

Lumpy Forecasts^{*}

Isaac Baley[†]

UPF, CREI, BSE and CEPR

Javier Turen[‡]

Pontificia Universidad Católica de Chile

December 23, 2024

Abstract

Professional forecasters adjust their inflation forecasts in a distinctly lumpy pattern, making infrequent but substantial revisions. Strategic concerns play a significant role—forecasters are more likely to adjust, and by larger amounts, when their forecasts deviate from the consensus. Using a fixed-event forecasting framework, we document the impact of lumpiness and consensus pressure on forecast adjustments. Our quantitative model, which integrates Bayesian belief updating with forecast revision costs and strategic concerns, not only replicates the observed lumpiness in survey data but also sheds light on forecasters’ apparent overreactions to new information. This structured framework enables us to “cleanse” forecasts, isolating the underlying inflation beliefs that drive these forecasts.

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

Keywords: forecasting, survey of professional forecasters, overreaction, consensus, forecast revision costs, forecast stability, strategic concerns, Bayesian learning, fixed-event forecasts

^{*}We largely benefited from conversations with João Victor Issler (discussant), Hassan Afrouzi, Philippe Andrade, Saleem Bahaj, Jeremy Boccanfuso, Benjamin Born, Giacomo Candian, Luca Gemmi, Paul Hubert, Alex Kohlhas, Francesco Lippi, Ricardo Reis, Rosen Valchev, Albert Marcet, Mirko Wiederholdt, Allan Timmermann, 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

Surveys of forecasts by households, firms, and professionals have become critical tools for understanding expectations, testing theories of belief formation, and guiding policy (Bachmann, Topa and van der Klaauw, 2022). Yet, a puzzling feature of survey forecasts is their tendency to “overreact” to news—forecast revisions often exceed what information changes would justify, leading to predictable forecast errors incompatible with rational expectations. Behavioral biases are frequently invoked to explain this overreaction (Bordalo, Gennaioli and Shleifer, 2022).

We offer a complementary perspective, arguing that much of this apparent overreaction can be explained by two rational mechanisms: lumpiness in forecast revisions and strategic pressures. Suppose forecasters face fixed costs when adjusting their predictions, whether due to a desire to signal stability, the effort required to process new information or the implicit costs of disclosing private information. In such cases, they would revise their forecasts infrequently, resulting in periods of inaction followed by substantial updates. Furthermore, if forecasters operate in environments where reputational concerns matter, they may be reluctant to diverge significantly from the consensus, further reducing the frequency of revisions. Our empirical and quantitative analysis shows that fixed costs and strategic pressures amplify the appearance of overreaction.

To explore how these mechanisms manifest in practice, we study inflation forecasts made by professional forecasters. Using high-frequency data from Bloomberg’s Economic Forecasts (ECFC) survey, which tracks monthly U.S. inflation expectations within a fixed-event framework (Nordhaus, 1987; Patton and Timmermann, 2011), we document three key empirical patterns. First, forecasts are *lumpy*: forecasters often remain inactive for extended periods, even when new information becomes available, before making discrete, substantial updates. Second, forecasts are *strategic*: revisions are influenced by the distance to the consensus, with forecasters closer to the consensus revising less frequently and, when they do, aligning more closely with it. Third, forecasts exhibit *overreaction*: revisions are often larger than what fundamentals alone would justify.

Motivated by these patterns, we develop a forecasting model with three core elements: (1) Bayesian belief updating that generates accurate predictions, (2) fixed revision costs that promote forecast stability, and (3) strategic complementarities that encourage alignment with the consensus. Fixed costs create an inaction region, where forecasters delay revisions until the expected benefit outweighs the cost. Strategic complementarities amplify this effect by introducing a trade-off between prioritizing private accuracy and maintaining reputational alignment with the consensus. The model’s parameters, which capture the relative weight of accuracy, stability, and strategic motives, are calibrated to match cross-sectional moments of forecast revisions in the survey data, including the frequency, size, and hazard rate.

The calibrated model successfully replicates the empirical patterns and provides a novel explanation for overreaction. Large revisions accumulate the effects of both new and prior information, amplifying the perceived overreaction in observed data. Inspired by these results, we propose a

two-stage procedure to refine survey forecasts as measures of true beliefs. The first stage isolates active adjustments by focusing on non-zero revisions. The second stage removes the influence of the consensus through a regression-based adjustment. This procedure reduces the measured overreaction, providing a more accurate signal of the underlying beliefs and enhancing survey forecasts’ interpretive and predictive value.

Overreaction may arise due to psychological biases, but our framework demonstrates how structural frictions can amplify it. For instance, [Bordalo, Gennaioli, Ma and Shleifer \(2020\)](#) attribute overreaction to diagnostic expectations, where recent information is given disproportionate weight in forecasts. Similarly, [Broer and Kohlhas \(2022\)](#) highlights how overconfidence in private signals can exaggerate overreaction. Our framework provides an alternative explanation by showing how adjustment costs and strategic pressures lead to lumpy behavior, which biases estimates upward. These findings complement work on measurement error ([Juodis and Kučinskas, 2023](#)) and information frictions ([Valchev and Gemmi, 2023](#)), emphasizing the importance of both structural and behavioral factors in understanding forecast dynamics.

Our analysis further uncovers evidence of forecasters’ preference for stability by exploiting the overlap between short- and long-term forecasts. Specifically, we observe that even when long-term forecasts are updated in response to new information, short-term forecasts often remain unchanged. This behavior suggests that forecasters actively maintain stability in their predictions despite the persistence of inflation processes that would typically warrant adjustments. These findings strongly support the role of fixed costs in driving forecast lumpiness.

In addition, we highlight significant heterogeneity in the strength of strategic concerns and preferences for stability across different types of forecasters—banks, financial institutions, consulting firms, universities, and research centers. This variation allows us to examine how institutional incentives influence forecast behavior. For example, forecasters in financial institutions may face stronger reputational pressures than those in academia, affecting their propensity to revise forecasts. These results complement studies on model heterogeneity ([Giacomini, Skreta and Turen, 2020](#)) and attention heterogeneity ([Boccanfuso and Neri, 2024](#)), providing a richer understanding of how professional forecasters’ reports are shaped by their environments and incentives.

Contributions We contribute to the literature on forecasting and macroeconomic expectations through multiple perspectives. Our empirical analysis complements evidence on professionals’ lumpy inflation forecasts in the Eurozone ([Andrade and Le Bihan, 2013](#)) and Brazil ([Gaglianone, Giacomini, Issler and Skreta, 2022](#)), and firms’ lumpy sales and price forecasts ([Born, Enders, Müller and Niemann, 2023](#)). Our findings on strategic forecasting provide evidence of reputational concerns as professionals revise forecasts closer to the consensus ([Marinovic, Ottaviani and Sørensen, 2013](#)). This relates to evidence of strategic complementarities in price-setting behavior ([Karadi, Schoenle and Wursten, 2024](#)). Finally, our results on forecast overreaction complement evidence from surveys ([Bordalo, Gennaioli, Ma and Shleifer, 2020](#); [Broer and Kohlhas, 2022](#);

Valchev and Gemmi, 2023) and experimental data (Afrouzi, Kwon, Landier, Ma and Thesmar, 2023). We emphasize the importance of focusing on updaters and correcting their correlation with the consensus to better isolate the true structural level of overreaction in agents’ beliefs.

Our focus on forecast updaters—forecasts with non-zero revisions—naturally connects to the literature on “resetters.” The “reset inflation” measures introduced by Bils, Klenow and Malin (2012) and Blanco and Cravino (2020) isolate the effects of monetary shocks and real exchange rate dynamics by focusing on updated prices. Similarly, Bandeira, Castillo-Martínez and Wang (2024) proposes a “frictionless inflation” measure designed to eliminate lumpiness in price data. Afrouzi, Flynn and Yang (2024) further demonstrate that price adjusters, as the most informed agents, provide sufficient statistics for understanding price dynamics through their “reset uncertainty”. Together, these studies underscore the importance of active adjustments, reinforcing the rationale for our two-stage procedure to refine forecasts.

Our theoretical model draws inspiration from the menu cost literature in price-setting (Barro, 1972; Golosov and Lucas, 2007). As with firms that adjust prices infrequently due to fixed costs, forecasters revise their predictions intermittently, balancing revision costs against the benefits of an accurate prediction. Another central feature of our model is the incorporation of strategic incentives, inspired by Ottaviani and Sørensen (2006), which show that forecasters may adjust their predictions not only based on private information but also to align with the forecasts of others due to reputational and professional considerations. By merging these works of literature, our framework generates a two-dimensional inaction region where inflation and consensus beliefs act as substitutes, akin to a multi-product price-setting model (Midrigan, 2011; Álvarez and Lippi, 2014). Additionally, it features strategic complementarities modeled as a mean-field game (Lasry and Lions, 2007; Alvarez, Lippi and Souganidis, 2023).

By introducing a fixed revision cost to account for observed forecast lumpiness, we complement existing theories in which agents adjust beliefs infrequently due to the costs associated with acquiring or processing information. These include sticky information (Mankiw and Reis, 2002; Reis, 2006a,b), rational inattention (Sims, 2003; Maćkowiak, Matějka and Wiederholt, 2023; Turen, 2023), observation costs (Alvarez, Lippi and Paciello, 2011, 2016), or communication costs (Bec, Boucekkine and Jaret, 2023). We emphasize that forecasts, rather than underlying beliefs, exhibit lumpiness, and provide suggestive evidence of this distinction in the model and the data.

Finally, a key innovation is using a restricted perceptions equilibrium (RPE) to address the computational challenges in solving heterogeneous agent models (Moll, 2024). Our model’s rational expectations equilibrium is infeasible, including aggregate inflation shocks, ex-post heterogeneity due to private signals, lumpy adjustments, and strategic concerns. The whole distribution of forecasts matters in determining current and future consensus. Our approach, inspired by the internally rational framework of Marcet and Nicolini (2003) and Adam and Marcet (2011), assumes forecasters treat the consensus as following a simple random walk, bypassing the need to model higher-order beliefs explicitly while preserving key empirical features.

2 The Anatomy of Inflation Forecasts

We begin by describing the data sources and the fixed-event forecasting framework. We document the evolution of forecast revisions and errors along the forecasting horizon, and we show that forecast revisions are lumpy, strategic, and overreactive.

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

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

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:¹

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

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. The annual inflation range is significant, varying from -0.3% in 2009 to 4.7% in 2022.

2.2 Survey Forecasts

We analyze 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. One of the most appealing features of the Bloomberg 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, making it ideal for studying strategic considerations.²


Sample Our sample covers 2010-2019.³ 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.

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

²[Giacomini, Skreta and Turen \(2020\)](#) compares the Bloomberg and other professional forecasters' surveys.

³We focus on low-volatility years 2010-2019. In [Baley and Turen \(2024\)](#), we analyze the Great Recession, 2008-2009, and the COVID-19 pandemic, 2020-21, years in which the inflation process was more volatile.

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 $x_{m,t}$ are in red dots.

We remove forecasters who fail to provide at least one annual inflation revision. This criteria leaves approximately 100 forecasters per year. The panel dataset contains the history of forecast updates for all forecasters over a 12-month horizon each year.⁴


Incentives This survey is not anonymous. Tracing the entire time series of predictions for any institution across years is possible. 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. Thus, 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 We denote the inflation forecast f_h^i of forecaster $i = 1, \dots, N$ at horizon $h = 12, \dots, 1$ each year. To save on notation, we do not explicitly add a year reference. We count the horizon backward so that the index h indicates that the forecast was produced h months before the end of each corresponding year (the fixed event). Forecasts are measured in percentage points and reported up to one decimal point.

⁴Although we have information on the precise dates when a forecast was revised, we analyze at 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. Although the survey is not anonymous, we remove the names of the institutions to disclose their actual forecasts.

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 during the next month. Figure II illustrates the fixed-event forecasting framework for a given year. The gray dash line accounts for the ultimate value of π during that year. 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.

