Reputational Information Cycles*

Renan Muta† Guillermo Ordoñez‡

November 17, 2026

Abstract

Reputational incentives can narrow disagreement among professional forecasters and make their consensus less accurate. In our model, skilled forecasters extract more information from a signal that contains a shared error. When evaluated relative to their peers, forecasters overweight that signal to demonstrate skill, amplifying the shared error in the consensus. In Bloomberg's U.S. inflation survey, forecast errors carry larger ranking penalties and forecasts cluster more tightly when the inflation risk premium is high. We interpret high-premium periods as times when reputational stakes are greater. As those stakes rise and fall, the distortion in what forecasters report changes, generating cycles in how much information the consensus conveys.

JEL Codes: E44, D83, D84, D82, E37, G17

Keywords: reputation, information aggregation, professional forecasters, information cycles.

*First draft: November 17, 2026. We thank Vincent Glode, Juan Cruz Llambias, and Alex Oktay for helpful comments and suggestions.

†The Wharton School, University of Pennsylvania. Email: renanm@wharton.upenn.edu

‡University of Pennsylvania and NBER. Email: ordonez@econ.upenn.edu

1 Introduction

Macroeconomic consensus forecasts are important tools to guide investment and policy decisions by bringing together information dispersed across professional forecasters. Ahead of a U.S. inflation announcement, for example, Bloomberg publishes the median of individual forecasts, providing a public benchmark for making decisions and later interpreting the realized inflation release. But these forecasts are not provided by machines; the economists who supply these forecasts depend on their reputations to attract and retain clients. Each submission contributes to the consensus and adds to the track record on which clients judge the forecaster's expertise. The forecasts that enter the consensus therefore reflect the information available to forecasters, their ability to interpret it, and their incentives to demonstrate that ability. We ask how the strength of those reputational incentives affects the quality of the information conveyed by the consensus, and how those incentives change over the business cycle, affecting the quality of information available for decision making.

Suppose some analysts have a greater ability to extract information from a signal whose errors have a component in common across the profession. A forecaster seeking to demonstrate ability may respond more strongly to that signal than is warranted by its predictive value. If many forecasters face the same incentive, their reports become more alike and the consensus becomes more exposed to the common error. Reputation thus explains why public measures of expectations can become less informative even when forecasters appear to agree more. We use the term *reputational information cycles* for this variation in the informational quality of an aggregate forecast as reputational incentives change with economic conditions.

The Bloomberg survey of monthly U.S. CPI inflation provides a setting in which to study this mechanism. The survey records individual submissions and their timing, provides a public consensus, and publishes rankings that make forecast performance visible. Three features of the evidence motivate the analysis. First, changes in the rank of an individual forecaster depend not only on the accuracy of such a forecaster but also on the accuracy of the consensus. Second, the association between an individual forecast error and a deterioration in rank is stronger when the inflation risk premium is high, which we take as a situation in which clients value and rely more

on the forecaster's reputation. Third, forecast dispersion is lower in these periods, when agreement with earlier submissions also has a weaker association with individual accuracy. Taken together, these patterns raise a question about how the prospect of being evaluated relative to one's peers changes the information that makes it into the consensus.

In the model, analysts draw on two private sources of information to predict an outcome such as inflation. Both sources are informative, and they differ in how their errors are distributed across forecasters. Errors from one source are specific to the analyst. Errors from the other include a component shared by everyone who uses it. More skilled analysts observe the second source with less individual noise and therefore give it more weight when forming their best prediction. How an analyst uses that information can reveal something about their ability.

Forecasters care about the accuracy of the numbers they submit and what it will lead others to infer about their skill. In the model, the evaluator knows how forecasters use their information. After the release, the average forecast and the realized outcome together allow the evaluator to infer the common error. This helps the evaluator judge whether an individual report is characteristic of a more or less skilled analyst. The same forecast miss can have different consequences for reputation depending on what other forecasters reported. Forecasters anticipate this evaluation when choosing how much of each signal to reflect in their reports.

We first study a model in which all forecasters make reports simultaneously. In the model, the desire to appear skilled changes how forecasters use the information they already have. As reputational stakes rise, both types of analysts place greater weight on the signal that skilled analysts naturally rely on more heavily. They give less weight to the signal with independent errors. Their forecasts draw closer together because they respond more strongly to the shared disturbance and rely less on the information that differs across forecasters. This form of agreement can arise even when every analyst submits before seeing anyone else's forecast.

Forecast errors that are particular to individual analysts tend to cancel when reports are averaged. A shared error survives because it pushes many forecasts in the same direction. When reputational concerns strengthen, the consensus therefore carries more shared noise relative to the information

it contains about the economy. The shared error remains embedded in the consensus, since a shift in the average could reflect either a change in the outcome that forecasters are trying to predict or a change in the error they have in common. Lower dispersion can therefore give forecast users a misleading impression of how much they can learn from the consensus.

We then allow for sequential forecasting, which is closer to how forecasts are submitted in our empirical case: a small group of forecasters submit first and the rest follow after observing those first forecasts. A follower learns about the earlier forecasts, but if they also care about reputation, can follow the early forecasts strategically. We examine how this incentive differs when early forecasters are known to be highly skilled or less skilled. We show that followers rely less on a low-ability leader as reputational concerns rise, but rely more on a high-ability leader. A small group can consequently add little additional information but have a substantial indirect influence through subsequent submissions. At high reputational stakes, access to a skilled leader's forecast improves the accuracy of the consensus relative to simultaneous reporting.

Our empirical analysis uses 19,583 individual forecasts from the Bloomberg survey for 260 releases of monthly U.S. CPI inflation from January 2005 through October 2026. The average release has approximately 75 forecasts. To study sequential reporting, we calculate a median for each forecaster using forecasts submitted on earlier dates. We use the previous month's estimate of the ten-year inflation risk premium as a proxy for variation in reputational stakes: a higher inflation risk makes forecasters' clients more concerned about their ability.

We first examine changes in Bloomberg rankings, which is a coarse proxy for reputation in this environment (highly ranked forecasters are deemed as accurate ones). We show that both the consensus's recent absolute errors and those of the individual forecaster raise the share of variation explained from 18.5 to 33.5 percent. An individual forecast miss is also associated with a larger ranking penalty when the premium is high (our proxy for reputation concerns).

We show indeed that a ten-basis-point increase in the lagged inflation risk premium is associated with a 1.3-basis-point decline in the cross-sectional standard deviation of forecasts, relative to a sample mean of 8.8 basis points. We also show that reports closer to the earlier median tend to

be more accurate, and this relationship weakens as the premium rises. The combination of tighter clustering and a weaker accuracy advantage is consistent with the mechanism in the model. The status of earlier forecasters helps explain where later forecasts are placed. When at least one member of the earlier pool was previously in Bloomberg's Top 10, a later submission is approximately 7.8 percentage points less likely to differ from that pool's median.

