covariance does, too.

5. **Weakly decreasing covariance between consensus and proxy**

$$\mathbb{C}ov(F_{t,m}, F_{t,s}) \geq \mathbb{C}ov(F_{t,l}, F_{t,s}). \tag{E.46}$$

Proof. This test relies on the previous one but builds on the fact that shorter predictions must better reflect the targeted variable. Then, short-term predictions are used as a proxy for π_t. In particular, as we know that $\mathbb{C}ov(F_{t,s}, \pi_t) \geq \mathbb{C}ov(F_{t,l}, \pi_t)$ it must hold that $\mathbb{C}ov(F_{t,m}, F_{t,s}) \geq \mathbb{C}ov(F_{t,l}, F_{t,s})$. Again, this test does not rely on having information about the target.

6. **Bounded variance of consensus revisions**

$$\mathbb{V}ar(\Delta F_{t,s|l}) \ \leq \ 2\mathbb{C}ov(\pi_t, \Delta F_{t,s|l}). \tag{E.47}$$

Proof. Under rationality $\mathbb{V}ar(E_{t,l}) \geq \mathbb{V}ar(E_{t,s}) \Rightarrow \mathbb{V}ar(\pi_t - F_{t,l}) \geq \mathbb{V}ar(\pi_t - F_{t,s})$. We can express the short-horizon average forecast as $F_{t,s} = F_{t,l} + F_{t,s} - F_{t,l} = F_{t,l} + \Delta F_{t,s|l}$. The left-hand side of the variance of forecast errors inequality is $\mathbb{V}ar(\pi_t - F_{t,l}) = \mathbb{V}ar(\pi_t) + \mathbb{V}ar(F_{t,l}) - 2\mathbb{C}ov(\pi_t, F_{t,l})$. Likewise, the right hand side of the inequality $\mathbb{V}ar(\pi_t - F_{t,s}) = \mathbb{V}ar(\pi_t) + \mathbb{V}ar(F_{t,l}) + \mathbb{V}ar(\Delta F_{t,s|l}) - 2\mathbb{C}ov(\pi_t, F_{t,l}) - 2\mathbb{C}ov(\pi_t, \Delta F_{t,s|l})$, where we use the expression for $F_{t,s}$ along with the fact that $\mathbb{C}ov(F_{t,l}, \Delta F_{t,s|l}) = 0$. Therefore, by removing terms we show that the variance of revisions is bounded, In other words, the variability of forecast revisions is limited by the covariance of revisions with the target variable.

7. **Bounded variance of forecast revisions with proxy**

As a result of the previous implications, it is possible to express the bound using the short-term forecast for π_t under rationality:

$$\mathbb{V}ar(\Delta F_{t,s|l}) \ \leq \ 2\mathbb{C}ov(F_{t,s}, \Delta F_{t,s|l})$$

E.5.2 Regression-based rationality tests

The three remaining regression-based tests build on Mincer and Zarnowitz (1969) rationality regressions. The original MZ test regresses the actual target variable on forecasts, $\pi_t = \alpha_h + \beta_h F_{t,h} + v_{t,h}$, and tests the null that the intercept is zero and the slope is one for each horizon h: $H_0^h : \alpha_h = 0 \cap \beta_h = 1$. The following three tests re-express this test using forecast revisions or eliminating the target variable.

8. **MZ Revision test (Univar opt revision)** Express the MZ test using forecast revisions:

$$\pi_t \ = \ \alpha + \beta_H F_{t,h=H} + \sum_{j=2}^{H-1} \beta_j \Delta F_{t,j-1|j} + v_t$$

and jointly testing $H_0 : \alpha = 0 \cap (\beta_1 = \cdots = \beta_H = 1)$.

9. **MZ Revision test with proxy (Univar opt revision with proxy)** Building on the previous specification, substitute the target variable with its short-term expectation ($\pi_t \approx F_{t,1}$):

$$F_{t,1} = \alpha + \beta_H F_{t,h=H} + \sum_{j=2}^{H-1} \beta_j \Delta F_{t,j-1|j} + v_t$$

and jointly testing $H_0 : \alpha = 0 \cap (\beta_1 = \cdots = \beta_H = 1)$.