Given the fixed-event scheme, 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 for the remaining horizons until the end of year $h, \dots, 1$:

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

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:

$$(3) \quad \Delta f_h^i \equiv f_h^i - f_{h+1}^i.$$

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, $\mathbb{E}[\Delta f] = -0.01$, which suggests a symmetric environment in which positive and negative revisions, on average, cancel out. The average revision size (in absolute value, excluding zeros) equals $\mathbb{E}[abs(\Delta f)|\Delta f \neq 0] = 0.25$. There are, on average, five forecast revisions per year, which means forecasts are inactive for 1.6 months on average. The adjustment frequency is 0.43, and downward revisions (0.23) are slightly more likely than upward revisions (0.19).

Table I – Summary Statistics of Forecast Revisions and Errors

Average revision	$\mathbb{E}[\Delta f]$	−0.01
Size non-zero revisions	$\mathbb{E}[abs(\Delta f) \Delta f \neq 0]$	0.25
Avg. number of revisions	$count[\Delta f \neq 0]$	5.06
Months of inaction	$\mathbb{E}[\tau]$	1.60
Adjustment frequency	$\Pr[\Delta f \neq 0]$	0.43
Upward	$\Pr[\Delta f > 0]$	0.19
Downward	$\Pr[\Delta f < 0]$	0.23
Average error	$\mathbb{E}[e]$	−0.05
Mean squared error	$\mathbb{E}[e^2]$	0.26
Observations	N	9,256

Notes: The stylized facts are computed using Bloomberg data from 2010-2019. Cross-sectional statistics are averaged across years and horizons.

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


$$(4) \quad e_h^i \equiv \pi - f_h^i$$

The bottom block of Table I provides summary statistics on individual and aggregate forecast errors, averaged across years and horizons. Individuals make small errors on average $\mathbb{E}[e] = -0.05$ but tend to overpredict inflation, as reflected in the negative error e_h^i .

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. 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 $h = 1$. We interpret this as a tell-tale sign that forecast accuracy is not the only driving force behind forecasters’ behavior.

Figure III – Term Structure of Forecast Revisions and Errors



Notes: Results computed using Bloomberg data from 2010-2019. Panel (a) plots the absolute value of non-zero revisions $E[|\Delta f|_{adjust}]$. Panel (b) plots the mean squared forecast error $E[(\pi - f_h^i)^2]$.

“Naive” 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 Δf_h^{rw} and forecast errors e_h^{rw} that evolve with the horizon according to

$$(5) \quad \Delta f_h^{rw} = (h+1)\Delta x_{h+1},$$

$$(6) \quad e_h^{rw} = \sum_{j=1}^h x_j - hx_{h+1}.$$


The previous figure shows the random-walk case 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 again suggests that forecasts may reflect motives other than accuracy. Building on this, we explore two potential explanations: forecast lumpiness and strategic concerns.

2.6 Forecasts are Lumpy

The first potential explanation for why average revisions and squared errors remain large even at short horizons is that forecasts are lumpy; they remain unchanged for some periods and then undergo a large adjustment. To explore this possibility, we study the frequency of revisions and the revision hazard.

Figure IVa shows the unconditional probability of updating a forecast within the fixed-event scheme across the horizon. Forecasts are updated infrequently as described by Table I. On average, only 43% of forecasters choose to update their predictions throughout the year. The share of updaters also drops as the fixed event approaches. The increasing inaction is puzzling as relevant

Figure IV – Term Structure of Forecast Revisions



Notes: Results computed using Bloomberg data from 2010-2019. 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 $h(age)$.

information accumulates, which could be used to improve the accuracy of the prediction further.

We consider the hazard rate of revisions to see the forecast lumpiness from a different but related angle. 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 sloping, implying that a recent or “young” forecast is more likely to be revised than an “old” forecast. For instance, the probability of adjusting a newly set forecast is 0.5; it drops below 0.3 for six-month-old forecasts and reaches 0.1 for eleven-month-old forecasts.⁵

Together, the adjustment frequency and hazard point towards infrequent forecast revisions.

2.7 Forecasts are Strategic

A second reason why forecast revisions and errors remain large is strategic considerations. Forecasters may care about what the “average” forecaster reports and, thus, may be reluctant to change a forecast that is close to the average, even if that means entertaining a significant forecast error or making large adjustments in the future to compensate for their past mistakes. To assess the role of the consensus forecast in shaping individual forecasting decisions, we adopt the empirical strategy from Karadi, Schoenle and Wursten (2024) that tests for strategic complementarities in the context of firms’ price-setting decisions.

⁵Appendix 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.

Consensus forecast To study the potential role of strategic concerns, we define the *consensus forecast* as the average forecast across the N participants at each horizon h :

$$(7) \quad F_h \equiv \frac{1}{N} \sum_{i=1}^N f_h^i.$$

Let c_h^i be the consensus gap, defined as the individual i 's forecast at horizon $h + 1$ minus the consensus at horizon h :

$$(8) \quad c_h^i \equiv f_{h+1}^i - F_h.$$

Forecasters observe the consensus in real-time through the Bloomberg terminal; thus, it is in their information set. We examine how the consensus gap c_h^i affects the probability of updating a prediction—the extensive margin—and the size of the revision—the intensive margin. We consider equally sized bins for gaps c_h^i , indexed by $b \in [B]$ with $B = 15$, and compute the revision frequency and magnitude in each bin. The two extreme bins include c_h^i below -1.3% or above 1.3% .


Extensive margin First, we run a linear probability model for the probability of revision against bin dummies $\mathbb{1}(c_h^i(b))$ that equal one if c_h^i falls inside bin b :

$$(9) \quad \Pr[\Delta f_{t,h}^i \neq 0] = \beta_0 + \sum_{b=1}^B \beta_j \mathbb{1}(c_{t,h}^i(b)) + \alpha_i + \alpha_t + \alpha_h + \epsilon_{t,h}^i.$$

Estimating different coefficients for each bin captures possible non-linearities in the relationship between the extensive margin of revisions and consensus gaps. We run separate regressions for upward and downward forecast revisions to account for potential asymmetries in the dependence on the consensus. We include forecaster (α_i), year (α_t), and horizon (α_h) fixed effects. The year-horizon fixed effects embed inflation realizations, allowing us to disentangle strategic concerns from the correlation between the consensus and actual inflation. Moreover, coefficients are robust when including the cumulative inflation in year t up to horizon h as a control.

Figure [Va](#) plots the estimated coefficients associated with each dummy, showing the effect of the consensus gap on the probability of revision relative to the omitted category (the middle bin $[-0.1\%, 0.1\%]$). Two interesting features arise. First, as the relative distance between the forecasts and either gap increases, the probability of a revision increases; 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 likelihood of revising downwards ($f_h^i < f_{h+1}^i$) significantly increases. Likewise, when gaps are negative, the likelihood of revising upward substantially increases, and revising downward decreases. Second, the extensive margin reaction appears asymmetric; the updating probability reacts differently depending on

Figure V – Consensus Triggers Revisions



Notes: The estimation relies on forecast data from Bloomberg between 2010-2019. The left panel shows the estimated coefficients of regression (9) where the dependent variable corresponds to a dummy variable taking the value of one if the forecasts was revised upwards (downwards) and zero, otherwise. The right panel repeats the estimation but using the magnitude of revisions (conditioning on updaters) instead. Standard errors are robust and clustered by time, horizon and forecaster.

whether the forecast is below or above the focal point. Thus, the evidence suggests that distance to the consensus is relevant as it triggers forecast revisions.

Intensive margin Conditioning on agent revisions, we now study the determinants behind the magnitude of revisions as a function of the consensus gap. We run a similar specification to equation (9), taking the magnitude of revisions Δf as the dependent variable. As before, we control for forecaster and year-horizon fixed effects. Figure Vb plots the average revision against the consensus gap c_h^i . Revisions, on average, close the gap: Positive deviations call for negative revisions, and negative deviations call for positive revisions. The strong negative correlation implies that larger deviations call for larger revisions.

2.8 Forecasts Exhibit Overreaction

Evidence of forecast overreaction comes from analyzing forecast error predictability at the individual level. Our preferred test, following Broer and Kohlhas (2022) and Valchev and Gemmi (2023), extends the work by Bordalo, Gennaioli, Ma and Shleifer (2020) by adding the consensus as a regressor. We chose this specification as it connects naturally to our empirical evidence by stressing the role of strategic concerns.

Test specification and interpretation Let $\pi_t - f_{t,h}^i$ be the individual ex-post forecast error at horizon h about annual inflation (known at time t), let $f_{t,h}^i - f_{t,h+1}^i$ be the forecast revision between consecutive horizons, and let $F_{t,h} - f_{t,h+1}^i$ be the gap to the consensus. Relying on the panel structure, we consider the following OLS regression with forecaster fixed effects α_i :

$$(10) \quad \underbrace{\pi_t - f_{t,h}^i}_{\text{error}} = \underbrace{\gamma_0^h}_{\text{bias}} + \gamma_1^h \underbrace{(f_{t,h}^i - f_{t,h+1}^i)}_{\text{revision}} + \gamma_2^h \underbrace{(F_{t,h} - f_{t,h+1}^i)}_{\text{consensus}} + \alpha_i + \epsilon_{t,h}^i$$


The key assumption is that the forecasters' information set contains the revision and the consensus. Thus, the rational expectations Bayesian benchmark implies that individual forecast errors are unpredictable: $\gamma_0^h = \gamma_1^h = \gamma_2^h = 0$. Deviations from rational expectations—Bayesian updating—would result in coefficients that differ from zero. In that case, coefficients are interpreted in the following way. If $\gamma_1^h > 0$, it indicates that, on average, forecasters underreact to their information as a positive revision correlates with the forecast being below the actual realization. In contrast, if $\gamma_1^h < 0$ indicates that the average forecasters overreact to his information as a positive revision correlates with the forecast being higher than the actual realization. Analogously, the sign of γ_2^h reflects how distance to the consensus affects forecast errors.

Overreaction to information Using our survey data, we run regression (10). Relative to the literature, which typically runs this regression at a fixed horizon h , we obtain different coefficients for all twelve horizons. The point estimates of the coefficients and their confidence intervals are plotted in Figure VI, where standard errors are robust and clustered by time and forecaster. Panel (a) shows the coefficient on forecast revisions γ_1 , which is negative at all horizons, with an average value of -0.76 . Panel (b) shows the coefficient on the consensus gap γ_2 , which is positive at all horizons, with an average value of 1.05 . According to the standard interpretation, forecasters in the Bloomberg survey show overreaction to information and underreaction to the consensus. These results are consistent with studies using other surveys of professional forecasters. We compare the estimated coefficients reported by Valchev and Gemmi (2023) using the quarterly Survey of Professional Forecasters (SPF), available at horizons $h = 9$ and $h = 6$ (shown as stars).

2.9 Taking Stock and Robustness

To summarize, forecasts are lumpy: they exhibit significant periods of inaction followed by large adjustments. Forecast errors and revision size fall with the forecasting horizon but at a slower pace than a naive random walk forecast would. Also, forecasts are strategic: the distance to the consensus forecast (the average prediction among participants) matters for revisions' extensive and intensive margins. Finally, consistent with existing evidence, forecasters overreact to private information and underreact to the consensus. We conclude this empirical section by discussing several robustness exercises supporting our stylized facts.

Figure VI – Forecast Rationality Tests



Notes: Results computed using Bloomberg data from 2010-2019. The figures show the estimated coefficients in equation (10). Standard errors are robust and clustered by time and forecaster.

Rounding Participants in the Bloomberg survey report their forecasts up to one decimal point. Could rounding artificially generate inaction and mask the underlying level of lumpiness? To assess rounding’s role, in Appendix A.5, we use another survey of professionals, Consensus Economics, where participants can report their inflation forecasts up to three decimal points, thus allowing us to construct counterfactual revision frequencies with various levels of rounding. We show that rounding naturally decreases the adjustment frequency (e.g., a revision below two decimal points gets lost when rounding to one decimal point). Still, it does not significantly alter the lumpy behavior in inflation forecasts across the term structure and the decaying pattern of the extensive and intensive margin of forecast revision along the horizon.