The benefit of our dataset is that the reports, their ordering, the evaluation of their authors, and the public use of their aggregate are observable. The mechanism can operate in other settings. Forecasts of employment, output, and policy rates are natural applications because they can serve as benchmarks for interpreting subsequent public announcements. In those settings, the incentives governing the production of a consensus can affect both what the public learns from expert forecasts and what it subsequently classifies as new information. Understanding the state dependence of those incentives is therefore part of understanding the state dependence of the informational value of macroeconomic releases.

Related Literature. Our paper contributes to two main strands of literature. First, we contribute to the literature on reputational incentives and herding among professional forecasters. Scharfstein and Stein (1990) show that career concerns can induce managers to disregard private information and follow the decision of others. Trueman (1994) shows that reputation concerns induce financial analysts to issue forecasts closer to prior expectations and to previously announced forecasts, with lower-ability analysts exhibiting stronger herding behavior. Graham (1999) further studies how reputation, forecasting ability, and correlated private signals affect incentives to herd, while Ottaviani and Sørensen (2006*a*) focus on reputational cheap talk, showing that concerns about appearing well-informed generate herding and incomplete learning in sequential settings. Others argue that professional forecasters have incentives to differentiate their predictions to increase the publicity associated with making an accurate forecast (Laster, Bennett and Geoum, 1999; Ottaviani and Sørensen, 2006*b*; Gemmi and Valchev, 2026). We build on this literature by studying how reputation concerns and the reputation of earlier forecasts shape the incentives to conform or

differentiate, especially when reputation is updated not only based on one's own accuracy but also on the accuracy of others.

Our paper is also related to recent work examining how strategic incentives distort professional forecasts and the implications for the aggregate information in the consensus. Lichtendahl, Grushka-Cockayne and Pfeifer (2013) and Gemmi and Valchev (2026) show that competition and strategic differentiation can induce forecasters to overweight idiosyncratic information, potentially improving the accuracy of the consensus despite distorting individual forecasts. In contrast, our mechanism induces greater reliance on commonly biased information, reducing the informativeness of the aggregate consensus.

Our empirical analysis also relates to the recent literature focused on understanding the behavior of professional forecasters in Bloomberg surveys. Baley and Turen (2024) study how forecast revision costs and strategic concerns generate lumpy adjustment in professional inflation forecasts. Qiu, Deschamps and Liu (2024) show that macroeconomic uncertainty is associated with greater forecast disagreement and lower accuracy, while financial uncertainty is associated with lower disagreement and higher accuracy[1]. They interpret the latter finding as consistent with forecasters remaining closer to the consensus when financial conditions are uncertain. We complement their analysis by providing a reputation-based mechanism through which economic conditions can affect forecasting strategies. Moreover, we explore the sequential nature of Bloomberg forecasts and its ranking system to further investigate whether the influence of previously submitted forecasts depends on reputation and the salience of inflation risk.

Second, we contribute to the literature on information cycles, which studies how information production, learning, and the informativeness of economic signals vary with aggregate conditions. Veldkamp (2005), Van Nieuwerburgh and Veldkamp (2006) and Ordoñez (2013) show how economic activity affects the production and precision of information, generating asymmetric learning over the business cycle. Gorton and Ordoñez (2014) and Chousakos, Gorton and Ordoñez (2023) study how endogenous information production interacts with credit and economic activity.

[1]The authors use the Jurado, Ludvigson and Ng (2015) metrics for macroeconomic and financial uncertainty.

More recently, Bostanci and Ordoñez (2026) show how liquidity conditions affect the information about fundamentals contained in stock prices, with consequences for capital allocation and output. Closely related to our mechanism, Ordoñez (2013) shows how reputation concerns depend on the behavior of other agents, generating strategic complementarities and clustering of risk-taking that can amplify small changes in economic fundamentals. We build on this insight by studying how reputation evaluations based on aggregate outcomes affect the use and aggregation of information. Rather than focusing on changes in information production or acquisition, we show that state-dependent reputational incentives distort how forecasters report their available information, changing the informativeness of the consensus even when the underlying precision of private signals remains unchanged. Sequentiality introduces an additional mechanism, as idiosyncratic errors from earlier forecasters can propagate through subsequent forecasts and become aggregate errors. Our framework therefore highlights a reputational channel through which the quality of aggregate information can fluctuate across economic states.

This paper is organized as follows. Section 2 introduces the baseline model and characterize how reputation concerns affect optimal forecasting decisions and information aggregation. Section 3 extends the model by introducing sequential forecasting, examining how the reputation of earlier forecasters influences subsequent forecasts and the information quality of the consensus. Section 4 describes the data, while Section 5 presents the empirical results. Section 6 concludes.

2 Baseline Model

The economy is static. There is a unit mass of forecasters indexed by $i \in [0, 1]$, whose job is to forecast the fundamental θ. They observe two private signals, a biased signal, y_i, and an unbiased signal, x_i, before submitting a forecast a_i. Forecasters can be either a high or low type, with high types being better at processing the biased signal. Each forecast has an initial reputation φ_i, which reflects the perceived probability that forecaster i is a high-type. All forecasts are submitted simultaneously and are aggregated into a consensus $\bar{a} = \int_0^1 a_i \, di$. After the realization of the

fundamental θ, an external evaluator updates the reputation of each forecaster through Bayes' rule using the actual realization θ, the consensus $\bar{a}$, and their individual forecast a_i.

Fundamental. The fundamental θ follows:

$$\theta \sim N(\mu_0, v(U)), \quad v'(U) > 0 \tag{1}$$

where U is a state variable representing the uncertainty about the underlying fundamental. We assume that the variance of the process increases with the uncertainty level.

Private Signals. All forecasters observe two private signals: a biased signal y_i and a private signal x_i.

$$y_i = \theta + \xi + \eta_i \tag{2}$$

$$x_i = \theta + u_i \tag{3}$$

where $\xi \sim N(0, \sigma_\xi^2)$ is a common bias and $u_i \sim N(0, \sigma_u^2)$.