10. **MZ Regression-based (Univar MZ short h)** Finally, running a MZ regression-based test without relying on the target variable

$$F_{t,1} = \alpha_h + \beta_h F_{t,h} + v_{t,h}$$

and testing $H_0^h : \alpha_h = 0 \cap \beta_h = 1$ for each h.

Figure F.9 – Cross-sectional moments across model configurations

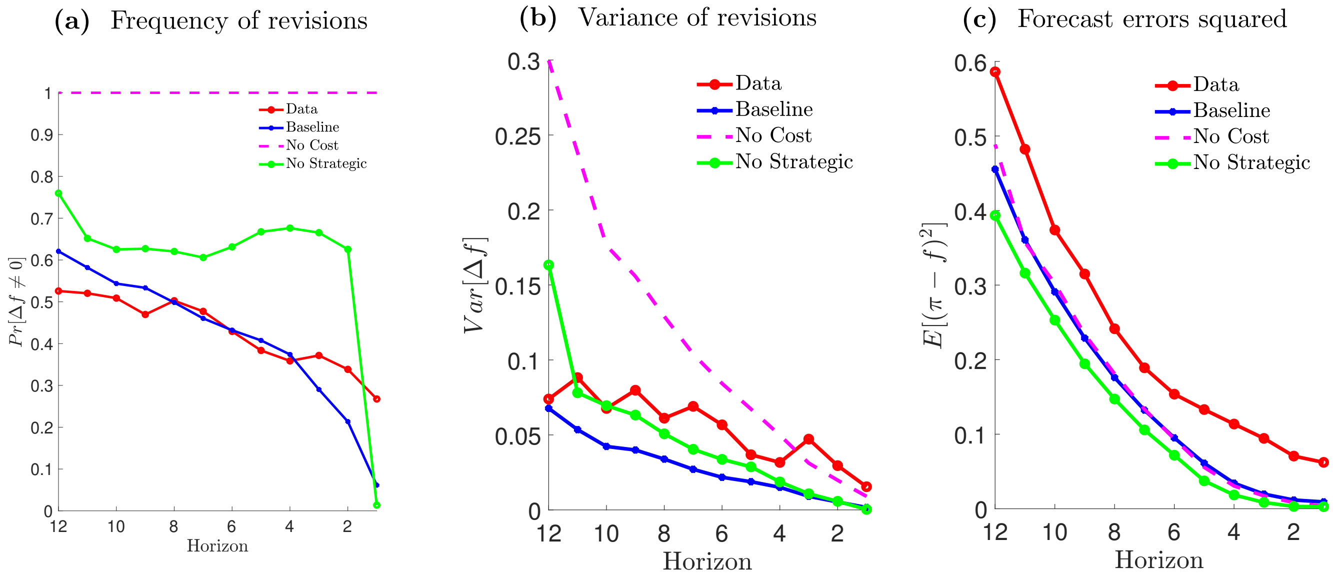

Notes: Bloomberg data for normal years = 2010-2019. Benchmark calibration uses parameters from Table III: $\kappa = 0.083, r = 0.263, \sigma_\zeta = 0.098$. No fixed costs: sets $\kappa = 0$ and re-estimates parameters. No strategic concerns: sets $r = 0$ and re-estimates parameters.

F On the role of fixed revision costs and strategic concerns

We explore two alternative model versions of the model to assess the role of fixed costs κ and strategic concerns r. In each panel of Figure F.9, we plot four lines: data (red), benchmark (blue), zero fixed costs $\kappa = 0$ (dashed pink), and zero strategic concerns $r = 0$ (dotted blue). In each alternative, we re-estimate the model's parameters to fit a subset of the target moments.

As expected, the model with zero fixed costs implies a frequency of revisions close to one for all horizons and thus fails dramatically in replicating the observed empirical patterns. The model with zero strategic concerns delivers steeper profiles in frequency and variance. Notably, all model configurations deliver very similar patterns for the square of forecast errors.

