# Lumpy Forecasts[*]

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

**Javier Turen** [‡]
Pontificia Universidad Católica de Chile

December 23, 2025

**Abstract**

Professional forecasters face frictions—the desire to project stability, reputational concerns about appearing erratic or misaligned, and the institutional cost of crafting and communicating revised narratives. These frictions distort the link between true beliefs and reported survey forecasts. We show that inflation forecasts are updated infrequently and strategically, leading to lumpy, consensus-aligned revisions and apparent overreaction to news. A model with Bayesian updating, fixed revision costs, and strategic alignment replicates these dynamics. Variation across institutions and forecast horizons validates the role of frictions and helps rule out alternative explanations. We propose a two-step correction that isolates underlying beliefs, offering a sharper lens into expectations.

**JEL**: D83, D84, E17, E37

**Keywords**: surveys of professional forecasters, fixed-event forecasting, consensus, forecast revision costs, inflation, strategic concerns, Bayesian learning, overreaction.

---

[*] We largely benefited from conversations with João Victor Issler (discussant), Hassan Afrouzi, Philippe Andrade, Alex Kohlhas, Joel Flynn, and Allan Timmermann. We also thank Saleem Bahaj, Jeremy Boccanfuso, Benjamin Born, Giacomo Candian, Leland Farmer, Laura Gati, Luca Gemmi, Raffaella Giacomini, Paul Hubert, Francesco Lippi, Albert Marcet, Ricardo Reis, Oliver Pfäuti, Rosen Valchev, Mirko Wiederholdt, and many seminar participants for their suggestions. Baley acknowledges financial support from Severo Ochoa Research Excellence Program (CEX2024- 001476-S) through Spain's State Research Agency (Agencia Estatal de Investigación) and the 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 professional forecasters are central tools for understanding expectations, testing models of belief formation, and guiding policy decisions (Coibion and Gorodnichenko, 2015; Bachmann, Topa and van der Klaauw, 2022). Yet, the process of generating forecasts involves frictions that distort the link between true beliefs and reported survey responses. Forecasters may be reluctant to revise frequently due to the desire to project stability and avoid appearing erratic, the implicit costs of disclosing private information, or the institutional burden of crafting and communicating a revised narrative to their clients or the public. These factors contribute to periods of inaction followed by large forecast revisions (Andrade and Le Bihan, 2013). Furthermore, if forecasters operate in environments where reputational concerns matter (Ottaviani and Sørensen, 2006), they may be reluctant to diverge significantly from others, further reducing the frequency of revisions.

We document two salient features of professional forecasts—lumpy forecast revisions and strategic pressures—and show they are central to the dynamics of inflation expectations. Using data from Bloomberg's Economic Forecasts (ECFC) survey, which tracks annual U.S. inflation forecasts in a fixed-event setting (Nordhaus, 1987; Patton and Timmermann, 2011), we find that forecast revisions are *lumpy*: forecasters often remain inactive for extended periods, even in the presence of new information, before issuing discrete and substantial updates. We also show that forecast revisions are *strategic*: the distance to the consensus influences revisions, with forecasters closer to the consensus revising less frequently and, when they do, aligning more closely with it. These patterns are difficult to reconcile with frictionless models of belief formation, stressing the importance of adding such features in shaping survey forecasts.

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. We calibrate parameters capturing the relative weights of accuracy, stability, and strategic motives to match the cross-sectional moments of forecast revisions in the survey data, including frequency, size, and hazard rate.

The model replicates the empirical patterns of lumpiness and consensus alignment, providing a foundation for understanding more subtle features of forecast behavior. A notable example is the pattern of forecast overreaction, where revisions appear too large relative to recent information. We also document this pattern in our data, consistent with findings in other survey datasets. While behavioral models have emphasized psychological drivers of overreaction—such as diagnostic expectations (Bordalo, Gennaioli, Ma and Shleifer, 2020) or overconfidence in private signals (Broer and Kohlhas, 2024)—other accounts attribute it to information frictions (Gemmi and

Valchev, 2023) or measurement error (Juodis and Kučinskas, 2023). We offer a complementary perspective: rational inaction. Infrequent revisions compress the response to multiple previously ignored signals into a single update, amplifying the adjustment and producing the appearance of overreaction. In this sense, our model complements alternative explanations, highlighting how forecast frictions can produce exaggerated responses even under Bayesian updating.

Inspired by these results, we propose a two-stage correction procedure to refine survey forecasts and recover underlying beliefs. The goal is to purge distortions arising from infrequent updates and coordination motives. The first stage isolates active adjustments by focusing on non-zero revisions ("reset forecasts"). The second stage removes the influence of the consensus through a regression-based correction. This procedure generates less biased and more volatile forecasts, while reducing the measured overreaction, yielding a more accurate measure of beliefs. We apply our proposed procedure to the Survey of Professional Forecasters (SPF), and show that our results remain valid. The resulting cleansed forecasts offer a sharper lens into expectations and enhance the informational value of surveys, particularly useful for practitioners and policymakers seeking measures untainted by lumpiness or strategic concerns.

We conclude by presenting additional evidence supporting our interpretation of frictions. First, the decreasing hazard rate of forecast revisions, the large size of updates, and the selective adjustment of long- vs short-term forecasts all favor a story of costly forecast stability rather than inattentiveness or limited information processing. This differentiates our model from alternatives such as sticky information (Mankiw and Reis, 2002; Reis, 2006a,b), rational inattention (Sims, 2003; Maćkowiak, Matějka and Wiederholt, 2023), and observation costs (Alvarez, Lippi and Paciello, 2011, 2016). Second, we uncover systematic heterogeneity across institutional types, including banks, investment firms, consultancies, and academia. These differences align with varying reputational incentives, regulatory environments, and communication practices, reinforcing the idea that strategic and organizational frictions shape forecast behavior. Our findings are related to disagreement, heterogeneity, and attention allocation in forecasting (Capistrán and Timmermann, 2009; Ahn and Farmer, 2024; Giacomini, Skreta and Turen, 2020; Boccanfuso and Neri, 2024).

**Contributions** Our paper contributes to the literature on macroeconomic expectations along three dimensions: empirically, by documenting new facts on frictions in professional forecasts; theoretically, by developing a model of lumpy and strategic forecasting; and methodologically, by proposing a correction procedure that recovers underlying beliefs.

Empirically, we complement a growing body of evidence on lumpy forecast behavior across countries and institutional settings. We document novel patterns in U.S. inflation forecasts, building on related findings from the ECB's survey (Andrade and Le Bihan, 2013), the Brazilian Focus survey (Gaglianone, Giacomini, Issler and Skreta, 2022), and firm-level expectations data (Born, Enders, Müller and Niemann, 2023). Our results on strategic revisions echo concerns about reputational incentives (Marinovic, Ottaviani and Sørensen, 2013) and relate to recent evidence on strategic complementarities in price setting (Karadi, Schoenle and Wursten, 2024). In addition, our results on forecast overreaction complement evidence from surveys (Bordalo, Gennaioli, Ma and Shleifer, 2020) and experimental data (Afrouzi, Kwon, Landier, Ma and Thesmar, 2023).

Theoretically, our model builds on the menu cost literature in price setting (Barro, 1972; Golosov and Lucas, 2007), adapting its logic to a forecasting problem. As in Bec, Boucekkine and Jardet (2023), we interpret the menu cost as the cost of revising the forecast. Just as firms adjust prices intermittently due to fixed costs, forecasters revise predictions selectively, balancing accuracy against the costs of revision. Another central feature is that forecasters adjust their predictions to align with the forecasts of others due to reputational considerations (Morris and Shin, 2002; Ottaviani and Sørensen, 2006). Finally, the model features Bayesian uncertainty dynamics (Baley and Veldkamp, 2025). By merging these different elements, our framework generates a two-dimensional inaction region as in multi-product price-setting models (Midrigan, 2011; Álvarez and Lippi, 2014); it resembles a stopping-time mean-field game (Lasry and Lions, 2007; Alvarez, Lippi and Souganidis, 2023); and generates a decreasing hazard rate typical of learning models with inaction (Baley and Blanco, 2019; Baley, Figueiredo and Ulbricht, 2022).

A central modeling innovation is the use of a restricted perceptions equilibrium to overcome the computational complexity of rational expectations in heterogeneous-agent models (Moll, 2024). Our model's rational expectations equilibrium is infeasible, including aggregate shocks, heterogeneity due to private signals, and lumpy adjustments. The entire forecast and belief distribution determines the current and future consensus needed to make individual forecasting choices. Our approach, inspired by the internally rational framework of Marcet and Nicolini (2003) and Adam and Marcet (2011), assumes forecasters treat the consensus as an exogenous process, bypassing the need to model higher-order beliefs explicitly while preserving key empirical features.

Methodologically, we connect to work on "resetters" in inflation and pricing. Our focus on non-zero forecast revisions parallels "reset inflation" measures (Bils, Klenow and Malin, 2012; Blanco and Cravino, 2020), "frictionless" pricing indices (Bandeira, Castillo-Martínez and Wang, 2024), and the "reset uncertainty" of price adjusters (Afrouzi, Flynn and Yang, 2024). These connections motivate our two-step correction procedure, which uses revision timing and consensus proximity to isolate structural belief changes.

Our findings provide a deeper understanding of how forecasting frictions extend beyond professional economic forecasting. In corporate finance, firms often do lumpy forecasts to signal stability and enhance credibility with investors, for instance, by discontinuing quarterly earnings guidance or strategically timing financial disclosures to manage market expectations (Chen, Matsumoto and Rajgopal, 2011). Surveys of CFOs further emphasize the importance of maintaining stable forecasts to build trust among stakeholders (Graham, Harvey and Rajgopal, 2005). Beyond economics, meteorologists adopt a lumpy approach to updating weather predictions to preserve their credibility, as frequent back-and-forth adjustments can create confusion among users (Griffiths, Marzocca and Michaelides, 2019).

# 2 The Anatomy of Inflation Forecasts

We begin by describing the data sources and the fixed-event forecasting framework. Survey participants are asked for their predictions about the year-on-year percentage change in the Consumer Price Index (CPI) in the United States, measured at the *yearly* frequency, that is, they are asked for the average annual inflation rate. Hence, we construct its actual counterpart to compute forecast errors. We then document the evolution of forecast revisions and errors as the forecasting horizon changes, showing that forecast revisions are lumpy, strategic, and prone to overreaction.

## 2.1 Inflation

For a given year $t$, let $cpi_j$ denote the CPI observed $j$ months before the end of year and define the average CPI in year $t$ as $\overline{cpi}_t = \frac{1}{12}\sum_{j=0}^{11} cpi_j$. The annual inflation $\pi_t$ rate is computed as the percentage difference in the average CPI across consecutive years, approximated with log differences:

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

Following Patton and Timmermann (2010) and Giacomini, Skreta and Turen (2020), we approximate annual inflation as the sum of year-on-year monthly inflation rates, $x_{m,t}$, defined as the log difference in the CPI in month $m$ relative to the same month of the previous year:

$$x_{m,t} = \frac{\log(cpi_{t,m}) - \log(cpi_{t-1,m})}{12}, \quad \forall m = 1, \ldots, 12. \tag{2}$$

This yields the following approximation, which holds particularly well in our sample:[1]

$$\pi_t \cong \sum_{m=1}^{12} x_{m,t}. \tag{3}$$

Figure Ia plots the time series of year-on-year monthly inflation $x_{m,t}$. Figure Ib plots the series of annual inflation $\pi_t$ (solid black line) and the partial sum of year-on-year monthly inflation $\sum_{m=1}^{M} x_{m,t}$ up to month $M$ in year $t$ (in red markers). By the relationship in (3), the two series coincide in the final month of each year ($M = 12$). In our sample, annual inflation averaged 1.8%, ranging from a low of $-0.3\%$ in 2009 to a high of 4.7% in 2021.

---

[1] Appendix A.1 provides theoretical and empirical support for this approximation. Specifically, the approximation holds when two consecutive years exhibit similar within-year CPI dispersion, a condition met in our sample. The resulting discrepancy between actual annual inflation (either in percentage differences or in logs) and the approximation with monthly rates is negligible throughout the analysis period.