Forecasters' Types. Forecasters have two types, either high or low. Let q_i be the type for forecaster i and $q_i \in \{H, L\}$. High-type forecasters are better at processing the biased signal.

$$\eta_k \mid H \sim N(0, \sigma_{\eta H}^2), \quad \eta_k \mid L \sim N(0, \sigma_{\eta L}^2), \quad \sigma_{\eta H}^2 < \sigma_{\eta L}^2$$

One can interpret the unbiased signal x_i as the result of forecasters' private statistical models and the biased signal y_i as coming from biased news. High type forecasters are the ones that can better interpret the common bias from the news outlet.

For convenience, we define precisions to be:

$$\tau_0(U) = \frac{1}{v(U)}, \quad \tau_{yq_k} = \frac{1}{\sigma_\xi^2 + \sigma_{\eta q_k}^2}, \quad \tau_u = \frac{1}{\sigma_u^2}$$

Fundamental Bayesian Posterior. Forecasters form posterior beliefs about the fundamental through Bayesian updating. A type-q forecaster's posterior belief about the fundamental is:

$$m_q^\theta = (1 - \alpha_q^B - \beta_q^B)\mu_0 + \alpha_q^B y_k + \beta_q^B x_k \tag{4}$$

where

$$\alpha_q^B \equiv \frac{\tau_{yq}}{\tau_0(U) + \tau_{yq} + \tau_u}, \quad \beta_p^B \equiv \frac{\tau_u}{\tau_0(U) + \tau_{yq} + \tau_u}$$

The uncertainty about the fundamental U dictates the extent to which the agents employ their private signals. However, it does not determine the composition between α_q^B and β_q^B.

Since $\tau_{yH} > \tau_{yL}$, high types place more weight on the biased signal relative to low types when forming the Bayesian beliefs:

$$\alpha_H^B > \alpha_L^B, \quad \beta_H^B < \beta_L^B$$

Common Bias Bayesian Posterior. Each forecaster also forms a belief about the realization of the common bias ξ after observing their own private signal y_i. Let $m_q^\xi \equiv \mathbb{E}[\xi \mid y_i, x_i, q_i = q]$ be the common bias posterior belief after observing the private signals.

Conditional on the realization of θ, forecasters would be able to infer the common bias with some noise:

$$\xi + \eta_i = y_i - \theta$$

By Gaussian signal-extraction:

$$\mathbb{E}[\xi \mid y_i, \theta, q_i = q] = \frac{\sigma_\xi^2}{\sigma_\xi^2 + \sigma_{\eta q}^2}(y_i - \theta)$$

Because forecasters do not actually observe θ, they use their posterior belief about the fundamental:

$$m_p^\xi = \frac{\sigma_\xi^2}{\sigma_\xi^2 + \sigma_{\eta p}^2}(y_j - m_p^\theta) \tag{5}$$

Reputation and external evaluator. There is an external evaluator who updates the reputation of each forecaster through Bayesian updating. Let the reputation of forecaster k be φ_k, which is the evaluator's prior probability that the forecaster is high-type:

$$\varphi_k \equiv \Pr(q_k = H) \tag{6}$$

For tractability, we define reputation log-odds ℓ_k as:

$$\ell_k \equiv \log \frac{\varphi_k}{1 - \varphi_k} \tag{7}$$

Preferences. Each forecaster provides a forecast a_k in order to optimize the following objective function:

$$\max_{a_k} \mathbb{E}\left[\rho \ell'_k - \frac{\gamma}{2}(a_k - \theta)^2 \mid \mathcal{I}_k\right] \tag{8}$$

where ρ is a state variable that captures reputation concerns, ℓ' is next period's log-odds reputation, γ is the accuracy cost, and $\mathcal{I}_k$ is the information set for forecaster k.

Forecasters care about their future reputation. Since reputation is public, one can imagine that high reputable forecasters secure more clients and receive higher remuneration. The degree to which they care about future reputation is determined by the state variable ρ. Forecasters also generate disutility from submitting a forecast that is too far from their posterior. One can interpret the second term in the utility as an auditing cost, with regulators scrutinizing the rationale behind a forecast. By making accuracy costly, we rule out cheap-talk games as in Ottaviani and Sørensen (2006*a*).

Evaluator Using the Average Forecast. The external evaluator conjectures that a type-s forecaster follows an affine strategy on the observed signals.

$$a_{i,s} = (1 - \alpha_s - \beta_s)\mu_0 + \alpha_s y_i + \beta_s x_i \tag{9}$$

where α_s and β_s are the evaluator's conjectured signal weights for $s \in \{H, L\}$.

Let χ denote the actual population share of high-type forecasters, which is publicly available. Hence, the weighted-average forecaster weight in each signal can be written as:

$$\bar{\alpha} = \chi\alpha_H + (1-\chi)\alpha_L, \quad \bar{\beta} = \chi\beta_H + (1-\chi)\beta_L$$

In the consensus $\bar{a}$, idiosyncratic errors from Equation 9 average out:

$$\bar{a} - \mu_0 = (\bar{\alpha} + \bar{\beta})(\theta - \mu_0) + \bar{\alpha}\xi \tag{10}$$

Hence, provided that $\bar{\alpha} \neq 0$, the external evaluator can use the consensus to invert 10 to recover the realization of the common bias ξ after observing the realization of θ:

$$\xi = \frac{\bar{a} - \mu_0 - (\bar{\alpha} + \bar{\beta})(\theta - \mu_0)}{\bar{\alpha}} \tag{11}$$

Therefore, the consensus $\bar{a}$ is used by the external evaluator to update reputation. This is an important feature of the model because forecasters' reputation is updated not only based on their own forecast a_k and the realized fundamental θ, but also on what other forecasters said.

Evaluator Updating Reputation. Plugging in the definitions of the private signals allows us to write the evaluator's conjecture for a type-s forecaster as:

$$a_{i,s} = (1 - \alpha_s - \beta_s)\mu_0 + (\alpha_s + \beta_s)\theta + \alpha_s\xi + \alpha_s\eta_i + \beta_s u_i \tag{12}$$

Conditional on the realized (θ, ξ), the evaluator expects a type-s forecast to report the following forecast:

$$\mu_s = (1 - \alpha_s - \beta_s)\mu_0 + (\alpha_s + \beta_s)\theta + \alpha_s\xi \tag{13}$$

with variance:

$$V_s = (\alpha_s)^2\sigma_{\eta s}^2 + (\beta_s)^2\sigma_u^2 \tag{14}$$

The evaluator then uses 13 and 14 as benchmarks for the high and low types to update the leaders' reputation log-odds using Bayes rule:

$$\ell_i' = \ell_i + \frac{1}{2}\log\left(\frac{V_L}{V_H}\right) - \frac{(a_i - \mu_H)^2}{2V_H} + \frac{(a_i - \mu_L)^2}{2V_L} \tag{15}$$

The higher the prior reputation ℓ_i, the higher is the future reputation ℓ_i'. A forecast a_i that is farther from the high-type benchmark μ_H reduces the updated reputation, while an a_i farther from the low-type benchmark μ_L improves future reputation. The noisier the benchmark, the less weight it has when updating reputation.