Longer horizons Participants in the Bloomberg survey often report forecasts for the end-of-year inflation in year t at longer horizons $h > 12$. In Appendix A.6, we study the evolution of adjustment frequency, adjustment size, and forecast errors eighteen months before the release $h = 18, 17, \dots, 1$. Regarding lumpiness, the adjustment probability remains relatively stable, around 45% on average between eighteen to thirteen months ahead. However, the magnitude of revisions is, on average, lower relative to its within-the-year counterpart. We interpret this as implying the absence of relevant public information outside the target year, leading forecasters to attenuate their revisions’ magnitude. In Section 6, we exploit the overlap between long and short-horizon forecasts to provide suggestive evidence of a preference for forecast stability.

3 A Structural Model of Lumpy Forecasts

We develop a horizon-dependent fixed-event Bayesian forecasting model with private information, frequent information revelation, fixed revision costs, and strategic concerns.

3.1 Forecasting Problem

Many forecasters, indexed by $i \in N$, generate forecasts of end-of-year inflation π . End-of-year 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:

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

The first term in the payoff function is the distance between the forecast and the actual end-of-year inflation, reflecting losses from the lack of *accuracy*.

The second term is the distance between the forecast and the consensus (the average) $F_h = N^{-1} \sum_{i=1}^N f_h^i$, multiplied by the parameter r that measures the strength of *strategic concerns*.⁶ 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*. Section 3.5 calibrates r and κ using the microdata, and Section 6 provides further suggestive evidence and alternative interpretations of the mechanisms they represent.

Inflation process Forecasters believe monthly inflation follows an autoregressive process:

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

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.

⁶We borrow the term “strategic” from the literature on global games (Morris and Shin, 2002) or mean-field games (Lasry and Lions, 2007), in which small agents consider the distance between their action and the *average* actions of others. We do not consider Cournot-style strategic games with finite and large agents.

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

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

The variance of the public signal is $\sigma_x^2 = \text{Var}[x_h | x_{h+1}] = \text{Var}[\varepsilon_h^x]$.

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

$$(14) \quad \tilde{x}_h^i = x_h + \zeta_h^i, \quad \zeta_h^i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_\zeta^2).$$

The idiosyncratic signal noise σ_ζ^2 reflects the heterogeneity in beliefs, private information, or models across agents. We do not explicitly include public (correlated) noise in this signal because the AR(1) signal plays this role.⁷

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

$$(15) \quad \mathcal{I}_h^i = \tilde{x}_h^i \cup \mathcal{I}_h = \tilde{x}_h^i \cup \{x_{h+1}, x_{h+2}, \dots, F_{h+1}, F_{h+2}, \dots\}.$$

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

3.2 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.⁸

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*

⁷[Valchev and Gemmi \(2023\)](#) explicitly introduce correlated noise.

⁸All proofs appear in Appendix C.

Figure VII – Model Timeline



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

following problem:

$$(16) \quad \min_{\{f_h^i\}_{h=12}^1} \sum_{h=12}^1 \Sigma_h + (f_h^i - \hat{\pi}_h^i)^2 + r(f_h^i - \hat{F}_h)^2 + \kappa \mathbb{1}_{\{f_h^i \neq f_{h+1}^i\}},$$

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

Next, we characterize individual beliefs about end-of-year inflation $\hat{\pi}_h^i$ and the consensus $\hat{F}_h$. To guide the characterization, Figure VII 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:

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

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

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 $\hat{x}_h^i$:

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

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 $\hat{\pi}_h^i$:

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

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:

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

The first term of Σ_h^π corresponds to the uncertainty driven by the AR(1) projection and the noisy signal (weighted by α) for the current release of monthly inflation. Likewise, the second part of (20) reflects the accumulated uncertainty caused by the remaining $(h - 1)$ unforecastable shocks that will hit the process until the release date.

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

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

It is useful to establish a relationship between individual beliefs $\hat{\pi}_h^i$ in (19) under the information set $\mathcal{I}_h^i$ and public beliefs z_h in (21) under the information set $\mathcal{I}_h$. The following relationship links individual and common beliefs:

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

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

Learning from the consensus? We assume that forecasters do not use past consensus realizations as inflation signals; solely, the strategic concern drives the relationship between forecasts and the consensus. Since monthly inflation and consensus realizations become public after one period, private information is short-lived.⁹ Thus, lagged monthly inflation is a superior signal to the lagged consensus. The consensus at $h + 1$ aggregates private information about x_{h+1} , but once x_{h+1} gets observed, the information in the consensus becomes outdated.

If, instead, private information was long-lived (e.g., actual monthly inflation was never released), then the consensus would contain helpful information. However, the problem becomes untractable as it enters the domain of higher-order beliefs and involves forecasting the forecasts of others (Townsend, 1983).

3.3 Equilibrium

We now define our notion of equilibrium. We focus on a restricted perceptions equilibrium (RPE), representing a slight deviation from rational expectations (Evans and Honkapohja, 1993). We posit that forecasters believe the consensus follows a random walk, and ex-post, they cannot distinguish the actual consensus process from a random walk. Forecasters are *internally rational* (Marcet and Nicolini, 2003), as they use an “internally consistent” learning model. This equilibrium concept delivers enormous tractability by eliminating the fixed point between the consensus and the aggregation of individual forecasts.¹⁰

⁹Our truncation of the information is a particular case of the algorithms developed for solving a class of dynamic models with higher-order expectations by Nimark (2008, 2014), or the assumption that shocks become common knowledge after a finite yet arbitrarily large delay, as in Hellwig and Venkateswaran (2009).

¹⁰The RPE 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 (2019).

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)

$$(23) \quad \hat{F}_h = g(\hat{F}_{h+1}, \delta) + \epsilon_h^{\hat{F}}, \quad \epsilon_h^{\hat{F}} \sim \mathcal{N}(0, \sigma_F^2)$$

(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 (23), forecasts $\{f_h^i\}$ are optimal and solve the forecasting problem (16);
2. parameters (δ, σ_F^2) are such that the forecast errors arising from predicting the actual consensus using the perceived law of motion, i.e., $\epsilon_h^F \equiv F_h - g(F_{h+1}, \delta)$, satisfy: $\text{Cov}[\epsilon_h^F, \epsilon_j^F] = 0$ $\forall h \neq j$ and $\text{Var}[\epsilon_h^F] = \sigma_F^2$.

In the restricted perceptions equilibrium, the actual consensus process given by the aggregation of individual forecasts, $F_h = N^{-1} \sum_{i=1}^N f_h^i$, differs from the prediction. However, in this equilibrium concept, agents are assumed to use the δ, σ_F that best predicts future prices given (23).

3.4 Optimal Forecasting 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 shared across agents, we include it in the aggregate state. We thus index value function with the horizon h to account for the aggregate state.

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


$$(24) \quad \mathcal{V}_h(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i) = \min \left\{ \underbrace{\mathcal{V}_h^I(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i)}_{\text{inaction}}, \underbrace{\mathcal{V}_h^A(\hat{\pi}_h^i, \hat{F}_h)}_{\text{action}} \right\}$$

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 (23).

Figure VIII – Forecast Revision Policy: Inaction and Reset



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

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

$$(25) \quad \mathcal{R}_h \equiv \{(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i) : \mathcal{V}_h^I(\hat{\pi}_h^i, \hat{F}_h, f_{h+1}^i) \geq \mathcal{V}_h^A(\hat{\pi}_h^i, \hat{F}_h)\},$$

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

$$(26) \quad \Delta f_h = \begin{cases} 0 & \text{if } f_{h+1}^i \in \mathcal{R}_h \\ f_h^{i*} - f_{h+1}^i & \text{if } f_{h+1}^i \notin \mathcal{R}_h. \end{cases}$$

Panel (a) in Figure VIII shows the forecast revision policy at horizon $h = 6$. We plot it in two dimensions by fixing the current forecast at $f_6^i = 2$ and varying inflation beliefs $\hat{\pi}_h^i$ in the x -axis and consensus beliefs $\hat{F}_h$ in the y -axis. Panel (b) plots the width of the inaction region measured on the 45-degree line, against the forecasting horizon. We use the parametrization in Table II.

The inaction region is the negatively sloped dark band centered around the current forecast. Outside the band, the lines correspond to the reset forecast $f_6^{i*}(\hat{\pi}_6^i, \hat{F}_6)$. For example, given the current forecast $f_6^i = 2$, the beliefs $(\hat{\pi}_6^i, \hat{F}_6) = (2.5, 1.5)$ fall inside $\mathcal{R}_6$ and the forecast remains inactive; instead, the beliefs $(\hat{\pi}_6^i, \hat{F}_6) = (1.5, 2)$ fall outside $\mathcal{R}_6$ and the forecast is reset to $f_6^{i*} = 1.7$, revising it by $\Delta f_6^i = 1.7 - 2 = -0.3$.

Several features of the optimal forecasting policy are worth explaining.

First, the negative slope in the inaction region arises because the two beliefs are “substitutes” in that a smaller distance to the consensus belief may compensate for a greater distance from the inflation belief, or vice versa. This band-type inaction region contrasts with the circular inaction regions typical in multiproduct menu cost pricing models (Midrigan, 2011; Álvarez and Lippi, 2014). The reason is that, in those models, the gaps between current and optimal prices are independent, and different instruments (the price of each good) are available to target different variables (the optimal price of each good). Instead, in our setup, one instrument (the forecast f_h^i) targets two variables (inflation and consensus beliefs).

A second feature is that the width of the inaction region shrinks with the horizon.¹¹ At long horizons, belief uncertainty is at its highest level; forecasters anticipate that their belief would hit the band very often and thus optimally widen the band to minimize adjustment cost payments. This is known as an *option effect*. As belief uncertainty falls, the option effect is smaller, and the band shrinks. A shrinking inaction region implies that adjustment size falls with the horizon. Still, the impact on the adjustment frequency is nuanced because frequency depends on the option effect and the *volatility effect*. The price-setting model with idiosyncratic cost uncertainty and learning in Baley and Blanco (2019) features a similar decreasing inaction region.

Finally, the frictions shape the inaction region in different ways. The strength of strategic concerns determines the slope of the inaction region. The inaction region would be vertical (only the inflation belief matters) if $r = 0$, horizontal if $r \rightarrow \infty$, and negative with $r > 0$. The inaction region widens with the revision cost κ and private noise σ_ζ^2 , as standard in Ss-type models.¹²

3.5 Calibration and Solution

Externally set parameters Frequency is monthly. 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 we estimate the parameters $(c_x, \phi_x, \sigma_x) = (0.013, 0.932, 0.036)$. These values imply an unconditional annual inflation of $\mu_\pi = 12c_x/(1 - \phi_x) = 2.23$ with annual volatility $\sigma_\pi^2 = \sigma_x^2 \sum_{h=1}^{12} (1 - \phi_x^h)^2 / (1 - \phi_x^2) = 0.49$.¹³

Internally calibrated parameters Using the simulated method of moments (SMM), we estimate values for the three remaining parameters by matching the cross-sectional moments across years: the fixed revision cost κ , the strength of strategic concerns r , and the private noise σ_ζ .

¹¹We see a widening of the inaction region at $h = 1$ arising from the finite-horizon nature of the problem.

¹²Appendix E for comparative statics on the optional forecast policy for different values of the fixed cost κ , strategic concerns r , and private noise σ_ζ^2 .

¹³The inflation process estimation details appear in Appendix B. The monthly process is highly persistent $\phi = 0.932$ because it refers to year-on-year monthly inflation, not between consecutive months, whose autoregressive coefficient typically ranges between 0.5 and 0.7.

Table II – Internally calibrated parameters

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

We target three moments: the frequency of revisions $\Pr[\Delta f \neq 0] = 0.43$, the average absolute value of revisions $\mathbb{E}[abs(\Delta f)|adjust] = 0.25$ and the slope of the hazard rate between horizons 12 and 6 equal to -0.04 . The hazard’s slope informs idiosyncratic signal noise. Learning is slow when signals are very noisy, and the hazard rate declines slowly. In contrast, learning is faster when signals are less noisy, and the hazard rate declines faster.¹⁴

Internal consistency of consensus beliefs Forecasters 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 . Under internal rationality, the consensus’s perceived and actual probability distributions should coincide. This assumption imposes structure and disciplines the value of σ_F . Starting with a guess for the volatility of the consensus process σ_F^2 , we compute individual decision rules for each horizon h using backward induction. We then simulate the model, calculate the volatility of the realized consensus, and iterate on σ_F^2 to ensure belief consistency.¹⁵ As further validation, we run a Dickey-Fuller test and cannot reject the null of a random walk when considering a sample of 10 years or less.