Table F.5 – Parameterization of alternative models

Parameters	Data	(1) Baseline	(2) Zero fixed revision costs	(3) No strategic concerns	(4) No signal noise
κ		0.11	0.00	0.07	
r		0.51	0.43	0.00	
σ_ζ^2		1.50	1.48	1.70	0.00
Moments					
$\Pr[\Delta f \neq 0]$	0.43	0.43*	1.00	0.42*	
$\mathbb{V}ar[\Delta f]$	0.05	0.05*	0.05*	0.04	
$\mathbb{E}[(\pi - f)^2]$	0.23	0.21*	0.18*	0.17*	

Notes: Bloomberg data for normal years = 2010-2019. Baseline uses parameters from Table III. Alternative models: No fixed costs ($\kappa = 0$), No strategic concerns ($r = 0$), No idiosyncratic noise ($\sigma_\zeta = 0$). Targeted moments in each estimation are marked with *.

G Accuracy of AR(1) forecast vs. consensus

Since yearly inflation π equals to the sum of year-on-year monthly inflation x_h: $\pi = \sum_{h=1}^{12} x_h$, then monthly inflation releases x_h are relevant information strictly related to the forecasted variable. We assume forecasters believe monthly inflation follows an AR(1) process:[13] As noticed by $\mathcal{I}_h$, the public belief is formed using the available information at each horizon, corresponding to the lagged values of x_h. By relying on the available information for forecasters in real time, we estimate the AR(1) process parameters using a rolling window over the sample years. The parameters are $(c_x, \phi_x, \sigma_x^2) = (0.013, 0.932, 0.036)$. The estimates are presented in Appendix B. Given these estimates, we compute our proxy for public beliefs each month.

Given the estimates, we compare the forecast errors from the AR(1) project and the average prediction of forecasters (consensus). We compute the evolution of forecast errors over the horizon using these two forecasts, where

$$e_h = \pi - f_h, \qquad \text{where} \quad f_h \in \{z_h, F_h\}. \tag{G.48}$$

The results are presented in Figure G.10.

Figure G.10 – Accuracy of AR(1) proxy and Consensus Forecasts

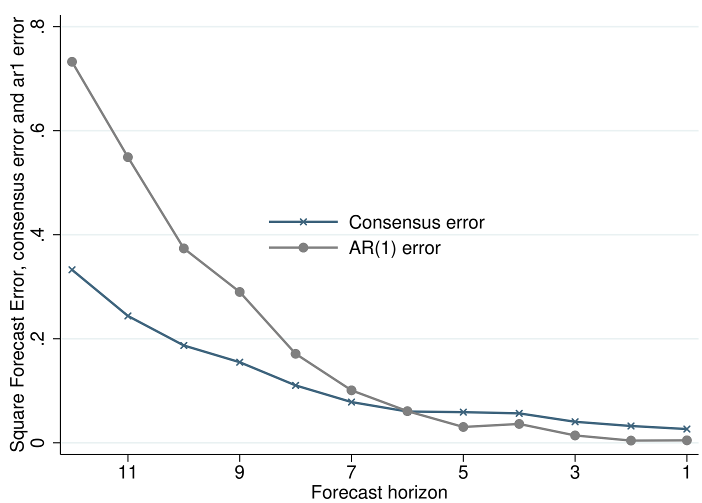

As expected, the accuracy of the AR(1) process increases monotonically as the release date of inflation approaches. This feature is particularly salient for normal years. As more relevant information is accumulated, the accuracy improvements are notorious. If we focus on normal years at longer and medium horizons, the forecast error of the consensus forecast is relatively better than the AR(1). This could be caused by the possible combination of public and private information from forecasters that makes predictions more accurate. However, at shorter horizons, the accuracy of the public belief is slightly better than the average prediction. Interestingly, the consensus is consistently more accurate during turbulent years.

[13] Although participants interpret public information differently, we argue that the prediction that builds on the AR(1) is a tractable and accurate proxy for a fixed-event forecast. We provide further discussion about the features and accuracy of the proxy in Appendix YY. See Giacomini, Skreta and Turen (2020).