Forecaster's Problem. Forecaster i chooses a forecast a_i to maximize the following maximization problem:

$$\max_{a_i} \mathbb{E}\left[\rho\ell_i' - \frac{\gamma}{2}(a_i - \theta)^2 \mid x_i, y_i, q_i = q\right] \tag{16}$$

Forecasters know that the evaluator updates future reputation using 15. However, since they do not know the realized pair (θ, ξ), they must form expectations over 13 to solve the maximization problem. Let $\mu_{s,q}^* \equiv \mathbb{E}[\mu_s \mid y_i, x_i, q_i = q]$ be the type-q forecaster's expectation of how the evaluator benchmarks a type-s forecaster. Hence, using the posteriors m_q^θ and m_q^ξ to take expectations over 13:

$$\mu_{s,q}^* = (1 - \alpha_s - \beta_s)\mu_0 + (\alpha_s + \beta_s)m_q^\theta + \alpha_s m_q^\xi \tag{17}$$

Forecaster's Optimal Forecast. The first-order condition of the problem in Equation 16 for a type-q forecaster is:

$$\gamma(a_{i,q} - m_q^\theta) = \rho\left[-\frac{a_{i,q} - \mu_{H,q}^*}{V_H} + \frac{a_{i,q} - \mu_{L,q}^*}{V_L}\right] \tag{18}$$

Defining $D \equiv \gamma + \rho\left(\frac{1}{V_H} - \frac{1}{V_L}\right)$ and re-arranging 18, we can write the optimal forecast for a

type-q forecaster as $a^*_{i,q}$:

$$a^*_{i,q} = m^\theta_q + \frac{\rho}{D}\left[\frac{\mu^*_{H,q} - m^\theta_q}{V_H} - \frac{\mu^*_{L,q} - m^\theta_q}{V_L}\right] \tag{19}$$

When there are no reputation concerns ($\rho = 0$), forecasters optimally report their Bayesian posterior belief. Similarly, when types have identical signal-processing abilities ($\sigma^2_{\eta H} = \sigma^2_{\eta L}$), both types follow the same reporting strategies in equilibrium, making their forecast distributions identical and not allowing the evaluator to infer type from her forecast. Reputation concerns then do not distort reporting (the bracketed term becomes zero) and forecasters optimally submit their Bayesian posterior beliefs. However, when reputation concerns are present ($\rho > 0$) and types differ in their signal-processing abilities ($\sigma^2_{\eta H} \neq \sigma^2_{\eta L}$), reputational incentives generally induce forecasters to deviate from their Bayesian posterior beliefs.

Equilibrium. We characterize an affine rational expectations equilibrium in which forecasters optimally choose their reporting strategies taking the equilibrium strategies of other forecasters, the evaluator's reputation-updating, and the public states as given.

Definition (Affine Rational Expectations Equilibrium). Given the public states (U, ρ), an affine rational expectations equilibrium consists of type-contingent strategies, their associated posterior beliefs, and an aggregate forecast satisfying the following conditions:

1. Each type-q forecaster submits a forecast that solves the optimization problem in Equation 16, conditional on her private signals and type, taking the evaluator's reputation-updating rule as given.

2. Forecasters form Bayesian posterior beliefs about the fundamental and the common bias conditional on their private signals and type.

3. The external evaluator's conjectured type-contingent reporting strategies coincide with forecasters' optimal reporting strategies. The evaluator updates reputations according to 15.

4. The aggregate consensus coincides with the population average of the equilibrium reporting strategies, as characterized by Equation 10.

An equilibrium is therefore characterized by a fixed point in the type-contingent reporting coefficients, such that the evaluator's conjectured strategies and forecasters' optimal reporting strategies coincide[2].

2.1 Reputational Distortions and Information Aggregation

In this section we discuss some theoretical results from the model. We restrict our attention to the equilibrium that is connected to the Bayesian equilibrium when $\rho = 0$, which is unique for every value of $\rho > 0$[3]. We first explain how changes in reputation concerns states (ρ) distort the optimal forecast away from the Bayesian posterior in equilibrium. We later discuss the consequences for the quality of information contained in the overall consensus and for forecasters' cross-sectional dispersion.

Forecasters Overweight the Biased Signal Due to Reputation Concerns. When leaders care about future reputation, they want to provide a forecast that would be perceived by the external evaluator as a high-type forecaster. In the Bayesian benchmark, high-type forecasters set a larger relative weight in the biased signal than low-type forecasters, since they are better at processing the bias. Low-type forecasters then overweight the biased signal to look like high-type forecasters. Similarly, high-type forecasters also overweight the biased signal to differentiate themselves from the low-type forecasters' strategy[4].

Proposition 1: Reputational Overweighting of the Biased Signal. *Suppose the common-bias variance is positive and high-type forecasters have strictly greater signal-processing ability than low-type forecasters. In the affine equilibrium locally connected to the Bayesian benchmark, an*

[2]See Appendix A.1 for an illustration of the fixed point problem.

[3]See Appendix B for details on computing the equilibrium.

[4]We do not have a pooling equilibrium because forecasters have a preference for being accurate relative to their posterior. This limits the extent to which their optimal forecast can deviate from what a type believes.

increase in reputation concerns from $\rho = 0$ raises the relative weight α_q/β_q placed on the biased signal for both forecaster types.

Proof. See proof in Appendix A.2.

Proposition 1 establishes that reputational incentives induce both forecaster types to increase their relative reliance on the biased signal. When reputation concerns arise, both types have an incentive to further emphasize their employment of the biased signal because doing so makes their forecasts more characteristic of high-ability forecasters. This distortion emerges even for arbitrarily small reputation concerns.