Estimated Parameters Table II 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.¹⁶

4 The Model in Action


This section explores various dimensions of the forecasting model.

¹⁴The idea that signal noise modulates the slope of the adjustment hazard underlies the calibration of information frictions in price-setting models by [Alvarez, Lippi and Paciello \(2011\)](#), [Baley and Blanco \(2019\)](#) and [Argente and Yeh \(2022\)](#) and labor market models by [Borovicková \(2016\)](#) and [Baley, Figueiredo and Ulbricht \(2022\)](#).

¹⁵Appendix D explains the solution algorithm and other computational details.

¹⁶Appendix F presents the details on the consistency of consensus beliefs.

Figure IX – Forecaster-level dynamics



Notes: The figure illustrates the beliefs, forecasts, revisions, and uncertainty dynamics of one forecaster for one year. Panel (a) shows the evolution of forecasts f_h^i and beliefs ($\hat{\pi}_h^i$, $\hat{F}_h^i$), and the end-of-year inflation. Panel (b) shows the magnitude of revisions (intensive margin) and the periods of inaction (extensive margin). Panel (c) shows the evolution of total uncertainty split between belief and consensus uncertainty.

4.1 Individual Dynamics

To explain the model's workings, Figure IX illustrates how one agent's beliefs, forecasts, revisions, and uncertainty evolve during one year. The first panel shows the agent's inflation beliefs $\hat{\pi}_h^i$ (green line) and consensus beliefs $\hat{F}_h^i$ (dash green line). Beliefs change from period to period but forecasts f_h^i (blue line) exhibit lumpy behavior, remaining fixed for some periods, followed by revisions that bring the forecasts closer to a linear combination of the two beliefs. Towards the year's end, the inflation belief meets actual inflation ($\pi = 3.77$), but the forecast is not.


The second panel shows the extensive margin of adjustment (gray areas), marking the periods of inaction between horizons 8 and 5 and 3 and 1. It also indicates the intensive margin of adjustment given by the revision size Δf (dark line), which shrinks with the horizon.

The third panel shows total uncertainty (solid pink line), equal to the weighted sum of conditional variance of inflation beliefs Σ_h^π , continuously decreasing and reaching zero at $h = 1$, and the conditional variance of consensus beliefs $r\sigma_F^2$, which is constant. While belief uncertainty dominates most of the prediction period, consensus uncertainty dominates in the last three periods.

4.2 Aggregate Dynamics

Next, we examine the dynamics of aggregate forecasts and aggregate beliefs. Figure Xa plots the realized value for the end-of-year inflation $\pi = 3.77$ (horizontal dash black line). We also show the

Figure X – Aggregate dynamics



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


average belief $\hat{\Pi}_h \equiv N^{-1} \sum_{i=1}^N \hat{\pi}_h^i$ (green solid line), the consensus forecast $F_h \equiv N^{-1} \sum_{i=1}^N \hat{f}_h^i$ (blue solid line) and 100 individual forecasts f_h^i (gray lines). The two aggregate series approximate the actual end-of-year inflation π as information becomes available throughout the year. The average belief converges to actual inflation, but the consensus does not. In this example, the consensus remains below actual inflation. Importantly, average beliefs are more volatile than the consensus, meaning the micro-level lumpiness does not fully wash out in the aggregate.

To highlight the difference between forecasts and beliefs, Figure Xb 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, pooled across all years and horizons. Both distributions are centered around zero. The belief distribution is unimodal, but the forecast distribution is bimodal, as in the data, resulting from the adjustment cost κ .

4.3 Untargeted Term Structures

We assess the model's capacity to generate the term structure of cross-sectional statistics. Figure XI shows the term structure of the revisions frequency, size, and hazard rate. While we only targeted average values (the dashed lines), the model matches the empirical patterns along the forecasting period. In particular, the model can quantitatively match the downward sloping patterns of the frequency of revisions (extensive margin, Figure XIa) and the size of non-zero revisions (intensive margin, Figure XIb). In addition, the model accurately matches the hazard's level despite only targeting its slope.

Figure XI – Cross-sectional statistics across horizons




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

The downward-sloping hazard is a feature of Bayesian learning models with fixed adjustment costs in actions, arising from a shrinking inaction region and decreasing uncertainty, with uncertainty falling faster (Baley and Veldkamp, 2025). Its slope serves as a discrimination device across models of inaction. 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 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) model). Alternatively, 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. Therefore, combining learning and fixed revision costs delivers a decreasing hazard.

4.4 Other Untargeted Moments

We provide further evidence of the model’s ability to replicate untargeted moments. Panel (a) in Figure XII shows mean squared forecast errors $\mathbb{E}[(\pi - f_h^i)^2]$, which are closely matched on average (0.15 in the model and 0.23 in the data) and in the horizon profile. Panel (b) shows the distribution of the final revision date. To compute it, we find forecasters’ average fraction (across years) that provides their final revision at horizon h . In the data, on average, 40% of participants do their last revision four months before the release date. The model matches the same qualitative pattern, suggesting that our theory can characterize the marginal benefits of waiting for an extra release of information relative to the revision costs throughout the horizons.

Figure XII – Untargetted Forecast Errors and Final Revision Date



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

Figure XIII shows the untargeted model's performance regarding the consensus gap. Panel (a) examines how the consensus gap c_h^i impacts the extensive margin of adjustment. The patterns are qualitatively consistent with the data. The model also quantitatively replicates the reductions in revision probability. Still, the increases in revision probability are more responsive in the model.¹⁷ Panel (b) examines how the consensus gap c_h^i impacts the intensive margin of adjustment. In this case, the model qualitatively matches the magnitude of adjustments in the data. Quantitatively, we also observe a slightly larger responsiveness in the model.


4.5 On the Role of Each Friction

We finish this section by exploring the role of fixed costs and strategic concerns in delivering empirical patterns. We shut down κ and r , one at a time, and recalibrate the model through the SMM procedure to match a subset of moments. The results are shown in Table III. Estimated parameters are shown in the first four rows, and targets are in the last three. Targeted moments are marked with stars. Columns (1) and (2) repeat the information in Table II with the data targets and the baseline calibration for reference.

Column (3) shows the results when shutting down the fixed costs ($\kappa = 0$), and only strategic concerns are present. Without fixed costs, forecasters continuously revise; thus, we cannot match the adjustment frequency and hazard slope. We estimate r to match the average size of non-zero revisions. Interestingly, we find a negative value of $r = -0.43$, indicating that strategic diversification is necessary to match the size of adjustments. In other words, including fixed costs shifts the data's implications for r from positive (complements) to negative (substitutes).

¹⁷Introducing free adjustment opportunities in the spirit of the CalvoPlus model, which combines state and time-dependent adjustment, or a generalized hazard model could bring the consensus response closer to data.

Figure XIII – Untargetted Extensive and Intensive Margins



Notes: Estimated coefficients are obtained by running equation (9), using the model’s simulation under the benchmark calibration.


Column (4) shows the results when shutting down strategic complementarities ($r = 0$), and we set κ and σ_η^2 to match the intensive margin and the hazard’s slope. Two main drawbacks appear in this case. First, the model is less effective in matching the data-implied probability of revisions relative to our baseline specification. Second, and most importantly, the model cannot reconcile the downward-sloping pattern of the hazard rate. The hazard rate is almost flat in this case. The fact that the probability of revisions becomes less “age-dependent” is a direct implication of removing the strategic concerns. Empirically, an “older” forecast is less likely to be revised than a recently updated prediction. Intuitively, this makes the consensus forecast more persistent as a function of age. When agents stop carrying about the relative distance between their predictions and the consensus, the updating probability becomes less sensible to age. This is precisely the result we get in this case. Further, untargeted moments and the intuition for these two alternative cases are discussed in Appendix G.

Table III – Shutting down frictions

	(1) Data	(2) Baseline	(3) Only strategic ($\kappa = 0$)	(4) Only fixed costs ($r = 0$)
Parameters				
κ		0.06	0.00	0.07
r		0.75	-0.43	0.00
σ_ζ^2		0.09	0.03	0.09
σ_F^2		0.12	0.23	0.18
Targets				
Frequency	0.43	0.40*	1.00	0.38*
Size	0.25	0.19*	0.25*	0.30*
Hazard Slope	-0.04	-0.04*	∞	-0.01*

Notes: In the table, * denotes a targeted moment.

Figure XIV – Forecast Rationality Tests



Notes: The figures show the estimated coefficients after running equation (10) using both the survey data from Bloomberg 2010-2019 (red line and stars) and the simulated data from the model (blue line).

5 Overreaction to Information

This section examines the potential role of lumpy and strategic forecasts in amplifying overreaction.

Lumpy forecasts overreact Using model-generated forecasts, we run regression (10), repeated here for convenience:

$$(27) \quad \pi_t - f_{t,h}^i = \gamma_0^h + \gamma_1^h (f_{t,h}^i - f_{t,h+1}^i) + \gamma_2^h (F_{t,h} - f_{t,h+1}^i) + \alpha_i + \epsilon_{t,h}^i$$

Figure XIV shows the results. As in the data, simulated forecasts feature (i) over-reaction to private information, $\gamma_1 < 0$ in Panel (a), and (ii) under-reaction to the consensus, $\gamma_2 > 0$ in Panel (b). These patterns hold across all horizons. Crucially, these results are untargeted; therefore, they also provide an additional layer of validation of our model for lumpy forecasters.

Intuition The measured overreaction to information arises directly from lumpy behavior, characterized by the coexistence of large and zero revisions. While information accumulates continuously, expanding the agent's information set, the reported forecast remains unchanged unless it exceeds the inaction region. When adjustments occur, they incorporate both current and previously accumulated information, resulting in substantial revisions. This creates a pattern where large revisions coexist with many zero revisions at each horizon, reducing the overall variance of revisions. This reduced variance mechanically inflates the estimated coefficient γ_1 , which captures the covariance of revisions and forecast errors relative to the variance of revisions. As lumpiness amplifies this effect, the econometrician may interpret it as a significant overreaction, with predictions appearing to overshoot their forecast.

Cleansing procedure To refine survey forecasts as measures of true underlying beliefs, we propose a two-stage procedure to cleanse individual forecasts in the Bloomberg survey from lumpy and strategic behavior. In the first stage, we isolate active adjustments by focusing exclusively on non-zero revisions. In the second stage, we remove the influence of the consensus using a regression-based adjustment. This approach effectively adds a preprocessing step to the standard overreaction test, where forecast errors are regressed on forecast revisions. By conditioning on updates and using our cleansed forecasts, this method provides a more accurate representation of the underlying belief dynamics. Since beliefs are not directly observable in the data, we use overreaction tests as a lens to examine them.

5.1 Stage I: Eliminate Lumpiness

We re-run regression (10) in the data but conditioning on updaters, i.e., $\Delta f_h^i \neq 0$. These “reset” forecasts account for the latest release of information and all the information accumulated since the last time they were revised. Therefore, forecast revisions are a much more accurate proxy for the amount of acquired information once we condition on updates.

The results of the overreaction test, conditioned on updaters $\Delta f_h^i \neq 0$ at each horizon, are illustrated by the black dashed line in Figure XV. The estimated coefficients γ_1^h shift closer to zero—the Bayesian rational expectation benchmark—when focusing solely on updates. A similar pattern emerges from the consensus-related coefficients, γ_2^h . While the evidence continues to indicate overreaction and underreaction, our analysis demonstrates that failing to account for lumpy behavior can obscure its underlying level. This oversight can lead to biased estimates of structural parameters in models attempting to replicate and explain such dynamics.

5.2 Stage II: Eliminate Strategic Concerns


After leveraging the relevance of lumpy behavior, we focus on the potential role that strategic concerns play in deviating (recently updated) expectations. Inspired by our theoretical model, we introduce a second stage to account for this feature, which builds on an auxiliary reduced form regression to pin down the parameter r using data alone.