**(a)** Year-On-Year Monthly Inflation

**(b)** US Annual Inflation

**Figure I** – US Annual Inflation and Year-On-Year Monthly Inflation. CPI inflation in the US for 2008-2021. Panel (a) plots the series of year-on-year monthly inflation $x_{m,t}$. Panel (b) plots the series of the annual inflation $\pi_t$ (in a solid black line) and the partial sum of year-on-year monthly inflation $\sum_{m=1}^{M} x_{m,t}$ up to month $M$ in year $t$ (in red dots). Recall that for $M = 12$, the cumulated sum coincides with annual inflation, i.e., $\pi_t \cong \sum_{m=1}^{12} x_{m,t}$.

## 2.2 Survey Forecasts

We analyze CPI annual inflation forecasts from Bloomberg's Economic Forecasts (ECFC) survey of professional forecasters. This survey is comparable to other surveys of professional forecasters in terms of the number of participants and their institutional backgrounds, including banks, financial institutions, consulting firms, and universities. In Section 2.9, we provide a comparative analysis using the Survey of Professional Forecasters (SPF) from the Philadelphia Federal Reserve Bank.[2]

**Sample** Our primary analysis focuses on the period from 2010 to 2019, a decade characterized by relatively stable inflation dynamics. This allows us to document novel empirical facts without the confounding effects of extreme macroeconomic volatility.[3] For each year, we include survey participants who provide inflation forecasts for all 12 months preceding the release of the official annual inflation figure. We exclude forecasters who do not report at least one annual inflation revision. This selection yields a panel of approximately 100 forecasters per year. The resulting dataset comprises 9,562 observations, capturing a 10-year history of monthly forecast revisions per participant.[4]

---

[2] See Giacomini, Skreta and Turen (2020) for a detailed comparison of the Bloomberg and other surveys.

[3] Appendix A.2 reports stylized facts for an extended sample covering 2008-2021. The evidence still indicates a robust presence of lumpy forecast revisions, consistent with the patterns documented for the 2010-2019 period, which is the primary focus of this paper. In related work (Baley and Turen, 2025), we examine episodes of heightened inflation volatility, such as the Great Recession (2008-2009) and the COVID-19 pandemic (2020-2021), where different mechanisms may prevail.

[4] While we observe the exact date of each forecast revision, we analyze a monthly frequency, as intra-month updates are rare. We use the forecast available on the last day of each month to construct our monthly panel.

**Survey Features and Incentives** The Bloomberg ECFC survey offers several features that make it especially well-suited for studying forecasting frictions. First, it allows multiple revisions within a given month, so there is no restriction on how frequently participants can revise. Second, the survey is not anonymous: users can track each institution's forecast history. Third, the Bloomberg terminal displays in real time the most recent forecasts from other participants, the date of the last update, and the consensus (average) forecast.[5]

These features create a highly transparent environment that may shape forecasters' incentives. Public visibility of forecasts—both one's own and others'—can generate reputational concerns, especially for financial analysts whose public forecasts are closely tied to their trading behavior (Croushore, 1997). Empirical evidence from Bahaj, Czech, Ding and Reis (2023) confirms that portfolio decisions closely follow published forecasts, highlighting the importance of credibility. In addition, Bloomberg occasionally highlights top-performing forecasters in its quarterly summaries, further reinforcing incentives to maintain accuracy and avoid appearing erratic.

## 2.3 Fixed-Event Forecasting

We consider a fixed-event forecasting setting, which leverages the rich information available in high-frequency revisions of forecasts for a variable observed at a lower frequency. In our case, this involves monthly revisions to annual inflation forecasts.

The fixed event is the annual inflation $\pi$. Let $f_h^i$ be the forecast of forecaster $i = 1, \ldots, N$ at horizon $h = 12, \ldots, 1$. To save notation, we abstract from the year subindex. 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.

Year-on-year monthly inflation rates ($x_h$) are publicly observed with a delay, typically published between the second and third week of the following month. Given the fixed-event scheme, the forecast consists of a "sunk" component given by the sum of observed past realizations $\sum_{j=h+1}^{12} x_j$ and a "projection" component $\mathcal{P}_h^i$ for the remaining horizons:

$$f_h^i = \underbrace{\sum_{j=h+1}^{12} x_j}_{\text{past realizations}} + \underbrace{\mathcal{P}_h^i}_{\text{projection}}, \qquad h = 12, \ldots, 1. \tag{4}$$

Below, we use this decomposition to motivate the presence of forecasting frictions.

---

[5] Appendix A.3 shows a snapshot of the Bloomberg terminal.

## 2.4 Forecasts Revisions and Forecast Errors

**Forecast revisions** For any year, we define the forecast revision at horizon $h$, denoted by $\Delta f_h^i$, as the difference between forecasts at consecutive horizons: $\Delta f_h^i \equiv f_h^i - f_{h+1}^i$. Since forecasts are in percentages, revisions are measured in percentage points. The average revision is very close to zero, indicating a symmetric and low-volatility period of analysis, where upward and downward revisions offset each other. The average revision size (in absolute value and excluding zeros) is 25 basis points. On average, forecasters revise five times per year (an adjustment frequency of 0.43), implying that forecasts remain constant for approximately 1.6 months, even if relevant information becomes available every month.

**Forecast errors** For any year, we define the ex-post forecast error $e_h^i$ at horizon $h$ as the difference between the actual annual inflation and the forecast: $e_h^i \equiv \pi - f_h^i$. Forecasters make minor errors on average $\mathbb{E}[e] = -0.05$, but tend to overpredict inflation, as reflected in the negative sign. The mean squared forecast error (MSFE) is 26 basis points.

## 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 $h$. Figures IIa and IIb show that the size of non-zero revisions and the mean squared forecast error decrease as the horizon shrinks (red solid lines). These decreasing term structures are unsurprising, as the accumulation of information mechanically drives them down. As the fixed event approaches and more information accumulates, predictions become more precise, and thus revisions are smaller.

What is surprising, however, is the speed at which this happens. To see this, imagine a "naive" random-walk automata that applies the last observation $x_{h+1}$ for the remaining horizons; its projection is $\mathcal{P}_h^a = hx_{h+1}$. We illustrate the automata's revisions and errors as the dashed black lines in Figure II. Under the random-walk automata, the size of non-zero revisions and mean squared errors decline significantly faster toward zero relative to the data. These patterns suggest that professional forecasters entertain motives beyond accuracy. We explore two 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 not adjusted continuously. If they remain unchanged for several periods, then they must undergo a significant adjustment to catch up. To explore this possibility, we investigate the revision frequency and hazard rate.

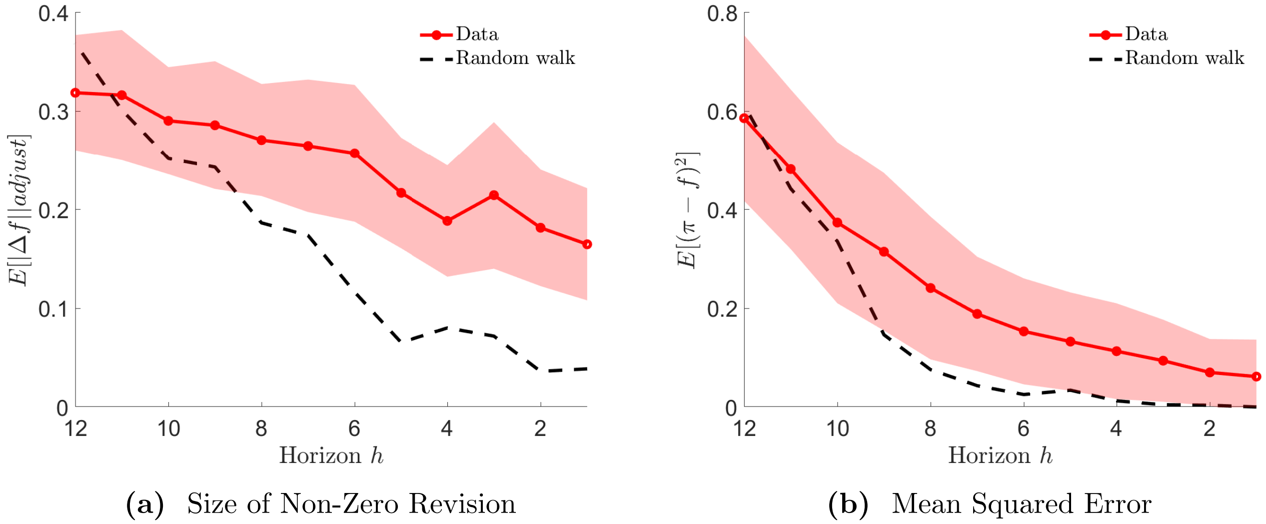

**(a)** Size of Non-Zero Revision **(b)** Mean Squared Error

**Figure II** – Term Structure of Forecast Revisions and Errors . Bloomberg data from 2010-2019. Panel (a) plots the absolute value of non-zero revisions, and Panel (b) plots the mean squared forecast error. Data in red solid lines; naive random-walk automata in dashed black lines. Error bands are computed across forecasters and years, and correspond to 10% confidence intervals.

Figure IIIa shows the unconditional probability of updating a forecast. On average, only 43% of forecasters update their predictions throughout the year. The share of updaters drops as the fixed event approaches. The increasing inaction is puzzling, as relevant information accumulates that could be used to improve the accuracy of the prediction further. Figure IIIb plots the revision hazard $h(age)$, given by the probability of a revision conditional on a forecast's age (i.e., the time elapsed since it was last updated). We estimate the hazard controlling for observed heterogeneity across forecasters and years. The hazard rate is downward sloping: newly set forecasts are more likely to be revised than older ones. For example, the probability of revising a one-month-old forecast is around 0.5; this falls below 0.3 after six months and reaches 0.1 for forecasts that are eleven months old. Importantly, this decreasing hazard is not an artifact of selection bias due to a heterogeneous population of forecasters.[6] Later in the paper, we utilize the decreasing hazard rate to inform learning dynamics and distinguish between alternative models of information frictions.

**Rounding** Participants in the Bloomberg survey report their forecasts up to one decimal point. Could rounding artificially generate inaction? To assess the role of rounding, in Appendix A.4 we use another survey of professionals, Consensus Economics, in which participants report forecasts to three decimal places, allowing us to construct counterfactual revision frequencies at various levels of rounding. Rounding naturally reduces the frequency of adjustments (e.g., a revision below two decimal places is lost when rounded to one decimal place). Still, it does not significantly alter the forecasting behavior across the term structure.

---

[6] The hazard is downward sloping even when conditioning on the number of revisions or the forecaster's type. Thus, the decline in the revision probability likely reflects genuine duration dependence, not composition effects.

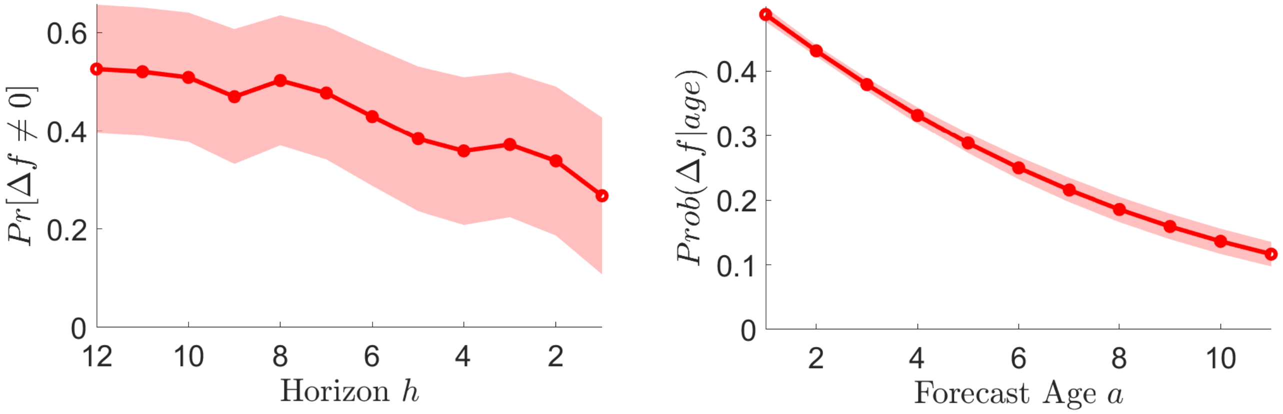

**(a)** Revision Frequency

**(b)** Hazard Rate

**Figure III** – Term Structure of Forecast Revisions. Results computed using Bloomberg data from 2010-2019. Panel (a) shows the frequency of non-zero revisions, and Panel (b) shows the revision hazard rate. Error bands are computed across forecasters and years, and correspond to 10% confidence intervals.