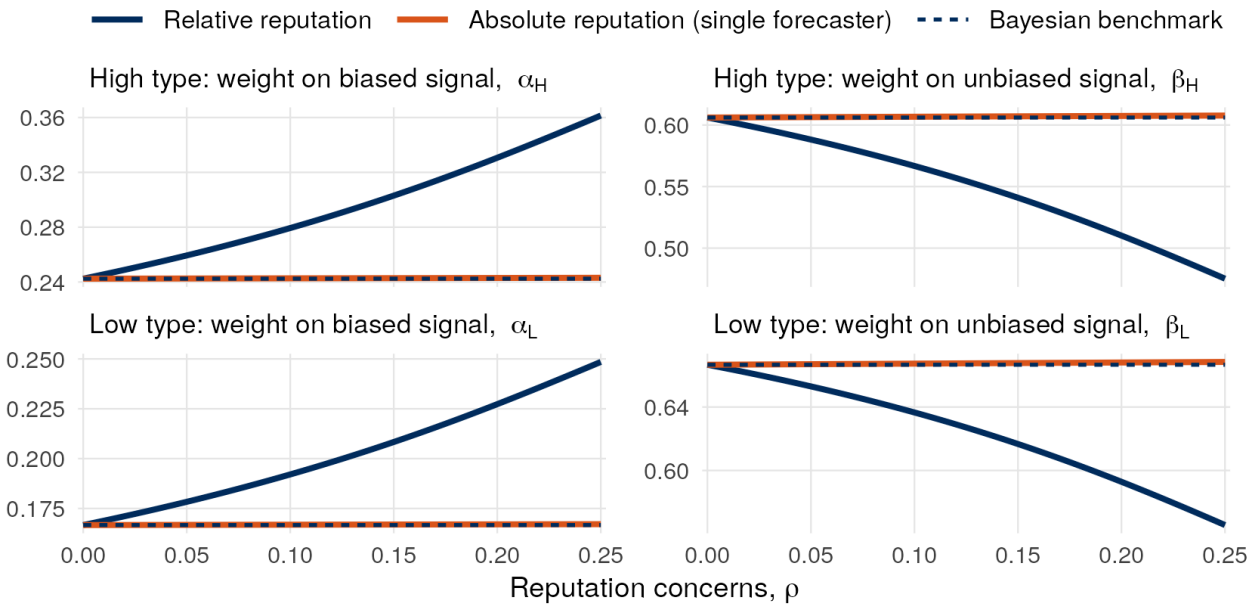


Figure 1: Reputation Concerns and Equilibrium Signal Weights

Notes: The figure illustrates how reputation concerns affect equilibrium reporting strategies under relative and absolute reputation evaluation. The top (bottom) panels correspond to high-type (low-type) forecasters, while the left (right) panels show the weights assigned to the biased (unbiased) private signal. The horizontal axis represents reputation concerns ρ. Solid blue curves depict equilibrium weights when the evaluator uses the aggregate consensus to update individual reputations, while solid orange curves represent the absolute-evaluation benchmark, in which reputation is updated using only the forecaster's own report and the realized fundamental. Dashed blue lines represent the Bayesian benchmark. Under relative evaluation, increasing reputation concerns induce both types to overweight the biased signal and underweight the unbiased signal. Under absolute evaluation, the relative allocation of weights across signals remains unchanged, and the equilibrium weights remain close to their Bayesian benchmarks.

Figure 1 illustrates Proposition 1. As reputation concerns increase, both types overweight the biased signal and underweight the unbiased signal relative to the Bayesian benchmark. Importantly, these distortions arise because the evaluator updates individual reputations using information from

other forecasters' reports. By observing the aggregate consensus and the realized fundamental, the evaluator can recover the common bias ξ, allowing her to distinguish common forecasting mistakes from individual signal-processing errors. Hence, the relative weight that a forecaster places on the biased signal becomes informative about her forecasting ability. Both types then have an incentive to increase their reliance on this signal to appear more likely to be high-type forecasters.

To highlight the importance of reputation being updated relative to others, Figure 1 also considers an absolute-evaluation benchmark (orange line) in which the evaluator observes only the forecaster's own report and the realized fundamental. Without the aggregate consensus, the evaluator cannot separately identify the common bias ξ. In this environment, reputation concerns do not distort the relative allocation of weights between the two private signals, since the evaluator cannot infer whether one was used more than the other. What changes is the usage of public information (the prior) relative to the private signals. Nonetheless, the optimal strategy is broadly close to the Bayesian benchmark. Thus, substantial overweighting of the biased signal in our baseline model is driven by the evaluator's ability to assess individual forecasting performance relative to information contained in other forecasts rather than by reputation concerns alone.

Information Aggregation and Cross-sectional Dispersion. Having established how reputation concerns distort the reported forecast away from the Bayesian benchmark, we now examine their implication for the quality of information contained in the overall consensus. The consensus is an important summary statistic for both financial market participants and policymakers, especially because surprises from macroeconomic news announcements are derived taking the forecasters' consensus as the benchmark (Elenev et al., 2024).

We measure the information aggregation through the mean-squared error of the overall consensus, with a larger MSE indicating a less accurate consensus.

$$\mathcal{M}(U,\rho) \equiv \mathbb{E}\left[(\bar{a}-\theta)^2 \mid U,\rho\right]$$

From Equation 10, the consensus depends on the fundamental and the common bias, since

idiosyncratic forecasting errors average out. The mean-squared errors can be written as:

$$\mathcal{M}(U,\rho) = (1-\bar{\alpha}-\bar{\beta})^2 v(U) + \bar{\alpha}^2\sigma_\xi^2 \tag{20}$$

The first term captures the extent to which the consensus does not fully respond to changes in the fundamental, while the second term reflects its exposure to the common bias. The biased signal is particularly costly for aggregate accuracy since its common bias component cannot be diversified away through aggregation.

We also measure forecast dispersion through the ex-ante expected cross-sectional variance of individual forecasts:

$$\mathcal{D}(U,\rho) \equiv \mathbb{E}\left[\int_0^1 (a_i-\bar{a})^2\, di \;\middle|\; U,\rho\right] \tag{21}$$

Both consensus MSE and forecast dispersion are ex-ante expected moments, integrating over realizations of the fundamental and signal disturbances conditional on the public states (U,ρ). The model therefore characterizes how changes in reputation concerns affect expected forecast dispersion and accuracy, which we can characterize analytically.

Proposition 2: Forecast Dispersion and Consensus Accuracy. *Fix the uncertainty state U. Let $\mathcal{M}(U,\rho)$ denote the mean-squared error of the overall consensus and $\mathcal{D}(U,\rho)$ be the expected cross-sectional variance of individual forecasts. Then,*

$$\mathcal{M}(U,\rho) - \mathcal{M}(U,0) = -\Big[\mathcal{D}(U,\rho) - \mathcal{D}(U,0)\Big] + \mathcal{R}(U,\rho)$$

where

$$\mathcal{R}(U,\rho) \equiv \mathbb{E}\left[\int_0^1 \left(a_i - m_{q_i}^\theta\right)^2\, di \;\middle|\; U,\rho\right]$$

measures the squared deviation of reported forecasts from their Bayesian posterior beliefs. When reputational incentives reduce forecast dispersion relative to the Bayesian benchmark, the mean-squared error of the consensus must increase.