Recover r from an OLS regression In the model, conditional on resetting the forecast, the first-order condition requires forecasts to be a convex combination of inflation and consensus beliefs plus a term reflecting the change in the continuation value:

$$(28) \quad f_h^{i*} = \frac{1}{1+r}(\hat{\pi}_h^i + r\hat{F}_h^i) + \frac{\partial \mathbb{E}[\mathcal{V}_h]}{\partial f_h^i} = \frac{1}{1+r}(z_h + rF_{h+1} + v_h^i),$$

where in the second equality, we have substituted the relationship between individual and average forecasts in (22), the random-walk consensus beliefs, and joined all idiosyncratic terms (including

Figure XV – Forecast Rationality: Updaters and Cleansed



Notes: The figures show the estimated coefficients after running equation (10) using survey data. All forecasters (solid red line), updaters (dashed black line), and cleansed (dotted pink line). The point estimates are plotted along with the corresponding confidence intervals for every possible horizon. Regressions include individual fixed effects, and the standard errors are robust and clustered by time and forecaster.

the continuation value) in a fixed effect term v_h^i . From this expression, we propose the following OLS regression to back out a data-implied value for strategic concerns r :

$$(29) \quad f_{h,t}^i = \underbrace{\beta_0}_{0.47(0.28)} + \underbrace{\beta_1}_{0.28(0.06)} z_{h,t} + \underbrace{\beta_2}_{0.44(0.11)} F_{h+1,t} + \alpha_h + \alpha_y + \alpha_i + \epsilon_{i,h,t}$$

To estimate it, we proxy z_h with the corresponding AR(1) projection using the *available* monthly releases of inflation at each horizon. As the estimation relies on the level of predicted inflation, we include forecasters, horizon, and year-fixed effects. In our theory, the r parameter is independent of agents and horizons, so we estimate polling all observations and horizons.¹⁸ The parameter r can be recovered from β_1 or β_2 . However, β_1 is more likely to be biased due to measurement error, model misspecification, or any omitted relevant variable that was part of the forecaster's i information set at horizon h . Thus, we chose to recover r through β_2 since we have a direct and observed measure of F_{h+1} in the data.¹⁹

Our preferred estimate is $\hat{r} = 0.79$. It is reassuring that the data-implied r closely resembles the SMM-implied value of 0.73 in Table II, further validating the calibration of our model.

¹⁸Detailed estimates appear in Appendix H. We run various specifications. We add macro controls conditioning to account for potential omitted variables, including first and second lags of annualized inflation, the growth of industrial production, and the three-month treasury yields. The estimated r remains stable across specifications.

¹⁹We see this as an advantage of our proposed procedure, as in almost all surveys of professional forecasters, participants can observe the lagged consensus through the survey's reports or newsletters. Hence, while we run our proposed test using the Bloomberg survey where the consensus forecast is available to all participants in real-time, this does not prevent any researcher from running the proposed first stage using any other expectations survey as long as the previous consensus is available at the time agents are asked to provide new forecasts.

Cleansing for strategic concerns Given the first-order condition in (28), and using the estimated value for $\hat{r}$, we back-out a proxy for agent’s beliefs “cleansed” from strategic concerns

$$(30) \quad \tilde{f}_h^i = (1 + \hat{r})f_h^i - \hat{r}F_{h+1}.$$

With the beliefs proxy $\tilde{f}_h^i$ we run the alternative specification:

$$(31) \quad \pi_t - \tilde{f}_h^i = \gamma_0^h + \gamma_1^h(\tilde{f}_h^i - \tilde{f}_{h+1}^i) + \alpha_i + \epsilon_{t,h}^i$$

The pink dash line in Panel (a) of Figure XV illustrates the results.

This analysis yields two key insights. First, the estimates of γ_1^h move closer to the rational expectations benchmark of zero when using our belief proxies and conditioning on updates. Specifically, the average value of γ_1^h goes from -0.76 (all data) to -0.68 (updates) and further to -0.29 (updates and cleansed), representing a 58% reduction in the magnitude of overreaction. This circles back to the fact that failure to account for both lumpy behavior and strategic motives in a professional survey may amplify the underlying level of overreaction. Second, while the estimated coefficient decreases, the level of overreaction remains significantly different from zero. This suggests that, even after removing strategic and lumpy behavior, agents still overshoot acquired information, aligning with previous findings (Bordalo, Gennaioli and Shleifer, 2022).

5.3 Additional Metrics

We close the analysis by studying the differences between (A) raw forecasts, (B) reset forecasts, and (C) cleansed forecasts in the data, with (D) forecasts and (E) beliefs in the model. Table IV showcases metrics such as volatility and autocorrelations of these variables.

Forecast volatility (line 1) increases from 0.6 to 0.8 when lumpy behavior (column B) and strategic considerations (column C) are removed. Reported forecasts typically incorporate a weighted consensus, which is highly persistent. Once strategic considerations are accounted for, individual forecasts become more volatile, reflecting agents’ evolving perceptions of the annual inflation process as new information emerges. This pattern aligns with the model, where beliefs are slightly more volatile than forecasts.

Two key patterns emerge regarding the autocorrelations of forecasts (line 2.1) and forecast errors (line 2.2). First, lumpiness amplifies the autocorrelation of forecasts (0.81) compared to cases where we condition only on updates (0.54). This is expected, as lumpy behavior causes forecasters to maintain the same predictions over time, increasing the persistence of individual forecast processes. Second, the autocorrelation of cleansed forecasts increases slightly (0.62) relative to the updating correction, reflecting the influence of the consensus process. Removing the variability of the consensus smooths out short-term fluctuations, resulting in more correlated residual forecasts over time. The model aligns with the observed drop in autocorrelation for both forecasts and

Table IV – Forecast Statistics

	Forecast Data			Model	
	(A) Raw f	(B) Updaters $f \Delta f \neq 0$	(C) Cleansed $\tilde{f} \Delta f \neq 0$	(D) Forecasts f	(E) Beliefs $\hat{\pi}$
(1) Volatility	0.60	0.64	0.80	0.92	0.95
(2) Autocorrelation					
(2.1) Forecasts	0.81	0.54	0.62	0.75	0.60
(2.2) Errors	0.81	0.51	0.65	0.75	0.60
(2.3) Revisions	−0.08	−0.19	−0.19	−0.05	−0.15

Notes: Bloomberg data for years 2010-2019. The model is solved under the benchmark calibration.

errors, and the autocorrelations of beliefs closely match their cleansed counterparts.

Regarding the autocorrelation of revisions (line 2.3), we notice that while it is negative, it remains close to zero across all forecasts. The magnitude increases when we condition on revised forecasts, which is consistent with the underlying inflation process being driftless.²⁰ Additionally, due to lumpiness, the coexistence of zero and non-zero forecast revisions results in low overall correlation. In a fixed-event forecast scheme, revisions serve as a proxy for the information set. The negative correlation observed in columns B and C suggests that upward revisions are often followed by downward revisions, consistent with a stationary inflation process where uncertainty decreases over time as more information becomes available (illustrated in Figure II). Finally, the model accurately replicates the autocorrelation of revisions, whether computed from forecasts or beliefs, further validating its consistency with the data.

6 Supportive Evidence and Interpretations


We conclude the paper by providing supporting evidence for the frictions central to our analysis and their implications for forecasting behavior.

6.1 Fixed Revisions Costs

Fixed costs and the resulting inaction in forecast revisions may arise from various sources. These could include the costs associated with acquiring and processing information or the opportunity cost of logging into the forecasting system. Another explanation is that forecasters may prefer to signal commitment to their predictions, favoring stability over frequent adjustments. For example, forecasters might refrain from revising their predictions when they believe a monthly inflation release reflects only transitory shocks that will dissipate in subsequent data. We favor this latter explanation and provide suggestive evidence of a preference for forecast stability.

²⁰A non-stationary inflation process would mechanically generate persistent revisions.

Figure XVI – Inaction in the short vs. long run



Notes: The figure relies on data from Bloomberg between 2010-2019. In this case, we rely on predictions reported up to 18 months before the release of end-of-year inflation. The figure illustrates the preference for forecast stability.

We construct a statistic that measures the overlap of short—and long-term revisions to provide suggestive evidence of forecasters’ preference for forecast stability. Besides the within-year forecasts used in our main analysis, the sample contains longer-term forecasts from 18 months to 13 months ahead.²¹ These long-term forecasts for year $t + 1$ overlap with short-term forecasts for year t . Given the overlap for any given year, we compute the probability that a short-term forecast remains inactive (i.e., $f_h^i = f_{h+1}^i$ for $h = 6, \dots, 1$) conditional on a long-term forecast being revised (i.e., $f_h^i \neq f_{h+1}^i$ for $h = 18, \dots, 13$). If a long-term forecast gets revised, it signals that some relevant information has been received, which causes the participant to log into the Bloomberg terminal to revise such prediction. If there is persistence in the inflation process, that information should also affect short-term forecasts. To the extent that this overlap probability is lower than one, it suggests that forecasters actively decide to maintain their short-term forecasts unchanged.

Figure XVI plots this statistic. Given that the corresponding long-term forecast was changed, the probability of keeping a 6-month ahead forecast unchanged is 0.2. This probability increases as the horizon shrinks. We take this as suggestive evidence for forecast stability.

Forecast stability in other contexts In corporate finance, firms often rely on lumpy forecasts to signal stability and enhance credibility with investors. For example, firms may discontinue quarterly earnings guidance or strategically time and present financial disclosures to minimize earnings volatility and manage market expectations (Chen, Matsumoto and Rajgopal, 2011; Barton and Simko, 2002). Surveys of CFOs further emphasize the importance of maintaining stable forecasts to build trust among stakeholders (Graham, Harvey and Rajgopal, 2005).

²¹See Appendix A.6 for cross-sectional statistics for long-term forecasts.

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	0.40	0.45	0.44	0.38	0.37	0.47	0.47	0.34	0.32
$\mathbb{E}[\Delta f adjust]$	0.25	0.19	0.25	0.18	0.26	0.23	0.27	0.18	0.29	0.27
Hazard Slope	-0.04	-0.04	-0.05	-0.05	-0.02	-0.02	-0.05	-0.05	-0.01	-0.01
Observations	12,355		5,366		2,567		2,982		1,440	

Notes: Own calculations based on Bloomberg data between 2010-2019. For each specification we re-calibrate the model to match the specific targeted moments for each of the subgroups.

Beyond economics, forecast lumpiness or anchoring is frequently motivated by a desire to maintain consistency and credibility. In meteorology, forecasters adopt a lumpy approach to updating weather predictions to preserve the credibility of their projections. Frequent back-and-forth adjustments can undermine trust and create confusion among users (Griffiths, Marzocca and Michaelides, 2019; Murphy and Winkler, 1987; Mullen and Buizza, 1993; Stewart and Lusk, 1994).

Our findings provide a deeper understanding of how these principles extend to professional economic forecasting.

6.2 Strategic Motives and Heterogeneity

Finally, we analyze the behavior of different forecaster types to provide evidence of strategic motives. The survey categorizes forecasters into four groups: (i) financial institutions, (ii) banks, (iii) consulting companies, and (iv) universities and research centers. Table V highlights substantial heterogeneity in average cross-sectional moments. The most notable differences emerge between consulting firms and universities; for example, consulting firms adjust their forecasts 30% more frequently than universities.

To account for these differences, we recalibrate the model to match type-specific moments. The variations in the cross-sectional moments of forecast revisions and errors suggest potential differences in fixed costs (κ), strategic concerns (r), private signal noise (σ_ζ), and the internally-rational consensus volatility (σ_F). Given the observed heterogeneity, we calibrate four distinct versions of the model, each tailored to match the moments specific to a forecaster type. The results of these recalibrations are presented in Table VI. We normalize values for financial institutions at unity to ease the comparison.

Through the lens of the calibrated models, universities face higher revision costs and the lowest strategic concerns among the four types. These differences underscore the importance of considering forecast heterogeneity across forecaster types when analyzing survey data. Our findings highlight that forecasters' ex-ante heterogeneity, shaped by revision costs, strategic concerns, and private signal noise, is critical in understanding professional forecasting behavior.