## 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 past mistakes. Alternatively, forecasters may gain from differentiating the average to compete in the market for forecasting expertise. To assess the role of the consensus forecast in shaping forecasting decisions, we adopt the empirical strategy from Campbell and Eden (2014) and Karadi, Schoenle and Wursten (2024), who test for strategic complementarities in firms' price-setting decisions by examining deviations from the average.

**Consensus gaps** To study the potential role of strategic concerns, we define the *consensus forecast* as the average forecast across the $N$ participants at each horizon $h$: $F_h \equiv \frac{1}{N}\sum_{i=1}^{N} f_h^i$. Then, we construct consensus gaps $c_h^i$ as individual forecasts at horizon $h+1$ minus the consensus forecast at horizon $h$: $c_h^i \equiv f_{h+1}^i - F_h$. Since forecasters observe the consensus in real-time in the Bloomberg terminal, it is part of 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 extreme bins include gaps below $-1.3\%$ or above $1.3\%$.

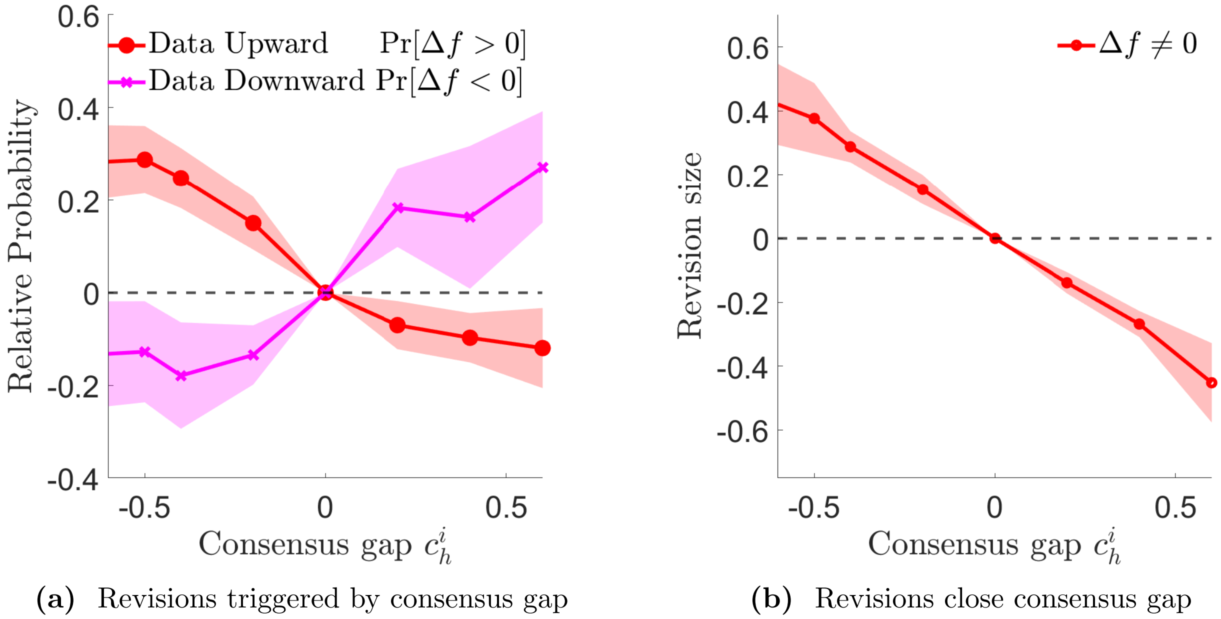

**(a)** Revisions triggered by consensus gap **(b)** Revisions close consensus gap

**Figure IV** – Consensus Triggers Revisions . Bloomberg data between 2010-2019. Panel (a) displays the estimated coefficients of regression (5), and Panel (b) displays the estimated coefficients for the size of revisions (conditioning on updaters). Standard errors are robust and clustered at the time, horizon, and forecaster levels. Confidence bands at 95%.

**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 the consensus gap $c_h^i$ falls within bin $b$:

$$\Pr[\Delta f_{t,h}^i \neq 0] = \beta_0 + \sum_{b=1}^{B} \beta_b \mathbb{1}(c_{t,h}^i(b)) + \alpha_i + \alpha_h + \alpha_t + \epsilon_{t,h}^i. \tag{5}$$

Estimating different coefficients for each bin captures non-linearities in the relationship between the extensive margin and the gaps. We run separate regressions for upward and downward revisions to account for potential asymmetries. We include forecaster ($\alpha_i$), horizon ($\alpha_h$) and year ($\alpha_t$) fixed effects. The year-horizon fixed effects absorb the actual inflation realizations, enabling us to separate strategic concerns from the correlation between the consensus and inflation. Coefficients are robust when including the cumulative inflation in year $t$ up to horizon $h$ as a control.

Figure IVa plots the estimated coefficients associated with each dummy, showing the effect of the consensus gap on the revision probability relative to the omitted category (the middle bin $[-0.1\%, 0.1\%]$). As the gap increases, the probability of a revision rises; 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 probability of revising upward increases, while the likelihood of revising downward decreases.

**Intensive margin** Conditioning on revisions, we examine how the consensus gap affects the size of revisions. We run a similar specification to equation (5), with the revision $\Delta f$ as the dependent variable. As before, we control for forecaster and year-horizon fixed effects. Figure IVb plots the average revision against the consensus gap $c^i_h$. 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.

**Strategic Concerns or Information Channel?** The strong association between forecast revisions and the distance to consensus is consistent with strategic motives in the form of complementarity. However, it could also be interpreted as reflecting an informational role for the consensus (Angeletos and La'O, 2009; Angeletos and Lian, 2018). The nature of the forecasting task is crucial here: in fixed-horizon settings, uncertainty persists, and aggregation across forecasters may be informative. In contrast, our focus on fixed-event forecasting involves a regular flow of public data that gradually resolves uncertainty, diminishing the value of aggregation and heightening the consensus's role as a reputational focal point. Thus, while we cannot entirely rule out an informational channel, the fixed-event structure and the monthly official data releases reduce its likely relevance. Section 6.5 discusses this further.

## 2.8 Forecasts Exhibit Overreaction

Forecast overreaction stems from analyzing the predictability of forecast errors at the individual level. Our test, following Broer and Kohlhas (2024) and Gemmi and Valchev (2023), extends the work by Bordalo, Gennaioli, Ma and Shleifer (2020) by including the consensus as a regressor. This specification naturally connects with the evidence by emphasizing the role of strategic concerns.

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

$$(6) \qquad \underbrace{\pi_t - f^i_{t,h}}_{\text{error}} = \underbrace{\gamma^h_0}_{\text{bias}} + \gamma^h_1 \underbrace{(f^i_{t,h} - f^i_{t,h+1})}_{\text{revision}} + \gamma^h_2 \underbrace{(F_{t,h} - f^i_{t,h+1})}_{\text{consensus}} + \alpha_i + \epsilon^i_{t,h}$$

The key assumption is that the forecasters' information set contains the revision and the consensus. Thus, the rational Bayesian benchmark implies that individual forecast errors are unpredictable: $\gamma^h_0 = \gamma^h_1 = \gamma^h_2 = 0$. Deviations from rationality result in coefficients that differ from zero. If $\gamma^h_1 > 0$, on average, forecasters underreact to information as positive revisions are associated with forecasts below realizations. If $\gamma^h_1 < 0$, on average, forecasters overreact to their information as positive revisions correlate with the forecast being higher than the actual realization. In turn,

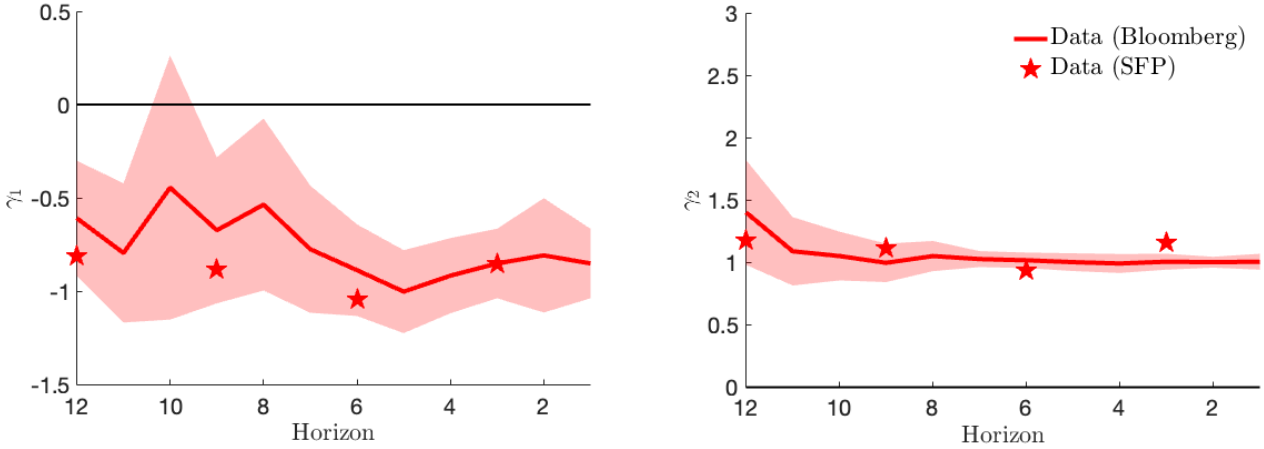

**(a)** Coefficient on Revisions ($\gamma_1^h$)

**(b)** Coefficient on Consensus ($\gamma_2^h$)

**Figure V** – Forecast Rationality Tests. Bloomberg data from 2010-2019. The figures show the estimated coefficients in equation (6). Standard errors are robust and clustered at the time and forecaster levels. Stars ($\star$) indicate the coefficients obtained using the Survey of Professional Forecasters. Confidence bands at 95%.

$\gamma_2^h$ reflects how distance to the consensus affects forecast errors. Deviations from $\gamma_2^h = 0$ indicate that the consensus matters, which should not be the case under rational expectations and in the absence of additional motives beyond accuracy.

Figure V plots the point estimates and confidence intervals (standard errors are robust and clustered by year and forecaster) for all horizons. Panel (a) shows the coefficient on forecast revisions $\gamma_1$, which is negative at all horizons, with an average value of $-0.76$. Panel (b) shows the coefficient on the consensus gap $\gamma_2$, which is positive at all horizons, with an average value of 1.05. According to these results, forecasters tend to overreact to information and respond to the distance from the consensus.

## 2.9 Beyond Bloomberg: Evidence from the SPF

The Bloomberg ECFC survey has two distinctive institutional features—forecasts are publicly attributable, and the platform displays the consensus in real time—which make it an excellent laboratory for detecting strategic behavior and revision costs. Still, they also raise the question of how important the documented distortions are in other environments. How much of the lumpiness, strategic alignment, and overreaction is shaped by these features? To address this point, we show that our results also hold in the Survey of Professional Forecasters (SPF), using fixed-event forecasts data from 1981 to 2024. This survey is quarterly, anonymous, and does not display the consensus in real time. Details are presented in Appendix F.3.

First, with respect to lumpy behavior, inaction remains sizable: the average frequency of revision is 0.8, implying that about 20% of forecaster-quarter observations exhibit no adjustment.

Consistent with the need to catch up at a quarterly frequency, the average size of non-zero revisions and the mean squared forecast error are roughly twice the magnitude in Bloomberg. Importantly, the SPF yields the same qualitative dynamics as the event approaches: we observe systematic declines in the frequency of revision, the size of revisions, and forecast errors with the horizon.

Second, regarding strategic alignment, we construct gaps to the SPF lagged consensus and document the same qualitative relationships between the probability of revision and the gaps. However, the effect is flatter in the SPF: for example, a consensus gap of +0.2 increases the probability of downward revision in 11 p.p. in the SPF, compared with 18 p.p. in Bloomberg. This attenuation is consistent with weaker strategic concerns, given the SPF's characteristics.

Finally, regarding overreaction, Figure V displays the estimated coefficients using the SPF marked with stars at $h = 12, 9, 6, 3$. The coefficients have the same sign and comparable magnitudes as in Bloomberg, and are consistent with existing evidence (Gemmi and Valchev, 2023).

**Taking Stock** 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 automata would. Additionally, forecasts are strategic: the distance to the consensus forecast matters for the extensive and intensive margins of revisions. Finally, forecasters tend to overreact to information and respond to the consensus. This evidence extends beyond the Bloomberg setting, in surveys with different frequencies, designs, and incentives.