Proof. See proof in Appendix A.3.

Proposition 2 establishes a direct relationship between individual forecast dispersion and the quality of aggregate information. Since Bayesian posterior beliefs minimize individual expected squared errors, reputation distortions cannot improve individual forecasting accuracy. Whenever reputation concerns reduce forecast dispersion relative to the Bayesian benchmark, the mean-squared error of the consensus must increase.

In our model, this deterioration can arise because reputational incentives induce forecasters to shift weight away from their unbiased private signal and toward the commonly biased signal. Although the unbiased signal contains idiosyncratic noise, these errors disappear through aggregation, making this information particularly valuable for the accuracy of the overall consensus. In contrast, greater reliance on the biased signal increases the exposure of the consensus to the common bias. Thus, the same reputational incentives that encourage forecasters to submit more similar predictions can also make their aggregate forecast less informative[5].

Figure 2 illustrates these results. As reputation concerns increase, forecast dispersion decreases while the mean-squared error of the consensus rises. The reduction in dispersion is driven primarily by forecasters' declining reliance on the unbiased private signal, which reduces the contribution of idiosyncratic noise to the dispersion of forecasts. Meanwhile, consensus becomes increasingly exposed to the common bias, causing a deterioration in aggregate accuracy.

These results highlight that greater agreement among forecasters does not necessarily reflect better information aggregation. In our model, reputation concerns can induce forecasters to place greater weight on commonly biased information, making their predictions more similar while reducing the accuracy of the overall consensus. Although Proposition 2 establishes a general relationship between forecast dispersion and consensus accuracy, the direction in which reputation concerns affect dispersion depends on equilibrium reporting strategies. Given our Propositions 1

[5]This result contrasts with Gemmi and Valchev (2026), where strategic incentives to differentiate from other forecasters induce greater reliance on idiosyncratic information and less reliance on common information. Although this behavior distorts individual forecasts, it improves the accuracy of the consensus relative to the average of the Bayesian beliefs, as idiosyncratic errors disappear through aggregation. Consistent with our Proposition 2, this improvement occurs when the increase in forecast dispersion more than compensates for the additional individual forecasting errors generated by strategic reporting. See also Lichtendahl, Grushka-Cockayne and Pfeifer (2013).

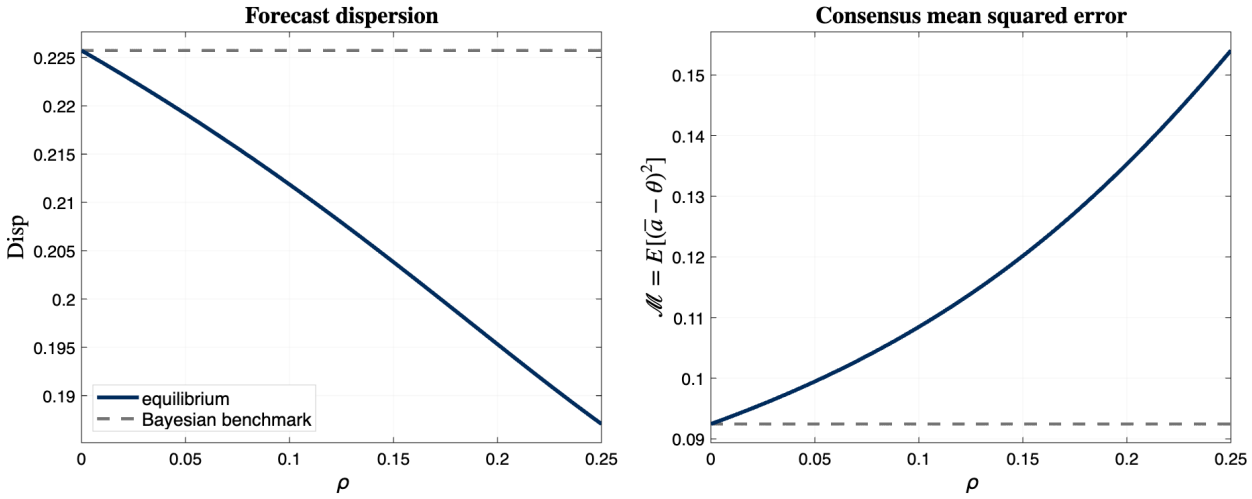


Figure 2: Reputation Concerns, Forecast Dispersion, and Consensus Accuracy

The figure illustrates how reputation concerns affect forecast dispersion and the accuracy of the aggregate consensus. The left panel reports the ex-ante expected cross-sectional variance of individual forecasts, while the right panel reports the mean-squared error (MSE) of the consensus. The horizontal axis represents reputation concerns (ρ). Solid blue curves depict equilibrium outcomes under relative reputation evaluation, while dashed lines represent the Bayesian benchmark ($\rho = 0$). As reputation concerns increase, forecast dispersion declines while consensus MSE rises, indicating that greater agreement among forecasters is accompanied by a deterioration in consensus information quality.

and 2, this motivates our first testable implication:

Testable Implication #1: Reputation Concerns and Lower Dispersion. *As reputation concerns increase, the cross-sectional dispersion of forecasts decreases.*

The prediction concerns expected forecast dispersion conditional on the fundamental uncertainty state, rather than dispersion in each individual realization. In the empirical analysis, we investigate whether periods associated with stronger reputation concerns exhibit lower dispersion across professional forecasts. This allows us to assess whether the observed patterns of forecast disagreement are consistent with the reputational distortions in information use characterized by the model.

3 Model with Sequentiality

In the baseline model, forecasters submit their predictions simultaneously, without observing the forecasts of others. However, professional forecasts submitted to Bloomberg are sequential,

allowing forecasters to observe earlier predictions before making their own. This introduces an additional channel through which reputational incentives can distort the use of information. Earlier forecasts provide information about the fundamental (Bikhchandani, Hirshleifer and Welch, 1992), but they may also influence how subsequent forecasters are evaluated. Forecasters then must decide not only how much weight to place on their own private signals, but also whether to conform or to differentiate themselves from previously submitted forecasts, which can vary with the reputation of the previous forecaster. To capture this mechanism, we extend the baseline model by introducing sequentiality, as we describe below.

Leader and Followers. A random forecaster from the continuum is selected to move first, which we denote as the leader and is indexed by j. The leader submits a forecast a_j and has reputation φ_j. Subsequent forecasters, denoted the followers and indexed by i, observe the leader's forecast and reputation before making their own predictions. Forecasts are aggregated into the consensus $\bar{a}$. Because the leader is mass zero, she does not mechanically influence the consensus, although she might indirectly affect the consensus through followers' optimal forecasts.