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

Parameter	Financial Inst.	Banks	Consulting	Universities
κ	1.00 (0.06)	1.08	0.94	1.29
r	1.00 (0.81)	0.62	0.89	0.50
σ_ζ	1.00 (0.04)	1.16	1.14	2.28
σ_F	1.00 (0.10)	1.13	1.08	1.33

Notes: Calibration that targets the group-specific moments reported in Table V. The last line reports the RPE-implied consensus volatility σ_F . Estimated parameters for banks, consulting companies, and universities are expressed relative to those estimated for financial institutions (reported in parenthesis)

7 Final Thoughts

Professional forecasters revise their predictions in a lumpy manner, influenced by both new information and strategic considerations. To capture this behavior, we develop a horizon-dependent forecasting model that integrates preferences for stability and strategic concerns, successfully replicating the observed data patterns. Importantly, while forecasters often appear to overreact to private information, our cleansing method reveals that much of this overreaction diminishes when accounting for these frictions, providing a clearer lens into their underlying beliefs.

Our findings have significant implications for the design and interpretation of expectations surveys. Forecasters’ incentives, such as preferences for stability or reputational concerns, can distort the measurement of beliefs, particularly when updates are infrequent or strategically motivated. For instance, evidence from the Brazilian FOCUS survey ([Gaglianone, Giacomini, Issler and Skreta, 2022](#)) shows that forecast accuracy and update frequency increase significantly around contests rewarding precision, while [Ottaviani and Sørensen \(2006\)](#) highlight how competitive environments influence the differentiation of forecasts. These insights suggest that better-designed incentives could promote more accurate and frequent updates, enhancing the reliability of survey-based measures of expectations.

In related work ([Baley and Turen, 2024](#)), we examine how lumpy forecasts respond to monetary policy and information shocks, offering a framework to understand their impact on macroeconomic expectations. This work underscores the importance of accounting for frictions in belief formation, highlighting how such shocks can influence learning patterns and potentially cause deviations in expectations from announced policy targets. Our findings shed light on the broader implications for central bank communication strategies and the alignment of beliefs with policy objectives.

References

- ADAM, K. and MARCET, A. (2011). Internal rationality, imperfect market knowledge and asset prices. *Journal of Economic Theory*, **146** (3), 1224–1252.
- AFROUZI, H., FLYNN, J. P. and YANG, C. (2024). *What Can Measured Beliefs Tell Us About Monetary Non-Neutrality?* Tech. rep., National Bureau of Economic Research.
- , KWON, S. Y., LANDIER, A., MA, Y. and THESMAR, D. (2023). Overreaction in expectations: Evidence and theory. *The Quarterly Journal of Economics*, **138** (3), 1713–1764.
- ÁLVAREZ, F. and LIPPI, F. (2014). Price setting with menu cost for multi-product firms. *Econometrica*, **82** (1), 89–135.
- ALVAREZ, F., LIPPI, F. and PACIELLO, L. (2016). Monetary shocks in models with inattentive producers. *The Review of Economic Studies*, **83** (2), 421–459.
- , — and SOUGANIDIS, P. (2023). Price setting with strategic complementarities as a mean field game. *Econometrica*, **91** (6), 2005–2039.
- 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. *The Review of Economic Studies*, **89** (6), 2992–3054.
- BACHMANN, R., TOPA, G. and VAN DER KLAUW, 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.
- , FIGUEIREDO, A. and ULBRICHT, R. (2022). Mismatch cycles. *Journal of Political Economy*, **130**, 2943–2984.
- and TUREN, J. (2024). *Lumpy Forecasts and Inflation Volatility*. Tech. rep.
- and VELDKAMP, L. (2025). *The Data Economy: Tools and Applications*. Princeton University Press.

- BANDEIRA, M., CASTILLO-MARTÍNEZ, L. and WANG, S. (2024). *Frictionless Inflation*. Tech. rep., Working Paper.
- BARRO, R. J. (1972). A theory of monopolistic price adjustment. *Review of Economic Studies*, **39** (1), 17–26.
- BARTON, J. and SIMKO, P. J. (2002). The balance sheet as an earnings management constraint. *The Accounting Review*, **77** (s-1), 1–27.
- 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.
- BILS, M., KLENOW, P. J. and MALIN, B. A. (2012). Reset price inflation and the impact of monetary policy shocks. *American Economic Review*, **102** (6), 2798–2825.
- BLANCO, A. and CRAVINO, J. (2020). Price rigidities and the relative ppp. *Journal of Monetary Economics*, **116**, 104–116.
- BOCCANFUSO, J. and NERI, L. (2024). Uncovering attention heterogeneity. Working paper.
- BORDALO, P., GENNAIOLI, N., MA, Y. and SHLEIFER, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, **110** (9), 2748–82.
- , — and SHLEIFER, A. (2022). Overreaction and diagnostic expectations in macroeconomics. *Journal of Economic Perspectives*, **36** (3), 223–244.
- BORN, B., ENDERS, Z., MÜLLER, G. J. and NIEMANN, K. (2023). Firm expectations about production and prices: Facts, determinants, and effects. In *Handbook of Economic Expectations*, Elsevier, pp. 355–383.
- BOROVICKOVÁ, K. (2016). *Job flows, worker flows and labor market policies*. Tech. rep., New York University.
- BROER, T. and KOHLHAS, A. N. (2022). Forecaster (mis-) behavior. *Review of Economics and Statistics*, pp. 1–45.
- CALVO, G. (1983). Staggered prices in a utility maximizing framework. *Journal of Monetary Economics*, **12**, 383–398.
- CHEN, S., MATSUMOTO, D. and RAJGOPAL, S. (2011). Is silence golden? an empirical analysis of firms that stop giving quarterly earnings guidance. *Journal of Accounting and Economics*, **51** (1-2), 134–150.

- 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.
- GAGLIANONE, W. P., GIACOMINI, R., ISSLER, J. V. and SKRETA, V. (2022). Incentive-driven inattention. *Journal of Econometrics*, **231** (1), 188–212.
- GIACOMINI, R., SKRETA, V. and TUREN, J. (2020). Heterogeneity, inattention, and bayesian updates. *American Economic Journal: Macroeconomics*, **12** (1), 282–309.
- GOLOSOV, M. and LUCAS, R. E. (2007). Menu costs and phillips curves. *Journal of Political Economy*, **115** (2), 171–199.
- GRAHAM, J. R., HARVEY, C. R. and RAJGOPAL, S. (2005). The economic implications of corporate financial reporting. *Journal of Accounting and Economics*, **40** (1-3), 3–73.
- GRIFFITHS, M., MARZOCCA, M. and MICHAELIDES, A. (2019). Flip-flop index: Quantifying the stability of fixed-event forecasts. *Meteorological Applications*, **26** (4), 580–589.
- HELLWIG, C. and VENKATESWARAN, V. (2009). Setting the right prices for the wrong reasons. *Journal of Monetary Economics*, **56**, S57–S77.
- JUODIS, A. and KUČINSKAS, D. (2023). On the role of measurement error in estimating beta in bgms regressions. *Quantitative Economics*, forthcoming.
- KARADI, P., SCHOENLE, R. and WURSTEN, J. (2024). Price selection in the microdata. *Journal of Political Economy Macroeconomics*, **2** (2), 000–000.
- 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.
- LASRY, J.-M. and LIONS, P.-L. (2007). Mean field games. *Japanese journal of mathematics*, **2** (1), 229–260.
- MANKIW, N. G. and REIS, R. (2002). Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *Quarterly Journal of Economics*, **117** (4), 1295–1328.
- MARCET, A. and NICOLINI, J. P. (2003). Recurrent hyperinflations and learning. *American Economic Review*, **93** (5), 1476–1498.

- MARINOVIC, I., OTTAVIANI, M. and SØRENSEN, P. (2013). Forecasting incentives: Strategic diversification and complementarity. In G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting*, vol. 2, Elsevier, pp. 1203–1251.
- MAĆKOWIAK, B., MATĚJKA, F. and WIEDERHOLT, M. (2023). Rational inattention: A review. *Journal of Economic Literature*, **61** (1), 226–273.
- MIDRIGAN, V. (2011). Menu cost, multiproduct firms, and aggregate fluctuations. *Econometrica*, **79** (4), 1139–1180.
- MOLAVI, P. (2019). Macroeconomics with learning and misspecification: A general theory and applications. *Unpublished manuscript*.
- MOLL, B. (2024). Heterogeneous agent macroeconomics: Eight lessons and a challenge. *Economic Journal Lecture*.
- MORRIS, S. and SHIN, H. S. (2002). Social value of public information. *American Economic Review*, **92** (5), 1521–1534.
- MULLEN, S. L. and BUIZZA, R. (1993). Forecasting with flip-flop indices: An application to weather forecasting. *Monthly Weather Review*, **121** (12), 3448–3457.
- MURPHY, A. H. and WINKLER, R. L. (1987). A general framework for forecast verification. *Monthly Weather Review*, **115** (7), 1330–1338.
- NIMARK, K. (2008). Dynamic pricing and imperfect common knowledge. *Journal of monetary Economics*, **55** (2), 365–382.
- NIMARK, K. P. (2014). Man-bites-dog business cycles. *American Economic Review*, **104** (8), 2320–67.
- NORDHAUS, W. D. (1987). Forecasting efficiency: concepts and applications. *The Review of Economics and Statistics*, pp. 667–674.
- OTTAVIANI, M. and SØRENSEN, P. (2006). The strategy of professional forecasting. *Journal of Financial Economics*, **81** (2), 441–466.
- PATTON, A. J. and TIMMERMAN, 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.
- REIS, R. (2006a). Inattentive consumers. *Journal of Monetary Economics*, **53** (8), 1761–1800.

- (2006b). Inattentive producers. *The Review of Economic Studies*, **73** (3), 793–821.
- SIMS, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, **50** (3), 665–690.
- STEWART, T. R. and LUSK, C. M. (1994). Seven components of judgmental forecasting skill: Implications for research and the improvement of forecasts. *Journal of Forecasting*, **13** (7), 579–599.
- TOWNSEND, R. M. (1983). Forecasting the forecasts of others. *Journal of Political Economy*, **91** (4), 546–588.
- TUREN, J. (2023). State-dependent attention and pricing decisions. *American Economic Journal: Macroeconomics*, **15** (2), 161–189.
- VALCHEV, R. and GEMMI, L. (2023). *Biased Surveys*. Tech. rep., National Bureau of Economic Research.

Lumpy Forecasts

Isaac Baley and Javier Turen

Online Appendix

Contents

A	Data	3
A.1	Inflation definitions	3
A.2	Inflation summary statistics	4
A.3	Cross-sectional statistics by year	4
A.4	Hazard rates by number of revisions	5
A.5	Extensive and Intensive Margin - Consensus Economics Survey	6
A.6	Cross-sectional statistics at long horizons	7
A.7	Cross-sectional statistics by forecaster type	8
B	Estimate of inflation process	9
C	Proofs	10
C.1	Inflation process	10
C.2	End-of-year inflation beliefs	11
C.3	Relationship between individual vs. aggregate beliefs	12
C.4	Martingale property of beliefs	13
C.5	Proof of Proposition 1	14
C.6	Proof of Proposition 2	15
D	Computational strategy	16
D.1	Initial forecast	16
D.2	Distributions of expected beliefs	16
D.3	Computing expectations	18
E	Comparative statics	19
F	Consistency of consensus process	21
G	On the role of fixed costs and strategic concerns	22
H	Cleansing forecasts in the data	24

A Data

A.1 Inflation definitions

We show how to approximate yearly inflation through year-on-year monthly inflation. As introduced in Section 2.1, 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:

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

This last condition is what we show in Section 2.1. From this, it is important to stress *when* is the sum of year-on-year monthly inflation a good approximation of annual inflation? Let us consider a second-order Taylor approximation of $\log(p)$ around $\mathbb{E}[p]$, which yields:

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

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

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

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

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

Hence, for similar within-year price dispersion ($\mathbb{C}V^2[p] \approx \mathbb{C}V^2[p']$) for two consecutive years, 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	$\text{Var}[\pi]$	1.621	0.622	4.267	2.478
Years		14	10	4	2

Notes: The CPI Index for the US is extracted from FRED. In this case we labelled the years 2010-2019 as Normal years, 2008-09 and 2020-21 as Turbulent years, and 2020-2021 as the Pandemic years.