# 3 A Structural Model of Lumpy Forecasts

We develop a fixed-event Bayesian forecasting model with revision costs and strategic concerns. We characterize the forecasting policy and the equilibrium, and calibrate the model's parameters by matching moments of the microdata.

## 3.1 Forecasting Problem

Many forecasters, indexed by $i \in [N]$, generate forecasts of annual inflation $\pi$. Annual inflation equals the sum of monthly year-on-year inflation $x_h$, 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 losses:

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

The first term is the distance between the forecast and the realization of annual inflation, reflecting losses from the lack of *accuracy*. The second term is the distance between the forecast and the consensus $F_h = N^{-1}\sum_{i=1}^{N} f_h^i$, multiplied by the parameter $r$ that measures the strength of *strategic concerns*.[7] If $r > 0$, there is strategic complementarity, as the payoff increases when the forecast is close to the consensus; forecasters value alignment with each other. If $r < 0$, there is strategic substitutability, as the payoff increases when the forecast is far from the consensus; forecasters follow a divergent forecasting strategy. The third term is the fixed cost $\kappa > 0$ paid for any forecast revision, capturing *forecast stability*.

We remain agnostic about the values of $r$ and $\kappa$. Section 3.5 utilizes the microdata to inform these parameters, and Section 6 provides further suggestive evidence of these frictions, discussing alternative interpretations and payoff structures.

**Inflation process** Year-on-year monthly inflation follows an autoregressive process:

$$x_h = c_x + \phi_x x_{h+1} + \varepsilon_h^x, \qquad \varepsilon_h^x \sim \mathcal{N}(0, \sigma_x^2), \tag{8}$$

where $c_x$ is the long-run mean, $\phi_x$ is the persistence parameter, and $\varepsilon_h^x$ is an *iid* normally distributed shock with volatility $\sigma_x^2$. The parameters $c_x$, $\phi_x$ and $\sigma_x^2$ are common knowledge.

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

$$x_h^{AR} \equiv \mathbb{E}[x_h | x_{h+1}] = c_x + \phi_x x_{h+1}. \tag{9}$$

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

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

$$\widetilde{x}_h^i = x_h + \zeta_h^i, \qquad \zeta_h^i \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_\zeta^2). \tag{10}$$

Private signals reflect heterogeneity in beliefs or models. We assume that their accuracy $\sigma_\zeta^2$ is identical across forecasters; Section 6.3 introduces heterogeneous noise across forecaster types.[8]

---

[7] We 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.

[8] We do not explicitly include public (correlated) noise in the private signal, as the AR(1) signal serves this role. See Gemmi and Valchev (2023) for an explicit introduction to correlated noise.

**Information dynamics** At the end of the period, and after $f_h^i$ is decided, the monthly inflation $x_h$ and the consensus forecast $F_h$ are observed by everyone. Therefore, there is no long-run uncertainty about their actual values. The individual information set $\mathcal{I}_h^i$ at the time of choosing the forecast is $\mathcal{I}_h^i = \widetilde{x}_h^i \cup \mathcal{I}_h = \widetilde{x}_h^i \cup \{x_{h+1}, x_{h+2}, \ldots, F_{h+1}, F_{h+2}, \ldots\}$. We denote the public information set at horizon $h$ as $\mathcal{I}_h \equiv \{(x_j, F_j) : j \geq h+1\}$, which includes only the releases of past inflation and past consensus (without the private signals).

## 3.2 Belief Formation

Proposition 1 writes the sequential problem in (7) as a function of inflation $\hat{\pi}_h^i$ and consensus beliefs $\hat{F}_h$, using the law of iterated expectations and conditioning payoffs on horizon-specific information.[9]

**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 annual 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. Given* $f_{13}^i$*, forecasters solve:*

$$\min_{\{f_h^i\}_{h=12}^{1}} \sum_{h=12}^{1} \Sigma_h \; + \; (f_h^i - \hat{\pi}_h^i)^2 \; + \; r(f_h^i - \hat{F}_h)^2 \; + \; \kappa \mathbb{1}_{\{f_h^i \neq f_{h+1}^i\}}, \tag{11}$$

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

**Monthly Inflation Beliefs** Forecasters combine the public signal $x_h^{AR}$ in (9) and their private signal $\tilde{x}_h^i$ in (10) to construct an individual monthly inflation belief $\hat{x}_h^i$:

$$\hat{x}_h^i \equiv \mathbb{E}[x_h | \mathcal{I}_h^i] = \frac{\sigma_x^{-2} x_h^{AR} + \sigma_\zeta^{-2} \tilde{x}_h^i}{\sigma_x^{-2} + \sigma_\zeta^{-2}} = (1-\alpha) x_h^{AR} + \alpha \tilde{x}_h^i. \tag{12}$$

The Bayesian weight on the private signal is $\alpha \equiv \sigma_\zeta^{-2} / (\sigma_x^{-2} + \sigma_\zeta^{-2})$; it increases in the private signal precision $\sigma_\zeta^{-2}$ and decreases in the inflation shock precision $\sigma_x^{-2}$.

**Annual Inflation Beliefs** Forecasters form annual 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 $\hat{\pi}_h^i$:

$$\hat{\pi}_h^i = \underbrace{h\left(\frac{c_x}{1-\phi_x}\right) + \frac{1-\phi_x^h}{1-\phi_x}\left(\hat{x}_h^i - \frac{c_x}{1-\phi_x}\right)}_{\text{AR(1) projection using } h \text{ info}} + \underbrace{\sum_{j=h+1}^{12} x_j}_{\text{realized, } j>h}, \quad h = 12, \ldots, 1. \tag{13}$$

---

[9] All proofs appear in Appendix B.

The first part of the expression (13) uses the AR(1) structure to project the monthly belief $\hat{x}_h^i$ into the future. The second part equals the sum of the realized monthly inflation values.

The conditional variance falls with the horizon and is independent of agents' identity:

$$\begin{aligned}\Sigma_h^\pi \quad &= \quad [(1-\alpha)^2\sigma_x^2 + \alpha^2\sigma_\zeta^2]\left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 \; + \\ &\frac{\sigma_x^2}{(1-\phi_x)^2}\left[(h-1) - \frac{2\phi_x(1-\phi_x^{h-1})}{1-\phi_x} + \frac{\phi_x^2(1-\phi_x^{2(h-1)})}{1-\phi_x^2}\right].\end{aligned} \tag{14}$$

The first term of $\Sigma_h^\pi$ corresponds to the uncertainty driven by the AR(1) projection and the noisy signal (weighted by $\alpha$) for the current release of monthly inflation. Likewise, the second part of (14) reflects the accumulated uncertainty caused by the remaining $(h-1)$ unforecastable shocks that will hit the process until the release date.

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

$$z_h \;=\; h\left(\frac{c_x}{1-\phi_x}\right) + \frac{\phi_x(1-\phi_x^h)}{1-\phi_x}\left(x_{h+1} - \frac{c_x}{1-\phi_x}\right) \;+\; \sum_{j=h+1}^{12} x_j, \qquad h = 12,\ldots,1. \tag{15}$$

It is useful to establish the following relationship between individual beliefs $\hat{\pi}_h^i$ in (13) under the information set $\mathcal{I}_h^i$ and public beliefs $z_h$ in (15) under the information set $\mathcal{I}_h$:

$$\hat{\pi}_h^i \;=\; z_h \;+\; \nu_h^i, \quad \text{with} \quad \nu_h^i \sim \mathcal{N}\left(0, \left(\frac{1-\phi_x^h}{1-\phi_x}\right)^2 \alpha^2(\sigma_x^2+\sigma_\zeta^2)\right). \tag{16}$$

In Section 5.1, we use the relationship in (16) to construct our cleansed version of forecasts.

**Consensus Beliefs** The consensus is an endogenous outcome that aggregates optimal forecasts $F_h = N^{-1}\sum_{i=1}^N f_h^i$. However, since it is observed with a one-period delay (e.g., at horizon $h$, $F_{h+1}$ is observed), forecasters must form expectations about the contemporaneous consensus when choosing their forecasts. Forecasters entertain random walk beliefs:

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

where volatility $\sigma_F^2$ is common knowledge.[10] Hence, the common consensus beliefs are $F_h|\mathcal{I}_h^i \sim \mathcal{N}(F_{h+1},\sigma_F^2)$. The definition of equilibrium specifies belief consistency.

---

[10] Allowing the innovation variance of the perceived consensus to vary by horizon, $\sigma_F^2(h)$, is a natural extension consistent with shrinking dispersion at shorter horizons. Still, it would require estimating 11 additional variance parameters, and we already closely match the horizon profiles of revision frequency and size under the constant-variance specification.

## 3.3 Equilibrium

Before we define our notion of equilibrium, let us discuss why the rational expectations equilibrium (REE) is computationally infeasible. Without the adjustment cost, and given quadratic payoffs, we could have conjectured symmetric linear forecasting policies on the signals and solved for the fixed point in higher-order beliefs encoded in the REE, as in Morris and Shin (2002). However, the adjustment cost makes the cross-sectional and time-varying joint distribution of beliefs and forecasts a state variable, as it is needed to predict the current and future consensus. A finite horizon and a finite number of agents do not eliminate the curse of dimensionality.

We focus instead 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. Instead of the whole joint distribution, the individual state-space uses the expected consensus $\hat{F}_h$ implied by the 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.[11]

**Definition 1.** *A restricted perceptions equilibrium (RPE) consists of (i) inflation beliefs* $\{\hat{\pi}^i_h\}$ *and forecasts* $\{f^i_h\}$ *for all agents* $i$ *and horizons* $h$ *and (ii) a perceived consensus process* $\{\hat{F}_h\}$ *given by* $\hat{F}_h = g(\hat{F}_{h+1}, \delta) + \epsilon^{\hat{F}}_h$ *with* $\epsilon^{\hat{F}}_h \sim \mathcal{N}(0, \sigma^2_F)$ *such that:*

- *(a) Given inflation beliefs* $\{\hat{\pi}^i_h\}$ *in (13) and the perceived consensus process* $\{\hat{F}_h\}$ *in (1), forecasts* $\{f^i_h\}$ *are optimal and solve the forecasting problem (11);*
- *(b) Parameters* $(\delta, \sigma^2_F)$ *are such that the forecast errors arising from predicting the actual consensus using the perceived law of motion, i.e.,* $\epsilon^F_h \equiv F_h - g(F_{h+1}, \delta)$*, satisfy:* $\mathbb{C}ov[\epsilon^F_h, \epsilon^F_j] = 0 \ \forall h \neq j$ *and* $\mathbb{V}ar[\epsilon^F_h] = \sigma^2_F$*.*

In the RPE, the actual consensus process given by the aggregation of individual forecasts, $F_h = N^{-1} \sum_{i=1}^N f^i_h$, differs from the prediction. However, in this equilibrium concept, agents use the parameters $(\delta, \sigma^2_F)$ that best predict future consensus given the perceived process.

## 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 encompasses past forecasts, inflation beliefs, and consensus beliefs. We index value functions with the horizon $h$ to account for forecast uncertainty $\Sigma_h$, which

---

[11] The 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).

evolves deterministically and is shared across agents, as well as the past realizations of monthly inflation and the consensus.