Timing. The model's timing protocol follows:

1. States U and ρ are publicly observed. The leader's reputation φ_j is also publicly observed.
2. Leader observes her type $p \in \{H, L\}$. Others do not know the leader's actual type.
3. The leader observes her private signals (x_j, y_j) and submits a forecast a_j.
4. Followers observe their private type $q_i = q$ and their own private signals (x_i, y_i). They use a_j as an additional signal about the fundamental θ and the common bias ξ. They take φ_j as the probability that the leader is a high-type.
5. Followers submit their own forecasts a_i, generating the consensus $\bar{a}$.
6. The fundamental θ is realized. The external evaluator uses θ, $\bar{a}$, and each individual forecast a_k where $k \in \{i, j\}$ to update the reputation of each forecaster to φ'_k.

Leader. Let the leader's type be $p \in \{H, L\}$. Although the leader's forecast affects subsequent reporting decisions and the overall consensus, this influence does not introduce an additional strategic incentive in the leader's problem. The evaluator observes both the leader's forecast and the aggregate consensus and can therefore recover the common bias after accounting for followers' response. Hence, changes in the leader's report do not affect the evaluator's inference about the common bias, even when the leader anticipates their influence on the followers. The leader's optimal reporting problem is therefore identical to that of a forecaster in the baseline economy, allowing us to solve the model recursively.

Lemma 1: Recursive Leader Optimization *Suppose followers use affine reporting strategies and the evaluator observes the realized fundamental, the leader's forecast, and the overall consensus. If the consensus identifies the common bias ($\bar{\alpha} \neq 0$), the leader's optimal reporting problem is independent of the followers' equilibrium strategies. A leader who anticipates the effects of her forecast on subsequent reporting decisions faces the same optimization problem as a forecaster in the simultaneous economy.*

Proof. See Appendix A.4.

Followers. Each follower i observes her own private signals (x_i, y_i), her own type q_i, the leader's forecast a_j, and the reputation of the leader φ_j before submitting a prediction a_i. Since there is an additional signal about the fundamental, we discuss next the new problem for the followers.

Leader Consensus as a Signal. Conditional on a leader's type $p \in \{H, L\}$, the leader's forecast can be written as:

$$a_{j,p} - \mu_0 = (\alpha_p^j + \beta_p^j)(\theta - \mu_0) + \alpha_p^j(\xi + \eta_j) + \beta_p^j u_j \tag{22}$$

where α_p^j and β_p^j are the equilibrium weights when solving the leader's problem.

Thus, conditional on p, a_j is an additional signal about the fundamental θ and the common bias ξ, with variance:

$$\text{Var}(a_j \mid \theta, q_j = p) = (\alpha_p^j)^2(\sigma_\xi^2 + \sigma_{\eta p}^2) + (\beta_p^j)^2\sigma_u^2 \tag{23}$$

However, the leader consensus and each follower's own biased signal are correlated conditional on θ, since both contain information about the common bias ξ:

$$\text{Cov}(y_i, a_j \mid \theta, q_i = q, q_j = p) = \alpha_p^j \sigma_\xi^2$$

Therefore, a_j cannot simply be treated as an independent third signal. However, it is possible to use the leader consensus incremental information as a third independent signal.

Incremental Information in the Leader Consensus. As in a linear regression, forecasters project the noise in a_j on the noise already observed in y_i in order to extract the orthogonalized error $e_{a,p}^{\perp}$. Let $e_{a,p} \equiv \alpha_p^j(\xi + \eta_j) + \beta_p^j u_j$ be the leader's forecast noise and $e_{y,q} \equiv \xi + \eta_i$ be the biased signal noise.

$$\lambda_{q,p} \equiv \frac{\text{Cov}(e_{a,p}, e_{y,q})}{\text{Var}(e_{y,q})} = \frac{\alpha_p^j \sigma_\xi^2}{\sigma_\xi^2 + \sigma_{\eta q}^2}, \quad e_{a,p}^{\perp} = e_{a,p} - \lambda_{q,p} e_{y,q} \tag{24}$$

The forecasters cannot directly observe $e_{a,p}$ and $e_{y,q}$ because they do not observe θ. However, they can replicate the regression structure with what they know. Conjecture $Z_{q,p}$ to be the orthogonalized leader's forecast:

$$Z_{q,p} \equiv \underbrace{(\bar{a}_j - \mu_0)}_{(\alpha_p^j + \beta_p^j)(\theta - \mu_0) + e_{a,p}} - \lambda_{q,p} \underbrace{(y_i - \mu_0)}_{\theta - \mu_0 + e_{y,q}} \tag{25}$$

Using the definitions in 24, Equation 25 can be re-written as:

$$Z_{q,p} = h_{q,p}(\theta - \mu_0) + e_{a,p}^{\perp}, \text{ where } h_{q,p} \equiv \beta_p^j + \alpha_p^j \frac{\sigma_{\eta q}^2}{\sigma_\xi^2 + \sigma_{\eta q}^2} \tag{26}$$

Finally, we normalize $Z_{q,p}$ into a standard signal of θ:

$$z_{j,qp} = \theta + v_{j,qp}, \quad \text{where } z_{j,qp} \equiv \frac{Z_{q,p}}{h_{q,p}} + \mu_0 \text{ and } v_{j,qp} \equiv \frac{e_{a,p}^{\perp}}{h_{q,p}} \tag{27}$$

Let $\Omega_{q,p} \equiv \text{Var}(e^{\perp}_{a,p})$[6], then the conditional variance of the incremental signal is $\text{Var}(z_{j,qp} \mid \theta) = \Omega_{q,p}/h^2_{q,p}$ and the precision is $\tau_{j,qp} = h^2_{q,p}/\Omega_{q,p}$.