A.3 Cross-sectional statistics by year

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




Notes: This figure reports the evolution of the extensive, intensive and the MSFE across years (gray lines), along with the average across years (solid black line). The data comes from Bloomberg between 2010 and 2019.

A.4 Hazard rates by number of revisions

In this section, we compute the hazard rates given the number of revisions the forecasters have made 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



Notes: Own calculations using Bloomberg data between 2010 and 2019. For each case, we compute the probability of a forecast revisions as function of age, i.e., the hazard rate, conditioning on the number of revisions made, over the term structure.


Independently of the revision, the decaying pattern of the hazard rates remains across specifications. While the probability of an immediate revision right at age one is roughly the same across the number of revisions (between 45% and 50%), we notice that 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 except for the first revision hazard.

We interpret the decaying probability of the number of revisions as an implication of the fixed-event scheme. Mechanically, as more revisions accumulate, the probability of doing additional revisions drops as there are only a few remaining horizons before the target variable is finally released. Overall, the fact that the adjustment decisions drop with the age of the forecast is also consistent with the decaying pattern of the average frequency of revisions over the term structure.

A.5 Extensive and Intensive Margin - Consensus Economics Survey

In this Section, we examine the robustness of our results using the *Consensus Economics Survey of Professional Forecasters*. Again, we focus on individual expectations at the monthly frequency for end-of-year inflation between 1995 and 2016 in the US. We have repeated the analysis as in our baseline data through the years. One of the key differences between Consensus and the Bloomberg survey is that in the former, participants can report predictions up to 3 decimal points. Thus, we will contrast the extensive and intensive margin dynamics using the raw forecasts (No Rounding) with the projections rounded up to the first decimal point (Rounding). Figure A.3 shows the evolution of the adjustment probability.


Figure A.3 – Extensive Margin - Consensus Economics Survey



Notes: Own calculations based on inflation expectations collected from *Consensus Economics Survey of Professional Forecasters* during 1995 and 2016.

As noticed, even when we rely on three decimal predictions, the evidence still supports lumpy behavior in inflation forecasts. During the last months before the variable was released, we also observed a drop in the frequency of non-rounded forecast revisions, consistent with our data. When we round the predictions up to the first decimal, the evolution of extensive margin resembles the dynamics of the Bloomberg data. Figure A.4 reports the evolution of the magnitude of revisions across the horizons. Again, there are no significant differences relative to our original results.

Figure A.4 – Intensive Margin - Consensus Economics Survey




Notes: Own calculations based on inflation expectations collected from *Consensus Economics Survey of Professional Forecasters* during 1995 and 2016.

A.6 Cross-sectional statistics at long horizons

Figure A.5 below shows the cross-sectional statistics after extending the forecast horizon to 18 months ahead. As discussed, although there is information about inflation between eighteen and thirteen months ahead, it is only twelve to one month ahead when there is relevant information about monthly inflation. While the frequency of revisions remains relatively similar to the updating probability between twelve and nine months ahead, the magnitude of revisions drops significantly when we are out of the target year. We interpret this reduction as implying the absence of relevant information in these longer terms. Consistent with this last intuition, the forecast error rises abruptly at longer horizons.

Figure A.5 – Long-term forecasts




Notes: The first row of the figure shows the extensive and intensive margins of forecast revisions. The second row shows the variance and the MSFE of forecasts. The figure plots the evolution of these variables for a forecast horizon of eighteen months. The data comes from Bloomberg between 2010 and 2019.

A.7 Cross-sectional statistics by forecaster type


Figure A.6 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.

Figure A.6 – Term Structure of Revisions and Errors: By Forecaster Type



Notes: The figure shows the evolution of the extensive and intensive margin of forecast revisions and the hazard rate for each of the four subgroups of survey participants. The data comes from Bloomberg between 2010 and 2019.

Figure A.7 – Forecast Errors by Groups



Notes: The figure shows the evolution of the MSFE for each of the four subgroups of survey participants. The data comes from Bloomberg between 2010 and 2019.

B Estimate of inflation process


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

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

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 .

Through OLS, 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.8 plots the resulting estimates and 95% confidence intervals.

Figure B.8 – Rolling Estimates for Inflation Parameters



Notes: Estimated results using official monthly inflation rate from the CPI between 2008-2021. The externally set parameters are given by the simple average across the studied sample 2010-2019.

We include the more Great recession 2008-2009 and the COVID pandemic 2020-2021 as a comparison. Regarding our range of years, it is clear that the parameters are fairly stable with respect to the more turbulent years counterpart.

C Proofs

C.1 Inflation process

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

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

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

$$(C.2) \quad x_h - \frac{c_x}{1-\phi_x} = \phi_x \left(x_{h+1} - \frac{c_x}{1-\phi_x} \right) + \varepsilon_h^x.$$

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} :

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

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

$$(C.3) \quad \pi = 12 \left(\frac{c_x}{1-\phi_x} \right) + \frac{1-\phi_x^{12}}{1-\phi_x} \left(x_{12} - \frac{c_x}{1-\phi_x} \right) + \sum_{j=1}^{11} \frac{1-\phi_x^j}{1-\phi_x} \varepsilon_j^x.$$

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:

$$(C.4) \quad \pi = h \left(\frac{c_x}{1-\phi_x} \right) + \frac{(1-\phi_x^h)}{1-\phi_x} \left(x_h - \frac{c_x}{1-\phi_x} \right) + \sum_{i=h+1}^{12} x_i + \sum_{j=1}^{h-1} \frac{1-\phi_x^j}{1-\phi_x} \varepsilon_j^x,$$

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:

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

$$(C.6) \quad \text{Var}[\pi] = \sigma_x^2 \sum_{j=1}^{h-1} \left(\frac{1 - \phi_x^j}{1 - \phi_x} \right)^2.$$

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

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

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

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

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 $\tilde{x}_h^i = x_h + \zeta_{ih}$. In addition, the historical values of lagged monthly inflation are observed without noise. Thus, the forecasters information set at each horizon $\mathcal{I}_h^i = \{\tilde{x}_h^i, x_{h+1}, x_{h+2}, \dots\}$.

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

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

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):

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

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 :

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

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

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

Substituting back into (C.11), we obtain the expression for the signal variance in (C.12)

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

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

C.3 Relationship between individual vs. aggregate beliefs

To construct individual belief about yearly inflation $\hat{\pi}_h^i$ in (19), forecasters combines the past release x_{h+1} with their noisy private signal $\tilde{x}_h^i$ to generate a monthly belief $\hat{x}_h^i$, which is then projected to obtain the yearly forecast

$$(C.13) \quad \hat{\pi}_h^i = h \left(\frac{c_x}{1 - \phi_x} \right) + \frac{1 - \phi_x^h}{1 - \phi_x} \left(\hat{x}_h^i - \frac{c_x}{1 - \phi_x} \right) + \sum_{j=h+1}^{12} x_j.$$

In contrast, the public belief about yearly inflation z_h in (21) only projects the past release x_{h+1} to obtain the yearly forecast (note the extra ϕ_x in the second term of the expression reflecting the timing of the information):

$$(C.14) \quad z_h = h \left(\frac{c_x}{1 - \phi_x} \right) + \frac{\phi_x(1 - \phi_x^h)}{1 - \phi_x} \left(x_{h+1} - \frac{c_x}{1 - \phi_x} \right) + \sum_{j=h+1}^{12} x_j.$$

Next, we establish a useful relationship between the private and public beliefs about yearly inflation. Starting from (C.13), we substitute the expression for $\hat{x}_h^i = (1 - \alpha)x_h^{AR} + \alpha\tilde{x}_h^i$. Then, we substitute $x_h^{AR} = \mathbb{E}[x_h | \mathcal{I}_h] = c_x + \phi_x x_{h+1}$ and the noisy signal $\tilde{x}_h^i = x_h + \zeta_h^i$. We also use $x_h = x_h^{AR} + \varepsilon_h^x$. Lastly, we define the noise term $\nu_h^i \equiv \frac{1 - \phi_x^h}{1 - \phi_x} \alpha(\varepsilon_h^x + \zeta_h^i)$, which includes idiosyncratic signal noise and the one-period ahead forecasting error arising from the different timing in the use

of information.

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

where α is the weight on private signals: $\alpha \equiv \sigma_\zeta^{-2} / (\sigma_x^{-2} + \sigma_\zeta^{-2})$. We can further simplify the noise term since:

$$\text{(C.16)} \quad \alpha^2 (\sigma_x^2 + \sigma_\zeta^2) = \left(\frac{\sigma_\zeta^{-2}}{\sigma_x^{-2} + \sigma_\zeta^{-2}} \right)^2 (\sigma_x^2 + \sigma_\zeta^2) = \left(\frac{\frac{1}{\sigma_\zeta^2}}{\frac{1}{\sigma_x^2} + \frac{1}{\sigma_\zeta^2}} \right)^2 (\sigma_x^2 + \sigma_\zeta^2)$$

$$\text{(C.17)} \quad = \left(\frac{\frac{1}{\sigma_\zeta^2}}{\frac{\sigma_\zeta^2 + \sigma_x^2}{\sigma_\zeta^2 \sigma_x^2}} \right)^2 (\sigma_x^2 + \sigma_\zeta^2) = \left(\frac{\sigma_x^2}{\sigma_\zeta^2 + \sigma_x^2} \right)^2 (\sigma_x^2 + \sigma_\zeta^2) = \frac{\sigma_x^4}{\sigma_\zeta^2 + \sigma_x^2}$$

Therefore, individual beliefs are decomposed as:

$$\text{(C.18)} \quad \hat{\pi}_h^i = z_h + \nu_h^i, \quad \text{where } \nu_h^i \sim \mathcal{N} \left(0, \left[\frac{1 - \phi_x^h}{1 - \phi_x} \right]^2 \frac{\sigma_x^4}{\sigma_\zeta^2 + \sigma_x^2} \right).$$

When signal noise is very large ($\sigma_\zeta^2 \rightarrow \infty$), the idiosyncratic component of beliefs has zero dispersion because private signals are ignored. When signal noise is tiny ($\sigma_\zeta^2 \rightarrow 0$), the idiosyncratic component of beliefs has dispersion equal to σ_x^2 . Beliefs become perfectly correlated, and the remaining noise comes from projecting x_{h+1} instead of x_h .

C.4 Martingale property of beliefs

We show that beliefs follow a martingale, that is, the expectation of future belief at $h + 1$ equals the current belief at h , i.e., $\mathbb{E}[\hat{\pi}_{h+1}^i | \mathcal{I}_h^i] = \hat{\pi}_h^i$. First, we use the relationship between public and private beliefs in (C.15) to set the expectation of future individual noise ν to zero.

$$\text{(C.19)} \quad \mathbb{E}[\hat{\pi}_{h+1}^i | \mathcal{I}_h^i] = \mathbb{E}[z_{h+1} + \nu_{h+1}^i | \mathcal{I}_h^i] = \mathbb{E}[z_{h+1} | \mathcal{I}_h^i].$$

Second, 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:

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

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

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

We conclude that $\mathbb{E}[z_{h-1}|\mathcal{I}_h^i] = \hat{\pi}_h^i$. 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:

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

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:

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

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 $\mathbb{E}[(\hat{\pi}_h^i - \pi) | \mathcal{I}_h^i] = \mathbb{E}[(\hat{F}_h^i - F_h) | \mathcal{I}_h^i] = 0$:

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

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:

As it is standard, the solution to the stopping time problem is characterized by solving the following problem. Let $(\hat{\pi}_h^i, \hat{F}_h^i, f_{h+1}^i)$ be the state of the forecaster i at horizon h . Then the value $\mathcal{V}_h(\hat{\pi}_h^i, \hat{F}_h^i, f_{h+1}^i)$ is given by

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

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

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

subject to the evolution of inflation beliefs in (19) and (20), and consensus beliefs in (17).