**Proposition 2.** *The value of a forecaster with state* $(\hat{\pi}^i_h, \hat{F}_h, f^i_{h+1})$ *equals*

$$\mathcal{V}_h(\hat{\pi}^i_h, \hat{F}_h, f^i_{h+1}) \; = \; \min\{ \underbrace{\mathcal{V}^I_h(\hat{\pi}^i_h, \hat{F}_h, f^i_{h+1})}_{inaction}, \; \underbrace{\mathcal{V}^A_h(\hat{\pi}^i_h, \hat{F}_h)}_{action} \} \tag{18}$$

*where the value of inaction* $\mathcal{V}^I_h$ *and the value of action* $\mathcal{V}^A_h$ *are, respectively,*

$$\begin{aligned} \mathcal{V}^I_h(\hat{\pi}^i_h, \hat{F}_h, f^i_{h+1}) \; &= \; \Sigma_h \; + \; (f^i_{h+1} - \hat{\pi}^i_h)^2 \; + \; r(f^i_{h+1} - \hat{F}_h)^2 \\ &+ \; \mathbb{E}[\mathcal{V}_{h-1}(\hat{\pi}^i_{h-1}, \hat{F}_{h-1}, f^i_{h+1}) | \mathcal{I}^i_h] \end{aligned} \tag{19}$$

$$\begin{aligned} \mathcal{V}^A_h(\hat{\pi}^i_h, \hat{F}_h) \; &= \; \kappa \; + \; \Sigma_h \\ &+ \; \min_{f^i_h} \left\{ (f^i_h - \hat{\pi}^i_h)^2 \; + \; r(f^i_h - \hat{F}_h)^2 \; + \; \mathbb{E}[\mathcal{V}_{h-1}(\hat{\pi}^i_{h-1}, \hat{F}_{h-1}, f^i_h) | \mathcal{I}^i_h] \right\} \end{aligned} \tag{20}$$

*subject to the evolution of inflation beliefs in* (13) *and* (14)*, and consensus beliefs in* (1)*.*

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

$$\mathcal{R}_h \equiv \{(\hat{\pi}^i_h, \hat{F}_h, f^i_{h+1}) : \mathcal{V}^I_h(\hat{\pi}^i_h, \hat{F}_h, f^i_{h+1}) \geq \mathcal{V}^A_h(\hat{\pi}^i_h, \hat{F}_h)\}, \tag{21}$$

and a reset forecast $f^{i*}_h(\hat{\pi}^i_h, \hat{F}_h)$. The current forecast remains unchanged if beliefs lie inside the inaction region and resets at horizon $h$ if beliefs fall outside it. Revisions are given by

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

Panel (a) in Figure VI shows the forecast revision policy at horizon $h = 6$. We plot it in two dimensions by fixing the current forecast at $f^i_6 = 2$ and varying inflation beliefs $\hat{\pi}^i_h$ 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. The inaction region is the negatively sloped dark band centered around the current forecast. Outside the band, the lines correspond to the reset forecast $f^{i*}_6(\hat{\pi}^i_6, \hat{F}_6)$. For example, given the current forecast $f^i_6 = 2$, the beliefs $(\hat{\pi}^i_6, \hat{F}_6) = (2.5, 1.5)$ fall inside $\mathcal{R}_6$ and the forecast remains inactive; instead, the beliefs $(\hat{\pi}^i_6, \hat{F}_6) = (1.5, 2)$ fall outside $\mathcal{R}_6$ and the forecast is reset to $f^{i*}_6 = 1.7$, revising it by $\Delta f^i_6 = 1.7 - 2 = -0.3$.

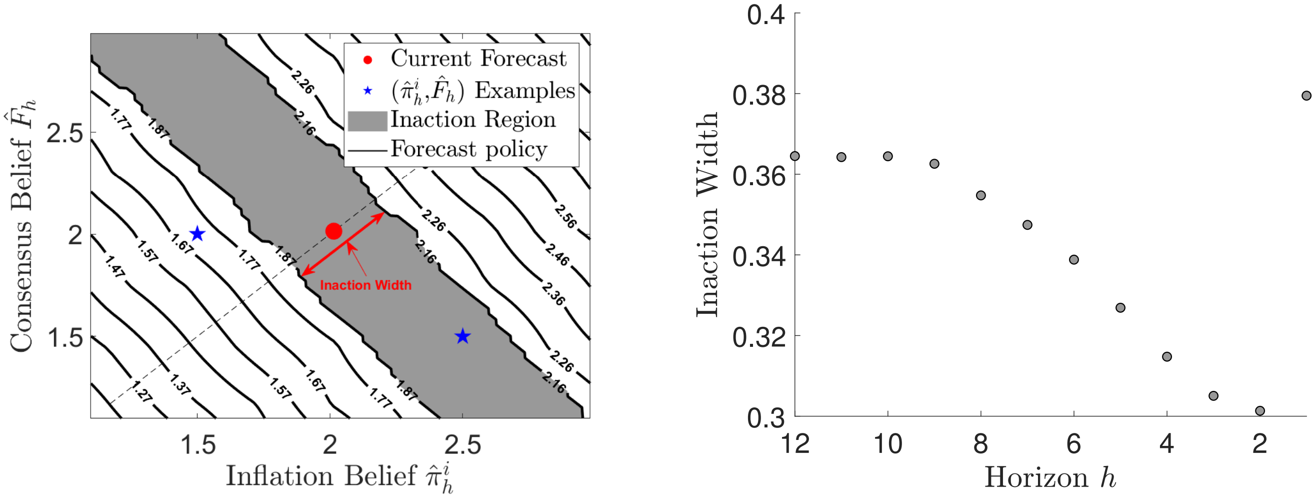

**(a)** Inaction Region and Reset Forecast $f_h^{i*}$

**(b)** Width of Inaction Region

**Figure VI** – Forecast Revision Policy: Inaction and Reset. Panel (a) illustrates the reset forecast $f_h^{i*}$ and inaction region $\mathcal{R}_h$ for $h = 6$ for a current forecast $f_h^i = 2$. We show two examples of beliefs (blue stars), 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. Parametrization from Table I.

Two 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). 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.[12] At long horizons, uncertainty is at its highest level; forecasters anticipate that their belief would hit the band very often and thus optimally widen the band to pay less revision costs (an *option effect*, increasing in $\kappa$). As uncertainty decreases, the option effect becomes smaller, and the band narrows. A shrinking inaction region implies that the revision size falls with the horizon (as documented in Figure IIa). Finally, the strategic concerns determine the slope of the inaction region. If $r = 0$, only accuracy matters and the region is vertical; if $r \to \infty$, only strategic concerns matter and the region is horizontal.[13]

## 3.5 Parametrization

**Externally set parameters** Frequency is monthly. The parameters of the AR(1) process for the year-on-year monthly inflation are estimated using a rolling window.

---

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

[13] Appendix C shows comparative statics on the optimal forecast policy for different values of $\kappa$ and $r$.

**Table I** – Internally calibrated parameters

| **Parameter** | | **Value** | **Moment** | **Data** | **Model** |
|---|---|---|---|---|---|
| $\kappa$ | Revision cost | 0.06 | $\Pr[\Delta f \neq 0]$ | 0.43 | 0.40 |
| $r$ | Strategic concerns | 0.73 | $\mathbb{E}[abs(\Delta f) \mid adjust]$ | 0.25 | 0.19 |
| $\sigma_\zeta$ | Signal noise | 0.03 | Hazard Slope $(h(6) - h(1))/6$ | $-0.04$ | $-0.04$ |
| $\sigma_F$ | Consensus volatility | 0.13 | Internal Consistency | — | — |

The estimates are $(c_x, \phi_x, \sigma_x) = (0.013, 0.932, 0.036)$, which imply an unconditional annual inflation of $\mu_\pi = 12c_x/(1-\phi_x) = 2.23$ with volatility $\sigma_\pi^2 = \sigma_x^2 \sum_{h=1}^{12}(1-\phi_x^h)^2/(1-\phi_x^2) = 0.49$.[14]

**Internally calibrated parameters** Using the simulated method of moments (SMM), we estimate the parameters related to the forecasting frictions $(\kappa, r, \sigma_\zeta)$ by targeting three cross-sectional moments: revision frequency, size, and hazard. Table I presents the parameterization, empirical targets, and model fit. While parameters are identified jointly, we provide a heuristic mapping from moments to parameters to guide intuition. The revision frequency disciplines the fixed revision cost $\kappa = 0.06$; without a fixed cost, the frequency would be one. Given $\kappa$, the size of non-zero revisions disciplines the strategic motive $r = 0.73$, as complementarities shape the responsiveness to private signals. Finally, the slope of the hazard rate (computed between forecast ages 1 and 6) disciplines the private signal noise $\sigma_\zeta = 0.03$. When signals are too noisy, learning is slow, and the hazard rate declines slowly; when signals are informative, learning is fast, and the hazard rate declines more rapidly.[15] Since the private signal noise $\sigma_\zeta$ is close to the public signal noise $\sigma_x$, both signals get a similar Bayesian weight $(\alpha = 0.5)$.

**Consistency of consensus beliefs** Under the internal rationality assumption, the consensus' perceived and actual processes must coincide. This assumption disciplines the value of $\sigma_F$. Starting with a guess for the volatility of the consensus process $\sigma_F$, 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 $\sigma_F^2$ to ensure consistency of beliefs. Setting $\sigma_F = 0.13$ delivers consistent consensus beliefs.[16] As further validation, we run a Dickey-Fuller test and cannot reject the null hypothesis of a random walk when considering a sample of 10 years or less.

---

[14] 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.

[15] Signal noise modulates the slope of the adjustment hazard in Alvarez, Lippi and Paciello (2011), Baley and Blanco (2019), Argente and Yeh (2022), and Baley, Figueiredo and Ulbricht (2022).

[16] Appendix D.1 presents the details on the consistency of consensus beliefs, and Appendix E explains the solution algorithm and other computational details.

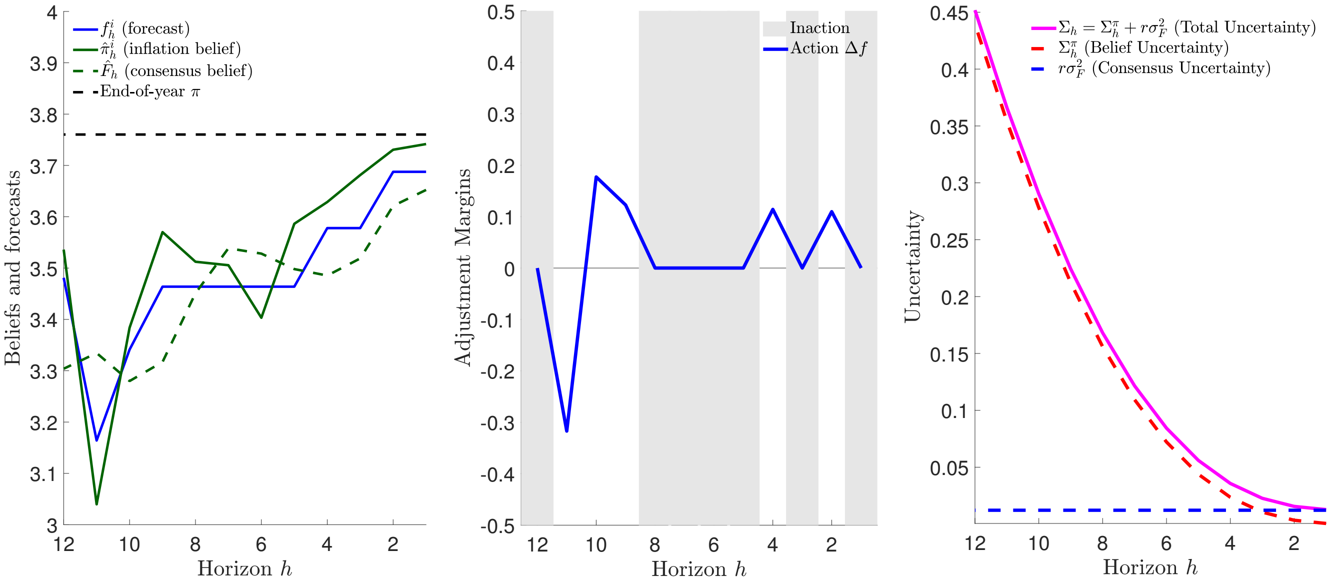

**(a)** Beliefs and forecasts

**(b)** Revisions

**(c)** Uncertainty

**Figure VII** – Forecaster-level dynamics. 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 annual 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 The Model in Action

This section explores various dimensions of the forecasting model. It illustrates the behavior of individual and aggregate beliefs and forecasts, as well as the model's ability to reproduce untargeted moments. We also show how overreaction to information arises as a natural consequence of lumpy forecasts.

## 4.1 Individual Forecast Dynamics

To explain the model's workings, Figure VII illustrates how the beliefs, forecasts, revisions, and uncertainty of a single agent evolve over one year. Panel (a) displays the agent's inflation beliefs, $\hat{\pi}_h^i$ (solid green line), and consensus beliefs, $\hat{F}_h^i$ (dashed 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 aligns with actual inflation, but the forecast does not. Panel (b) 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 $\Delta f$ (dark line), which shrinks with the horizon. Panel (c) shows total uncertainty (solid pink line), equal to the weighted sum of the conditional variance of inflation beliefs $\Sigma_h^\pi$, reaching zero at $h = 1$, and the conditional variance of consensus beliefs $r\sigma_F^2$, which is constant.

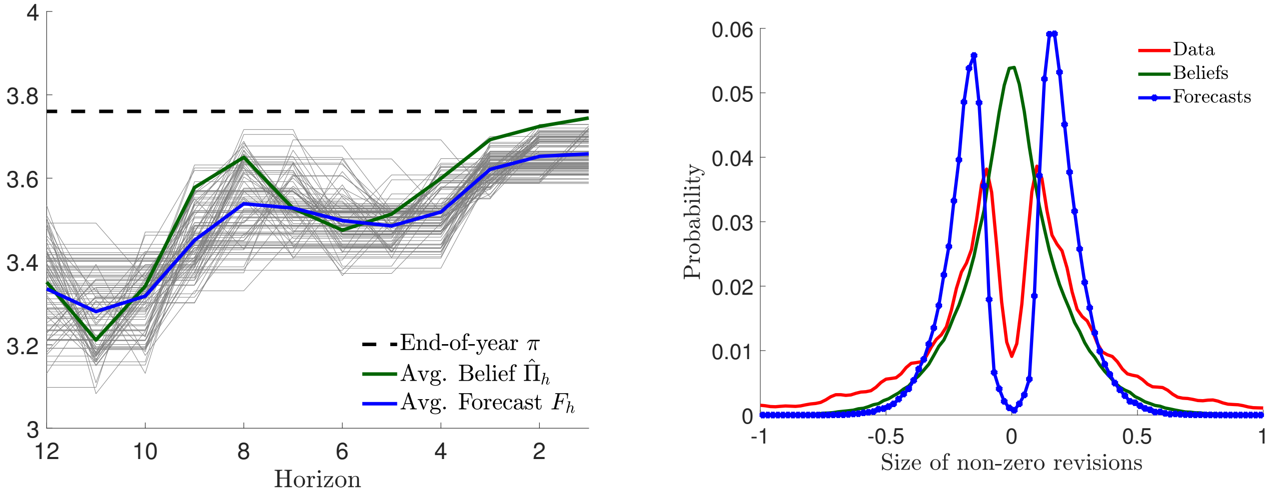

**(a)** Aggregate time series

**(b)** Cross-sectional distributions

**Figure VIII** – Aggregate dynamics. Model simulation for 100 forecasters with baseline parameterization. Sample pools together forecasters, years, and horizons.

## 4.2 Aggregate Forecast Dynamics

Next, we examine the dynamics of aggregate forecasts and aggregate beliefs. Figure VIIIa plots the realized value for the annual inflation $\pi = 3.77$ (horizontal dashed black line). We also show the 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 annual inflation $\pi$ as information becomes available throughout the year. The average belief converges to the actual inflation rate, but the consensus does not. In this example, the consensus remains below the actual inflation rate. 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 VIIIb shows the distribution of *non-zero* forecast revisions $\Delta f_h^i$ 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 $\kappa$.

## 4.3 Untargeted Term Structures

The calibration targeted average moments across horizons. However, the model accurately matches the whole term structure of revision frequency and size. The model also matches the hazard's level despite only targeting its slope (Appendix D.2). We also examine the model's ability to replicate untargeted moments, including the mean squared forecast errors, the distribution of the final revision date, and the effects of the consensus gap on the extensive and intensive margins of adjustment (Appendix D.3).

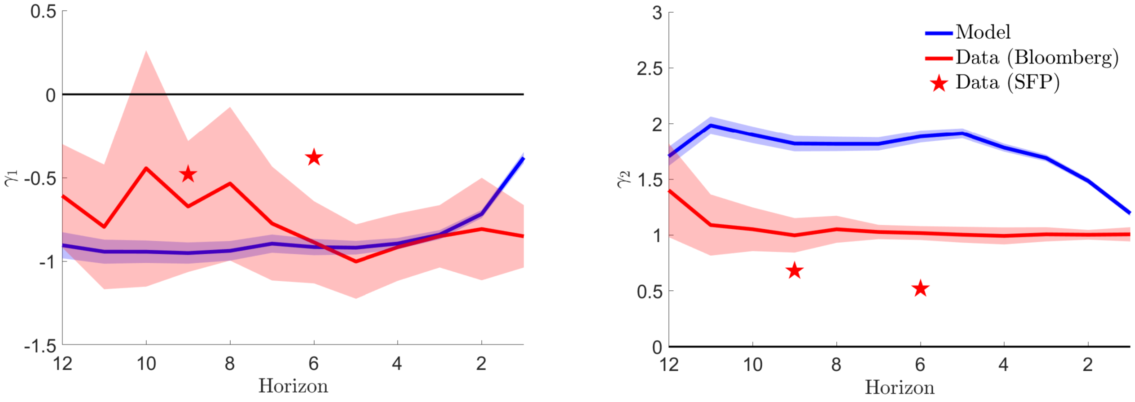

**(a)** Coefficient on Revisions ($\gamma_1^h$) **(b)** Coefficient on Consensus ($\gamma_2^h$)

**Figure IX** – Forecast Rationality Tests. The figures show the estimated coefficients after running equation (6) using Bloomberg data (red line and stars) and model-simulated data (blue line). Confidence bands at 95%.

## 4.4 The Role of Each Friction

We also examine the role of frictions in generating the observed empirical patterns. To do so, we shut down $\kappa$ and $r$ one at a time, recalibrating the model to match a subset of moments (Appendix D.4). Two key takeaways emerge: without fixed costs, agents revise continuously, and the estimated strategic motive flips from positive to negative—incorporating lumpiness shifts the estimated strategic interactions from complementarity to substitution. Removing strategic concerns flattens the hazard rate, making revision behavior less sensitive to forecast age and weakening the model's ability to match empirical dynamics.

## 4.5 Overreaction to Information

We next examine how forecasting frictions can lead to an apparent overreaction to new information. Using model-generated forecasts, we run regression (6). Figure IX shows the results. Forecasts feature (*i*) overreaction to private information, $\gamma_1 < 0$ in Panel (a), and (*ii*) a systematic response to the consensus, $\gamma_2 > 0$ in Panel (b), despite the consensus carrying no additional information. These patterns hold across all horizons. Crucially, these results are untargeted, providing an extra layer of validation for our model.

The measured overreaction to information arises from lumpy behavior. When adjustments occur, they incorporate 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 $\gamma_1$, 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.

# 5 Recovering Underlying Beliefs

To refine survey forecasts as measures of true underlying beliefs, we propose a two-stage procedure that isolates the active component of forecast revisions and removes the influence of the consensus. First, we isolate active adjustments by focusing exclusively on non-zero revisions. Second, we remove the influence of the consensus using a regression-based adjustment. This yields a sharper proxy for underlying beliefs: cleansed forecasts exhibit greater volatility, accuracy, and dispersion, but lower autocorrelation and overreaction.

## 5.1 Cleansing Forecasts

Conditional on resetting, the first-order condition requires forecasts to be a weighted average of the beliefs about inflation and the consensus, plus the expected change in the (non-differentiable) continuation value:

$$f^{i\ *}_{h,t} \;=\; \frac{1}{1+r}(\hat{\pi}^i_{h,t} \;+\; r\hat{F}^i_{h,t}) \;+\; \underbrace{\mathbb{E}[\text{change in } \mathcal{V}_{h-1,t}(\hat{\pi}^i_{h-1,t}, \hat{F}^i_{h-1,t}, f^{i\,*}_h)]}_{\approx 0}\,. \tag{23}$$

The expected change in the continuation value is nearly zero upon reset, as we discuss below. Given $r$, and using the assumption of random walk consensus beliefs $\hat{F}^i_{h,t} = F_{h+1,t}$, we back-out a proxy for individual beliefs "cleansed" from lumpiness and strategic concerns, denoted by $\widetilde{f}^i_{h,t}$, by rearranging expression (23):

$$\widetilde{f}^i_{h,t} \;\equiv\; (1+r)f^{i\ *}_{h,t} \;-\; rF_{h+1,t} \;\approx\; \hat{\pi}^i_{h,t}. \tag{24}$$

We compute the cleansed consensus by averaging the individual cleansed forecasts, which by (24), are approximately equal to the average beliefs $\widetilde{F}_{h,t} \equiv \mathbb{E}[\widetilde{f}^i_{h,t}] \approx \mathbb{E}[\hat{\pi}^i_{h,t}]$.

**Validating Cleansing Procedure** To validate our cleansed measure as a proxy for agents' beliefs, we examine potential discrepancies between $\widetilde{f}^i_h$ and $\hat{\pi}^i_{h,t}$ using model-simulated data.[17] We find that the two series are highly aligned (with a correlation of almost one), which reinforces the validity of our cleansing procedure and suggests that the expected change in the continuation value upon reset is negligible.

## 5.2 Recovering strategic concerns with data alone

One could estimate a structural model to recover the parameter $r$ and apply the correction in (24) to proxy beliefs. Alternatively, we propose estimating $r$ directly from the data using the following

---

[17] Appendix F.1 presents the details.

OLS regression based on (23), which uses survey data of resetters:

$$f^{i}_{h,t}{}^{*} = \underbrace{\beta_0}_{0.48\,(0.24)} + \underbrace{\beta_1}_{0.28\,(0.06)} z_{h,t} + \underbrace{\beta_2}_{0.44\,(0.13)} F_{h+1,t} + \alpha_i + \alpha_h + \alpha_t + \epsilon_{i,h,t} \tag{25}$$

We exploit the relationship in (16) to substitute private beliefs $\hat{\pi}^{i}_{h,t}$ for public beliefs $z_{h,t}$ plus an idiosyncratic term captured in the fixed effects. As the estimation relies on the level of predicted inflation, we include forecasters ($\alpha_i$), horizon ($\alpha_h$), and year ($\alpha_t$) fixed effects. In our theory, the $r$ parameter is independent of agents and horizons, so we pool all observations. We cluster standard errors at the horizon-year level (in parentheses). We also include controls to account for potential omitted variables, including lags of annualized inflation, the growth of industrial production, and three-month Treasury yields.

The parameter $r$ can be recovered from $\beta_1$ or $\beta_2$. However, $\beta_1$ 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 $\hat{r}$ from $\hat{\beta}_2 = \hat{r}/(1+\hat{r})$ since we have a direct and observed measure of $F_{h+1}$ in the data. We 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 use the Bloomberg survey, in which the consensus forecast is available to all participants in real time, this does not prevent any researcher from running this regression using other surveys, as long as the lagged consensus is available when agents are asked to provide new forecasts.

Our preferred estimate is $\widehat{r} = 0.79$.[18] It is reassuring that the data-implied $r$ closely resembles the SMM-implied value of 0.73 in Table I, further validating our calibration.[19]

## 5.3 Dynamic Properties of Cleansed Forecasts

We now examine how our correction procedure alters the dynamic behavior of forecasts and how these changes compare to the model's predictions. Table II summarizes three key statistics—forecast volatility, autocorrelation of forecasts, and autocorrelation of revisions—for four series: raw and cleansed forecasts from the data (columns A and B), and raw and cleansed forecasts from the model (columns C and D).

Starting with volatility (first row), we observe an increase when moving from raw to cleansed forecasts. In the data, volatility rises from 0.60 to 0.80, in the model from 0.84 to 0.95. This reflects the fact that raw forecasts are smoothed by two main forces: infrequent updates due to revision costs, and strategic alignment with the consensus. Once these frictions are removed,

---

[18] We run various specifications in Appendix F.2 and the estimated $r$ remains stable across them.

[19] Alternatively, and to capture the continuation value further, we can proxy the fixed effect $v^{i}_{h}$ in (25) using an individual × horizon interaction $\alpha_{i,h}$. Our results change marginally when we run this alternative specification, yielding an estimate of the strategic concerns of $\widehat{r} = 0.71$.

**Table II** – Dynamic Properties of Cleansed Forecasts

| | Data | | Model | |
|---|---|---|---|---|
| **Moment** | (A) Raw | (B) Cleansed | (C) Raw | (D) Cleansed |
| Volatility | 0.60 | 0.80 | 0.84 | 0.95 |
| Autocorr. Forecasts | 0.81 | 0.62 | 0.51 | 0.32 |
| Autocorr. Revisions | -0.08 | -0.19 | -0.12 | -0.23 |

Notes: Bloomberg data for 2010-2019. Model simulation under the benchmark calibration.

forecasts respond more directly to incoming information, resulting in greater volatility.