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 ($\kappa = r = 0$):

$$(D.21) \quad f_{13}^i = \hat{\pi}_{13}^i = z_{13} + \nu_{13}^i, \quad \nu_{13}^i \sim \mathcal{N}(0, \sigma_{13}^2)$$

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

$$(D.22) \quad z_{13} = 12 \left(\frac{c_x}{1 - \phi_x} \right) + \phi_x \frac{1 - \phi_x^{12}}{1 - \phi_x} \left(\hat{x}_{13} - \frac{c_x}{1 - \phi_x} \right)$$

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

D.2 Distributions of expected beliefs

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

$$(D.23) \quad \hat{\pi}_{h-1}^i = \left(\frac{\Sigma_{z,h-1}}{\sigma_o^2 + \Sigma_{z,h-1}} \right) \mu_o + \left(1 - \frac{\Sigma_{z,h-1}}{\sigma_o^2 + \Sigma_{z,h-1}} \right) \hat{\pi}^i \hat{\pi}_{h-1}^i$$

$$(D.24) \quad \hat{F}_{h-1} = c_F + \phi_F F_h$$

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:

$$(D.25) \quad \mathbb{E}[\hat{F}_{h-1} | \mathcal{I}_h^i] = c_F + \phi_F \mathbb{E}[F_h | \mathcal{I}_h^i] = c_F(1 + \phi_F) + \phi_F^2 F_{h+1}$$

$$(D.26) \quad \mathbb{V}ar[\hat{F}_{h-1} | \mathcal{I}_h^i] = \phi_F^2 \mathbb{V}ar[F_h | \mathcal{I}_h^i] = \phi_F^2 \sigma_F^2$$

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

$$(D.27) \quad \mathbb{E}[\hat{\pi}_{h-1}^i | \mathcal{I}_h^i] = \left(\frac{\Sigma_{z,h-1}}{\sigma_o^2 + \Sigma_{z,h-1}} \right) \mu_o + \left(1 - \frac{\Sigma_{z,h-1}}{\sigma_o^2 + \Sigma_{z,h-1}} \right) \mathbb{E}[\hat{\pi}_{h-1}^i | \mathcal{I}_h^i]$$

$$(D.28) \quad \mathbb{V}ar[\hat{\pi}_{h-1}^i | \mathcal{I}_h^i] = \left(\frac{\sigma_o^2 \Sigma_{z,h-1}}{\sigma_o^2 + \Sigma_{z,h-1}} \right)^2 \mathbb{V}ar[\hat{\pi}_{h-1}^i | \mathcal{I}_h^i]$$

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

Expected signals We evaluate the formula for $\hat{\pi}_h^i$ in (C.18) at $h-1$, and separate the observation x_h from the sum yields:

$$(D.29) \quad \hat{\pi}_{h-1}^i = (h-1) \left(\frac{c_x}{1-\phi_x} \right) + \frac{1-\phi_x^{h-1}}{1-\phi_x} \left(\tilde{x}_{h-1}^i - \frac{c_x}{1-\phi_x} \right) + x_h + \sum_{j=h+1}^{12} x_j.$$

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

$$(D.30) \quad \mathbb{E}[\hat{\pi}_{h-1}^i | \mathcal{I}_h^i] = (h-1) \left(\frac{c_x}{1-\phi_x} \right) + \frac{1-\phi_x^{h-1}}{1-\phi_x} \left(\mathbb{E}[\tilde{x}_{h-1}^i | \mathcal{I}_h^i] - \frac{c_x}{1-\phi_x} \right) + \mathbb{E}[x_h | \mathcal{I}_h^i] + \sum_{j=h+1}^{12} x_j$$

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

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

Rearranging, we obtain:

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

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

$$(D.31) \quad \mathbb{E}[\hat{\pi}_{h-1}^i | \mathcal{I}_h^i] = h \left(\frac{c_x}{1-\phi_x} \right) + \phi_x^2 \frac{1-\phi_x^{h-1}}{(1-\phi_x)^2} x_{h+1} + \phi_x \left(x_{h+1} - \frac{c_x}{1-\phi_x} \right) + \sum_{j=h+1}^{12} x_j.$$

For the variance, we apply the variance operator to (D.29) and note that the terms in the sum disappear because they are known at h . Thus we are left with two terms.

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

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

$$(D.32) \quad \tilde{x}_{h-1}^i = x_{h-1}^i + \zeta_{h-1}^i + \eta_{h-1} = c_x + \phi_x x_h + \varepsilon_{h-1}^x + \zeta_{h-1}^i + \eta_{h-1}.$$

D.3 Computing expectations

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

$$(D.33) \quad \mathbb{E}[\mathcal{V}_{h-1}(\hat{\pi}_{h-1}^i, \hat{F}_{h-1}, f_{h+1}^i) | \mathcal{I}_h^i] = \sum_{\hat{\pi}_{h-1}^i} \sum_{\hat{F}_{h-1}} \mathcal{V}_{h-1}(\hat{\pi}_{h-1}^i, \hat{F}_{h-1}, f_{h+1}^i) \omega(\hat{\pi}^i) \omega(\hat{F})$$

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 $\hat{F}_{h-1} | \mathcal{I}_h^i \sim \mathcal{N}(\mathbb{E}[\hat{F}_{h-1} | \mathcal{I}_h^i], \text{Var}[\hat{F}_{h-1} | \mathcal{I}_h^i])$ with

$$\begin{aligned} \mathbb{E}[\hat{F}_{h-1} | \mathcal{I}_h^i] &= F_{h+1} \\ \text{Var}[\hat{F}_{h-1} | \mathcal{I}_h^i] &= \sigma_F^2 \end{aligned}$$


Integration weights $\omega_{\hat{\pi}^i}$ are such that $\hat{\pi}_{h-1}^i | \mathcal{I}_h^i \sim \mathcal{N}(\mathbb{E}[\hat{\pi}_{h-1}^i | \mathcal{I}_h^i], \text{Var}[\hat{\pi}_{h-1}^i | \mathcal{I}_h^i])$, with

$$\begin{aligned} \mathbb{E}[\hat{\pi}_{h-1}^i | \mathcal{I}_h^i] &= h \left(\frac{c_x}{1 - \phi_x} \right) + \phi_x^2 \frac{1 - \phi_x^{h-1}}{(1 - \phi_x)^2} x_{h+1} + \phi_x \left(x_{h+1} - \frac{c_x}{1 - \phi_x} \right) + \sum_{j=h+1}^{12} x_j \\ \text{Var}[\hat{\pi}_{h-1}^i | \mathcal{I}_h^i] &= \sigma_x^2 + \left(\frac{1 - \phi_x^{h-1}}{1 - \phi_x} \right)^2 (\phi_x^2 \sigma_x^2 + \sigma_x^2 + \sigma_\zeta^2) \end{aligned}$$


E Comparative statics

We show how parameters affect the optimal forecasting policy.


Preference for stability κ A higher fixed revision cost makes the inaction band wider.




Strategic concerns r Strategic concerns shape the slope of the inaction region.



Idiosyncratic noise σ_ζ^2 More noise increases the option effect, widening the inaction region.



Horizon h As the horizon shrinks, uncertainty falls and thus it is a similar effect to σ_ζ^2 .



F Consistency of consensus process


In this section, we further validate the consistency of the consensus's assumed random walk. According to the estimation, the perceived law of motion for consensus is $F_h = F_{h+1} + \varepsilon_h^F$ with $\varepsilon_h^F \sim \mathcal{N}(0, 0.11^2)$. Thus, the perceived process is

$$(F.34) \quad \hat{F}_t = \hat{F}_{t-1} + \varepsilon_t^{\hat{F}}, \quad \varepsilon_h^F \sim \mathcal{N}(0, 0.11^2)$$

The actual law of motion is

$$(F.35) \quad F_h = -0.03 + 1.01F_{h+1} + \varepsilon_h^F, \quad \varepsilon_h^F \sim \mathcal{N}(0, 0.11^2).$$

Figure F.13 – Consensus belief consistency



Notes: The figure shows the evolution of the consensus forecasts F and the beliefs $\hat{F}$ over time. The time series is constructed using our proposed model using the benchmark calibration.


Dickey-Fuller test We run a Dickey-Fuller test on the simulated series of the actual consensus process F to test the null hypothesis that a unit root is present. The estimate of interest is ρ in the expression

$$(F.36) \quad F_{t+1} = \alpha + \rho F_t + \varepsilon_{t+1}.$$

The estimation uses N_{years} randomly drawn from the model. Figure F.14 shows the average estimate and the 95% confidence interval obtained by bootstrap when $N_{years} \in \{1, \dots, 10\}$ are employed in the estimation.

Given the test, we can not reject the null that the consensus process F_{t+1} follows a unit root process.

Figure F.14 – Dickey-Fuller Tests




Notes: Average estimate of ρ and the 95% confidence interval obtained by bootstrap when $N_{years} \in \{1, \dots, 10\}$.

G On the role of fixed 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 G.15, 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.


Figure G.15 – Cross-sectional moments across model configurations



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

As expected, the model with zero fixed costs implies a frequency of revisions equal to one for all horizons and thus fails dramatically in replicating the observed empirical patterns. Similarly, as shown in Figure G.15c, the hazard rate concentrates all the probability at age zero in this case. Likewise, as agents constantly revise their predictions, the distribution of final revisions is always at zero, as depicted by Figure G.16b. Finally, while the intensive margin and the mean square error resemble the baseline model and the data, the absence of the stability cost makes the model sensible to the consensus gap regarding the positive and negative relative probability.

Figure G.16 – Untargeted Moments




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

The model with no strategic concerns only marginally decreases the frequency of revisions relative to the baseline model and, consequently, slightly increases the size of revisions. What is interesting is the behavior of the hazard rate. As discussed in the data, the hazard rate is downward sloping, meaning that an “older” forecast is much less likely to be revised than a recently revised one. Intuitively, this means that the consensus forecast becomes more persistent as a function of age. Therefore, when we remove the concern for being close to the average, the updating probability becomes less “age-dependent”, delivering a flatter hazard rate, which is precisely what Figure G.15c shows.²²

²²This result holds even after we specifically target such slope as described by Table III.

Figure G.17 – Untargeted Moments - Extensive and Intensive Margins



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

H Cleansing forecasts in the data

Here we present the results of our auxiliary regression to back-out the data-implied parameter r . The results are shown in Table H.2.

As noticed, it is relevant to estimate the regression conditioning on updaters only as a further validation of our calibrated theory. As shown in column (2) of Table H.2, when we condition on updaters, $r = 0.79$, entirely in line with our estimated parameters. The precision of the estimation improves significantly when we account for updates only. The estimated r remains relatively stable when we add further macro controls such as the lagged inflation rate, industrial production, and the 3-month T-Bill rate. In all the estimations we included forecasters, horizon, and year-fixed effects. The standard errors are robust and clustered by forecaster and time, while the reputation concern parameter $\hat{r}$ was estimated through $\hat{\beta}_2$ using the Delta-Method.

Table H.2 – Individual forecast determinants

	All	Updaters	
β_1	0.1998 (.0444)	0.2807 (.0646)	0.2668 (.0566)
β_2	0.5791 (.1070)	0.4411 (.1113)	0.4846 (.0888)
Constant	0.4033 (.2036)	0.4798 (.2765)	0.5215 (.2868)
Macro Controls	×	×	✓
Horizon, Year FE	✓	✓	✓
Forecasters FE	✓	✓	✓
N	9,562	3,898	3,898
R^2	0.7674	0.8398	0.8501
$\hat{r} = \hat{\beta}_2/(1 - \hat{\beta}_2)$	1.3760 (.6041)	0.7891 (.3561)	0.9401 (.3343)

Notes: The table shows the estimated coefficients from equation (29). The first column includes all observation, while the second and third columns conditions on non-zero revisions. We include forecasters, horizon, and year-fixed effects. Standard errors are robust and clustered by forecaster and time. The $\hat{r}$ is estimated through $\hat{\beta}_2$ using the Delta-Method.