The second row shows how these frictions affect persistence. In the data, the autocorrelation of the raw forecasts is 0.81 and declines to 0.62 for the cleansed forecasts. This drop reflects more frequent and responsive updating in the cleansed series. In the model, the autocorrelation similarly declines from 0.51 to 0.32, again confirming that the persistence of raw forecasts primarily reflects frictions in reporting rather than sluggish belief formation.

Finally, the third row examines the autocorrelation of forecast revisions. In the data, raw revisions are mildly negatively autocorrelated ($-0.08$), and this negative autocorrelation becomes stronger after cleansing ($-0.19$). The model again mirrors this pattern: raw revisions exhibit autocorrelation of $-0.12$, which deepens to $-0.23$ after cleansing. These negative values are consistent with a stationary inflation process, where revisions reflect transitory surprises that do not persist. In such an environment, an upward revision is often followed by a downward correction, especially as the end-of-year target becomes clearer.

The model's close match to these revision dynamics further reinforces its empirical validity. While the model replicates the patterns observed when moving from raw to cleansed measures, a perfect fit was not expected a priori. These simple statistics may mask substantial unobserved heterogeneity across forecasters, horizons, and years—factors that are more relevant in the data than in the model and that are not fully accounted for by these measures. Our thought experiment assesses whether the model can qualitatively reproduce the patterns uncovered by our cleansing procedure. The model's ability to closely match the autocorrelation of revisions further supports this intuition, as differencing at the forecaster level effectively removes much of the unobserved heterogeneity at the individual level.

## 5.4 Accuracy of Cleansed Forecasts

We now assess the predictive accuracy of cleansed forecasts. At the aggregate level, forecast errors constructed with cleansed forecasts, $\pi - \widetilde{f}_h^i$, are, on average, smaller than those from the raw forecasts, reflecting reduced bias. However, because the cleansed series strips away the consensus component, it is also more volatile. This leads to larger variance in individual predictions and, as a result, greater forecast dispersion and revision size—even when conditioning on updaters.

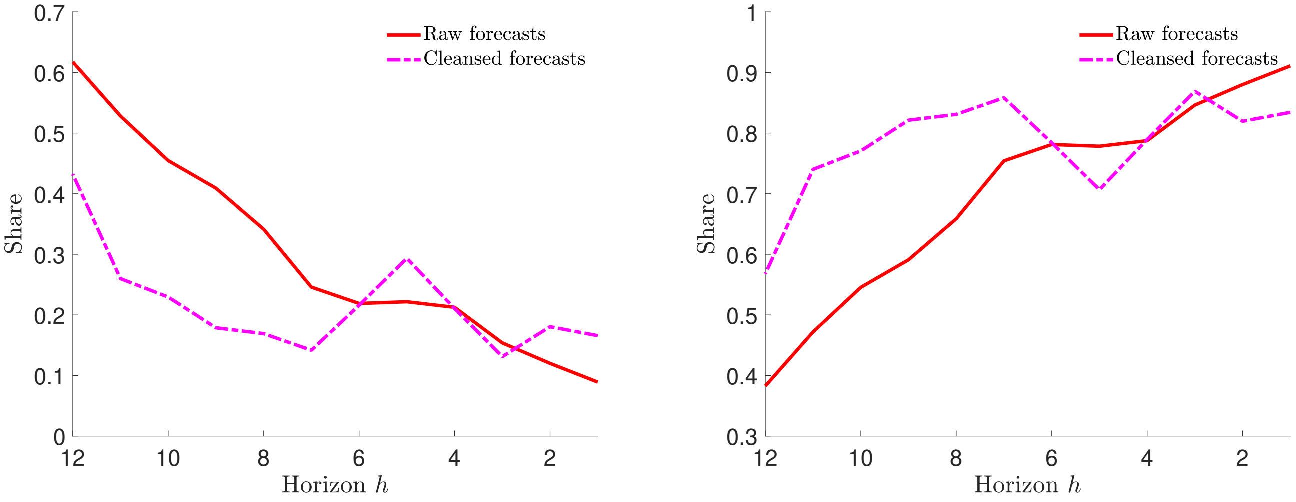

**(a)** Bias Share **(b)** Variance Share

**Figure X** – Bias and Variance Shares of the MSFE. This figure decomposes the mean squared forecast error (MSFE) into its bias and variance components for both raw and cleansed forecasts. Panel (a) shows the bias share, computed as $(\pi - \widetilde{F}_h)^2/\text{MSFE}_h$, while Panel (b) reports the variance share, computed as $\text{Var}(\widetilde{f}_h^i)/\text{MSFE}_h$. Cleansed forecasts exhibit lower bias and higher dispersion, especially at longer horizons.

Overall, the mean squared forecast error across horizons is marginally larger for cleansed forecasts than for raw forecasts (0.37 vs. 0.26).

To assess the trade-off between a lower bias and a higher dispersion, we decompose the mean squared forecast error into bias and dispersion components: $\text{MSFE}_h = (\pi - \widetilde{F}_h)^2 + \mathbb{V}ar(\widetilde{f}_h^i)$ where $\widetilde{F}_h$ is the average cleansed consensus. Figure X plots the average relative contribution of these two components—bias and dispersion—to total MSFE, separately for the original and cleansed forecast series.

As expected, the cleansed forecasts exhibit a substantially lower bias share (left panel), especially at longer horizons, indicating improved forecast accuracy of the cleansed consensus $\widetilde{F}_h$. In turn, the variance share is larger for cleansed forecasts (right panel), reflecting the removal of the common component and the greater dispersion in individual beliefs. These differences are most pronounced for horizons beyond six months, while for $h \leq 6$, the relative contributions of bias and variance remain similar across both forecast series.

This empirical validation in the data complements our earlier test in the model: we confirm that the exact cleaning procedure recovers model-consistent expectations from simulated data. Together, these exercises bolster the case for using cleansed forecasts as a robust diagnostic tool in applied work.

## 5.5 Revisiting Overreaction

Our cleansing procedure offers a refined lens for interpreting overreaction estimates. By conditioning on updates and removing strategic components, the cleansed forecast series provides a closer

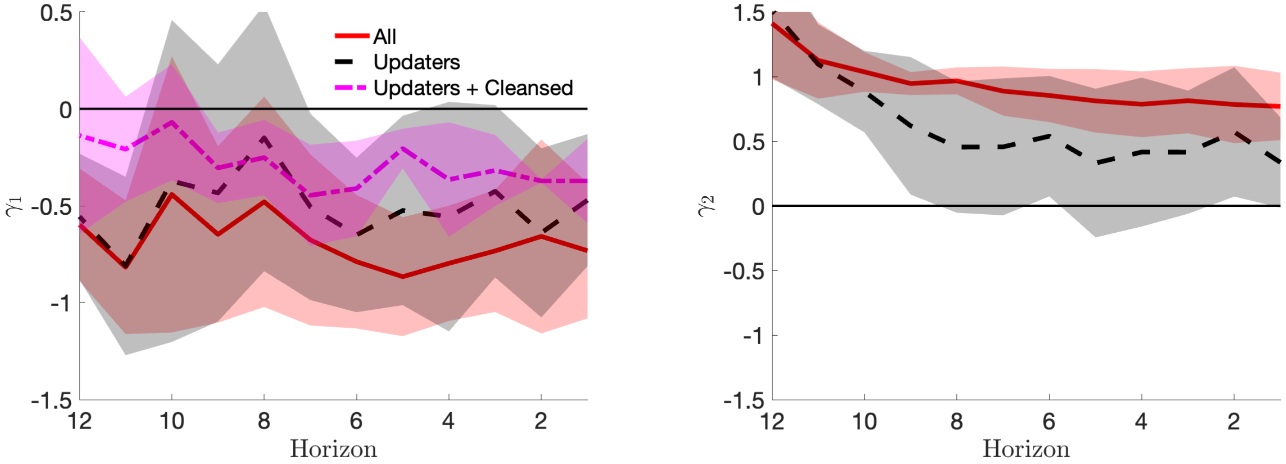

**(a)** Coefficient on Revisions ($\gamma_1^h$) **(b)** Coefficient on Consensus ($\gamma_2^h$)

**Figure XI** – Forecast Rationality: Updaters and Cleansed. The figures show the estimated coefficients after running equation (6) 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. Confidence bands at 95%.

proxy for underlying beliefs. This allows us to revisit the standard overreaction test with a more accurate measure of agents' responsiveness to information.

Figure XI, Panel (a), plots the estimated coefficients $\gamma_1^h$ from regression (6) under three specifications: using all forecasts, using only updates, and using cleansed updates. The black dashed line corresponds to conditioning on updaters ($\Delta f_h^i \neq 0$), while the pink dashed line uses our cleansed belief proxy $\widetilde{f}_h^i$. In both cases, the estimates move closer to the rational expectations benchmark of zero. Panel (b) of Figure XI displays the estimated coefficients $\gamma_2^h$, which capture the responsiveness of errors to the consensus. These coefficients also decline as we move from raw to updaters (naturally, we cannot run the test for the cleansed forecasts).

This refinement yields three key insights. First, the magnitude of overreaction falls sharply when we control for lumpy behavior and strategic incentives. The average coefficient $\gamma_1^h$ declines from $-0.76$ (all forecasts) to $-0.68$ (updates), and further to $-0.29$ (cleansed updates), implying a 58% reduction. This confirms that raw revisions confound true belief updating with frictions and incentives present in the survey environment. Second, the greater variance in $\widetilde{f}_h^i$ increases the informativeness of forecast revisions and raises the precision with which coefficients are ultimately estimated, as reflected in the tighter confidence bands in Figure XI. This is particularly valuable in small samples, where data limitations are acute. Third, despite the reduction, the coefficient remains significantly different from zero. This suggests that some degree of overreaction persists even after removing reporting distortions. In line with recent evidence from Bordalo, Gennaioli and Shleifer (2022), this pattern points to inherent behavioral features in belief formation.

To emphasize that the specific Bloomberg survey does not drive our results, we repeat the proposed two-step cleansing procedure on the Survey of Professional Forecasters and confirm that the coefficient on revisions moves progressively toward zero (see Figure F.13 in the Appendix).

These regressions also serve as a diagnostic for the quality of belief recovery. As the cleansed forecasts yield coefficients closer to zero, the implied expectation errors become more consistent with rational expectations. This convergence supports the interpretation of our cleansing method as not only filtering out frictions, but also recovering a latent belief process that is more internally consistent and economically meaningful.

# 6 Supportive Evidence, Interpretations, and Extensions

Our parsimonious benchmark model assumes constant parameters, equal treatment of forecast errors across horizons, and homogeneous private signals. It abstracts from discounting, heteroskedasticity, and signal heterogeneity, allowing us to isolate the roles of fixed revision costs and strategic motives. Despite its simplicity, the model successfully replicates untargeted features of the data, including the whole term structure of frequency and size of revisions, the hazard rate, and the patterns of overreaction.

Here, we present additional empirical evidence that supports the model's core mechanisms and examine extensions that build upon its structure. We consider richer cost functions, heterogeneous information, institutional differences, and time-varying behavior. We also present suggestive evidence that forecasters exhibit a preference for forecast stability, even when new information is available, shedding light on the reputational and institutional factors behind forecast inaction. These extensions highlight the flexibility of our framework and suggest directions for future work.

## 6.1 Forecast Stability

Fixed costs and the resulting inaction in forecast revisions may arise from various sources and could be interpreted in many ways. A first possibility is that inaction stems from the cost of processing information. However, given the sophistication of professional forecasters, such costs—along with the minor inconvenience of accessing the forecasting system—are unlikely to be prohibitive relative to the benefits of maintaining accuracy.

A second possibility involves the infrequent observation of information, as in models of sticky information and lumpy data flows (Mankiw and Reis, 2002; Andrade and Le Bihan, 2013; Bec, Boucekkine and Jardet, 2023). In such frameworks, forecasters update only occasionally because signals are received infrequently. However, this assumption is at odds with the decreasing hazard rate of forecast revisions observed in our data. If signals arrived with constant probability, as in Calvo-style models, the hazard would be flat. Moreover, since Bloomberg participants can revise forecasts at any point during the month, technological constraints or deadline effects are unlikely explanations for the observed pattern.

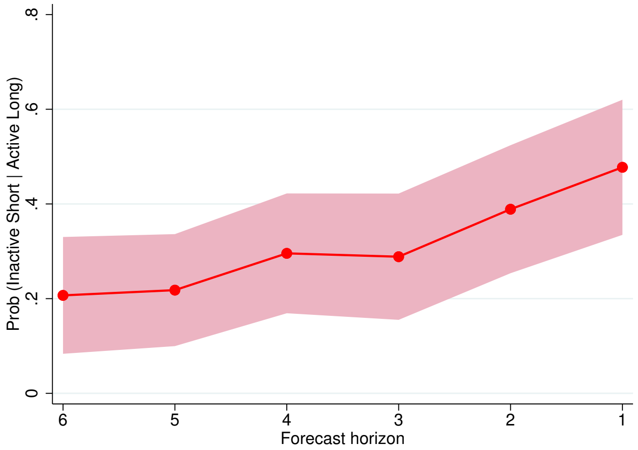

**Figure XII** – Inaction in the short vs. long run. Predictions reported up to 18 months before the release of annual inflation.

A third possibility is that forecasters behave similarly to firms in Matějka (2015), which set prices in a lumpy way due to consumers' information acquisition costs. While a similar mechanism could apply to professional forecasters and their clients, the high sophistication and strong incentives on both sides make this explanation less convincing in our context.

A fourth, more plausible explanation is the timing of macroeconomic data releases—forecasters may delay updates until key variables, such as GDP, are published at lower frequencies than inflation. Accounting for multiple data sources with different release schedules and measurement errors would significantly complicate our model, so we leave this for future research.

Finally, our preferred explanation is that forecasters prefer to signal commitment to their predictions, favoring stability over frequent adjustments and minimizing the reputational burden of appearing erratic. We provide suggestive evidence through selective adjustment of long vs. short-term forecasts. We construct a statistic that measures the overlap between short- and long-term revisions. Long-term (13 to 18 months ahead) forecasts for year $t+1$ overlap with short-term (1 to 6 months ahead) forecasts for year $t$. Given the overlap, we compute the probability that a short-term forecast remains inactive (i.e., $f_h^i = f_{h+1}^i$ for $h = 6, \ldots 1$) conditional on a long-term forecast being revised (i.e., $f_h^i \neq f_{h+1}^i$ for $h = 18, \ldots 13$). If a long-term forecast is revised, it signals that relevant information has been received, prompting the participant to log into the Bloomberg terminal to update the 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 XII plots this statistic. Given that the 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.

## 6.2 Hazard Rate as a Model-Discrimination Device

The shape of the hazard rate provides a powerful lens to distinguish between competing models of forecast inaction. In our data, the hazard rate of forecast revisions is decreasing with age: the longer a forecast has been in place, the less likely it is to be updated. The negative slope serves as a discrimination device across different models.

Consider first a benchmark model with no revision costs, in which forecasters are inattentive and update their forecasts randomly with a constant probability over time, as in Andrade and Le Bihan (2013) or the classic Calvo (1983) model. These models predict a flat hazard rate, since the likelihood of revision does not depend on the forecast's age. By contrast, models with fixed costs but no learning—where uncertainty is constant—imply an increasing hazard rate: a newly set forecast starts at the center of the inaction region and only gradually reaches the threshold for revision. In that setting, the longer a forecast has been in place, the more likely it is to be revised.

Neither of these polar cases aligns with our findings. The decreasing hazard we observe is instead consistent with models that feature both fixed costs and Bayesian learning (Baley and Veldkamp, 2025). In such models, agents update more aggressively early in the year when uncertainty is high and news is more informative, but update less as their uncertainty declines and their beliefs stabilize over time. This same pattern arises endogenously in our structural model, consistent with the empirical evidence.

## 6.3 Heterogeneity

Our baseline calibration assumes that agents have identical fixed revision costs, share the same strategic concerns, and have symmetric private information. Yet empirical patterns suggest systematic differences in revision behavior across four forecaster types: (*i*) financial institutions, (*ii*) banks, (*iii*) consulting companies, and (*iv*) universities and research centers.

Table III (in columns labeled Data) documents these differences. The most notable contrast emerges between consulting firms and universities: for instance, consulting firms revise their forecasts roughly 30% more frequently than universities, with non-zero revisions that are 40% smaller. We interpret these patterns as evidence of forecasters' ex ante heterogeneity in their institutional type. Such heterogeneity may stem from internal decision rules, reputational concerns, or variation in perceived information quality.

To explore these differences more formally, we recalibrate the model separately for each forecaster type. Table III (in columns labeled Model) reports the model fit. Table IV reports the parameter estimates relative to financial institutions (we normalize the values for financial institutions to unity to facilitate comparison).

The calibrated models indicate that universities face higher revision costs and display the weakest strategic concern for consensus alignment. In contrast, consulting firms face the lowest fixed costs and the most significant strategic concerns. Banks exhibit intermediate behavior, with

**Table III** – 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: Bloomberg data between 2010-2019 and model-generated moments under alternative calibrations.

**Table IV** – Calibration by forecaster type, relative to financial institutions

| **Parameter** | **Financial Inst.** | **Banks** | **Consulting** | **Universities** |
|---|---|---|---|---|
| $\kappa$ | 1.00 (0.06) | 1.08 | 0.94 | 1.29 |
| $r$ | 1.00 (0.81) | 0.62 | 0.89 | 0.50 |
| $\sigma_\zeta$ | 1.00 (0.04) | 1.16 | 1.14 | 2.28 |
| $\sigma_F$ | 1.00 (0.10) | 1.13 | 1.08 | 1.33 |

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

moderate levels of signal noise and costs. These differences reinforce the idea that expectations formation varies systematically across institutional settings, reflecting both informational, stability, and strategic asymmetries.

More broadly, heterogeneity is unlikely to be confined to institutional categories. Even within groups, we observe dispersion in revision frequencies and response profiles that point to underlying variation in signal precision or internal belief structures. The literature has highlighted the importance of modeling such within-type heterogeneity: Patton and Timmermann (2010) and Ahn and Farmer (2024) document and decompose forecast disagreement among experts, while models such as Giacomini, Skreta and Turen (2020) allow for variation in agents' priors, signal updating, or decision rules. These approaches offer promising directions for extending our framework.

## 6.4 Alternative cost structures

In our model, agents face a fixed cost whenever they revise their forecast, regardless of the magnitude of the change. This assumption is crucial for generating lumpy adjustment behavior, characterized by long periods of inaction punctuated by sizable revisions. An alternative approach would be to introduce costs that vary with the size of the revision, such as $|f_h^i - f_{h+1}^i|$ or its square. These structures may reflect a different perspective on reputational concerns, where large or erratic changes have a greater influence on credibility than gradual sequences of minor revisions.

A model that includes both fixed and magnitude-dependent costs would separate the extensive and intensive margins of adjustment, capturing more nuanced revision behavior. However, such models tend to induce continuous rather than lumpy adjustments, making it harder to match the frequency and size distribution of revisions. The fixed-cost framework provides a parsimonious benchmark that replicates the key patterns in the data. At the same time, more flexible cost specifications remain a promising avenue for future research.

Additionally, the model treats all forecast errors symmetrically, regardless of how far in the future the forecasted event is. Yet in practice, errors in near-term forecasts may carry more reputational or operational weight. Allowing error penalties or revision costs to vary with the horizon, $\kappa_h$, could help explain behaviors such as delayed corrections, last-minute adjustments, or the greater stability of short-term forecasts.

## 6.5 Learning from the consensus

In our framework, the consensus forecast influences forecasters through a strategic motive: agents incur a utility cost when deviating from the consensus but do not treat it as an informative signal. The exclusion of learning from the consensus is motivated by the structure of the forecasting environment (recall the discussion at the end of Section 2.7). Inflation data is released monthly, and each release supersedes prior private and consensus signals. Because the consensus at horizon $h+1$ aggregates beliefs about $x_{h+1}$, its informational content becomes redundant once $x_{h+1}$ is observed. Thus, lagged inflation is a more direct and reliable signal than lagged consensus.

If private signals were longer-lived or actual inflation remained unobserved, the consensus could play a more informative role. A natural extension would therefore endow the consensus with informational content—as an additional noisy signal—while maintaining its payoff relevance through reputational concerns. The downside is that a consensus that is simultaneously informative and payoff-relevant typically brings higher-order expectations about how others interpret the consensus, substantially complicating the analysis.[20] In that richer environment, estimates that abstract from learning would tend to load part of the consensus-response onto the reputational channel.

That said, the empirical relationship between revisions and the consensus gap could reflect both channels. If forecasters use the consensus to infer information held by others, the estimate of the strategic complementarity parameter $r$ may be upward-biased. Two features of our application mitigate this concern. First, the fixed-event structure of the forecasting task, combined with the regular release of official inflation data, likely attenuates the incremental informational role of the consensus. Second, the value $r = 0.73$ obtained from the estimation of the structural model, which incorporates only strategic considerations, is close to the empirical OLS estimate $\hat{r} = 0.79$ reported later in Section 5.1, which could in principle reflect both mechanisms. This similarity is reassuring and supports interpreting $r$ as capturing reputational motives in our setting.

---

[20] See, for example, Nimark (2008, 2014) or Hellwig and Venkateswaran (2009) for tractable frameworks that limit higher-order beliefs.

# 7 Final Thoughts

Professional forecasters revise their predictions in a lumpy manner, influenced by noisy information, adjustment costs, and strategic considerations. Our findings have significant implications for the design of expectation surveys. For instance, evidence from the Brazilian FOCUS survey (Gaglianone, Giacomini, Issler and Skreta, 2022) shows that forecast accuracy and revision frequency increase around contests rewarding precision, while Ottaviani and Sørensen (2006) highlights how competitive environments influence the differentiation of forecasts. These insights suggest that better-designed incentives could promote more frequent updates and less strategic forecasts, enhancing the reliability of survey expectations.

Our results open several avenues for further research on heterogeneity across time and agents. Regarding time variation, in companion work, we document that periods of high inflation volatility are associated with more frequent and larger forecast revisions, resulting in state-dependent forecast flexibility (Baley and Turen, 2025). Further, exploiting the panel dimension to disentangle aggregate and idiosyncratic news shocks could shed light on the shocks forecasters respond to (Capistrán and López-Moctezuma, 2014).

Turning to cross-sectional heterogeneity, it would be valuable to examine whether households and firms follow similar patterns to those of professional forecasters (Weber, d'Acunto, Gorodnichenko and Coibion, 2022; Born, Enders, Müller and Niemann, 2023). While strategic concerns may be less relevant for these groups, adjustment costs are presumably higher and likely driven by information processing costs. Understanding how these factors shape forecast behavior across different agents, and their macroeconomic implications, remains an open question.

# 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. Working Paper 32541, NBER.

—, 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.

Ahn, H. J. and Farmer, L. E. (2024). *Disagreement About the Term Structure of Inflation Expectations*. Tech. rep., FEDS Working Paper.

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.

Angeletos, G.-M. and La'O, J. (2009). Incomplete information, higher-order beliefs and price inertia. *Journal of Monetary Economics*, **56**, S19–S37.

— and Lian, C. (2018). Forward guidance without common knowledge. *American Economic Review*, **108** (9), 2477–2512.

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 Klaauw, W. (2022). *Handbook of Economic Expectations*. Elsevier.

Bahaj, S., Czech, R., Ding, S. and Reis, R. (2023). *The Market for Inflation Risk*. Tech. rep., Bank of England.

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

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

— and Turen, J. (2025). *Lumpy Forecasts in Volatile Times*. 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.

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.

Broer, T. and Kohlhas, A. N. (2024). Forecaster (mis-) behavior. *Review of Economics and Statistics*, **106** (5), 1334–1351.

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

Campbell, J. R. and Eden, B. (2014). Rigid prices: Evidence from us scanner data. *International Economic Review*, **55** (2), 423–442.

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

— and Timmermann, A. (2009). Disagreement and biases in inflation expectations. *Journal of Money, Credit and Banking*, **41** (2-3), 365–396.

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.

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

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

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

Gaglianone, W. P., Giacomini, R., Issler, J. V. and Skreta, V. (2022). Incentive-driven inattention. *Journal of Econometrics*, **231** (1), 188–212.

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

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, S. (2023). Quantifying noise in survey expectations. *Quantitative Economics*, **14** (2), 609–650.

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.

Matějka, F. (2015). Rigid pricing and rationally inattentive consumer. *Journal of Economic Theory*, **158**, 656–678.

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.

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 Timmermann, A. (2010). Why do forecasters disagree? lessons from the term structure of cross-sectional dispersion. *Journal of Monetary Economics*, **57** (7), 803–820.

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

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.

Weber, M., d'Acunto, F., Gorodnichenko, Y. and Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, **36** (3), 157–184.