Followers' Fundamental Bayesian Posterior. Each forecaster i uses her private signals (x_i, y_i) and the incremental leader's forecast signal $z_{j,qp}$ to update her belief about θ. Conditional on the leader's type being p, the follower's Bayesian posterior belief about the fundamental is:

$$m^{\theta}_{q,p} = (1 - \omega^{\alpha,\theta}_{q,p} - \omega^{\beta,\theta}_{q,p} - \omega^{\delta,\theta}_{q,p})\mu_0 + \omega^{\alpha,\theta}_{q,p} y_i + \omega^{\beta,\theta}_{q,p} x_i + \omega^{\delta,\theta}_{q,p} z_{j,qp} \tag{28}$$

where $\omega^{\alpha,\theta}_{q,p}$, $\omega^{\beta,\theta}_{q,p}$, and $\omega^{\delta,\theta}_{q,p}$ are the Bayesian weights.

$$\omega^{\alpha,\theta}_{q,p} = \frac{\tau_{yq}}{\tau_0(U) + \tau_{yq} + \tau_u + \tau_{j,qp}}, \; \omega^{\beta,\theta}_{q,p} = \frac{\tau_u}{\tau_0(U) + \tau_{yq} + \tau_u + \tau_{j,qp}}, \; \omega^{\delta,\theta}_{q,p} = \frac{\tau_{j,qp}}{\tau_0(U) + \tau_{yq} + \tau_u + \tau_{j,qp}}$$

It is possible to re-write the posterior belief in terms of a_j rather than the incremental signal $z_{j,qp}$[7], which is helpful in the later aggregation.

$$m^{\theta}_{q,p} = (1 - \alpha^{\theta}_{q,p} - \beta^{\theta}_{q,p} - \delta^{\theta}_{q,p})\mu_0 + \alpha^{\theta}_{q,p} y_i + \beta^{\theta}_{q,p} x_i + \delta^{\theta}_{q,p} \bar{a}_j \tag{29}$$

where

$$\alpha^{\theta}_{q,p} \equiv \omega^{\alpha,\theta}_{q,p} - \omega^{\delta,\theta}_{q,p} \frac{\lambda_{q,p}}{h_{q,p}}, \quad \beta^{\theta}_{q,p} = \omega^{\beta,\theta}_{q,p}, \quad \delta^{\theta}_{q,p} \equiv \frac{\omega^{\delta,\theta}_{q,p}}{h_{q,p}} \tag{30}$$

Followers' Common Bias Bayesian Posterior. Each follower generates a belief about ξ using both the biased signal y_i and the leader's forecast $\bar{a}_j$. Through Gaussian signal extraction using

[6] The derivation of $\Omega_{q,p}$ is in Appendix A.5.

[7] See Appendix A.6 for the derivation of Equations 29 and 30.

(y_i, a_j) and the posterior belief about the fundamental $m^{\theta}_{q,p}$:

$$\begin{aligned} m^{\xi}_{q,p} = & \frac{\sigma^2_{\xi} V^j_p}{(\sigma^2_{\xi} + \sigma^2_{\eta q}) V^j_p + (\alpha^j_p)^2 \sigma^2_{\xi} \sigma^2_{\eta q}} (y_i - m^{\theta}_{q,p}) \\ & + \frac{\alpha^j_p \sigma^2_{\xi} \sigma^2_{\eta q}}{(\sigma^2_{\xi} + \sigma^2_{\eta q}) V^j_p + (\alpha^j_p)^2 \sigma^2_{\xi} \sigma^2_{\eta q}} [a_j - \mu_0 - (\alpha^j_p + \beta^j_p)(m^{\theta}_{q,p} - \mu_0)] \end{aligned} \tag{31}$$

Because $m^{\theta}_{q,p}$ is affine in $(y_i, x_i, z_{j,pq})$, $m^{\xi}_{q,p}$ is affine as well.

$$m^{\xi}_{q,p} = \alpha^{\xi}_{q,p}(y_i - \mu_0) + \beta^{\xi}_{q,p}(x_i - \mu_0) + \delta^{\xi}_{q,p}(z_{j,qp} - \mu_0) \tag{32}$$

Unknown Leader Consensus Type. The previous analysis was conducted conditional on knowing the leader's type p. However, followers do not know the type of the leader. Instead, they observe the probability that the leader is a high-type φ_j. Thus, followers first calculate beliefs under each possible leader signal-processing ability and then combine them using the public reputation weights[8].

$$m^{\theta}_{q} = \varphi_j m^{\theta}_{q,H} + (1 - \varphi_j) m^{\theta}_{q,L} \tag{33}$$

$$m^{\xi}_{q} = \varphi_j m^{\xi}_{q,H} + (1 - \varphi_j) m^{\xi}_{q,L} \tag{34}$$

Since the conditional Bayesian posteriors are affine in the observed signals, their reputation-weighted averages are also affine. The resulting posterior beliefs can therefore be expressed as functions of the followers' private signals and the leader's forecast. The corresponding coefficients are derived in Appendix A.8.

External Evaluator and Follower's Problem. As in the baseline model, the external evaluator uses the consensus to recover the common bias ξ and to update reputation. Followers use their

[8] In the baseline model, we assume that followers do not update φ_j using a_j or (x_i, y_i). For our testable implications, we restrict attention to economies with $\varphi_j \in \{\varepsilon, 1 - \varepsilon\}$, where $\varepsilon > 0$ is small. When prior beliefs are sufficiently concentrated, Bayesian posterior probabilities remain close to their initial values on average. Thus, fixed reputation weights provide a limiting approximation to Bayesian updating as $\varepsilon \to 0$.

posterior beliefs to form expectations about how each type is going to be benchmarked by the external evaluator. See Appendix A.9 for the details.

Definition (Sequential Affine Rational Expectations Equilibrium). Given the public states (U, ρ, φ_j), an affine rational expectations equilibrium consists of type-contingent reporting strategies for leaders and followers, their associated posterior beliefs, and aggregate forecasts satisfying the following conditions:

1. Each leader's reported forecast solves the optimization problem in Equation 16, taking the evaluator's reputation-updating rule as given.

2. Each follower's reported forecast solves the optimization problem in Equation 51, conditional on her private information and the observed leader's forecast, taking the evaluator's reputation-updating rule as given and the aggregate consensus as given.

3. Leaders form Bayesian posterior beliefs about the fundamental and common bias conditional on their private signals and type. Followers form type-conditional Bayesian posterior beliefs, combining them using the fixed reputation-weighted inference rule in Equations 33 and 34.

4. The external evaluator's conjectured type-contingent reported strategies coincide with forecasters' optimally reported strategies in equilibrium.

5. The overall consensus coincides with the population averages of the equilibrium reporting strategies.

An equilibrium is therefore characterized by a fixed point in the strategic reporting coefficients, such that conjectured and optimal reporting strategies coincide[9].

3.1 Following a Reputable Leader

Introducing a leader distorts followers' optimal forecast with the reputation channel of the distortion depending on how agents perceive the leader. When the leader has a low reputation

[9] In Appendix B, we discuss in detail the solution technique implemented to solve for the fixed point.

($\varphi = 0$), both follower types underweight the leader and overweight the biased signal relative to the Bayesian benchmark. Under this scenario, the best way to signal quality is to differentiate themselves from the leader and increase the reliance on the biased signal.

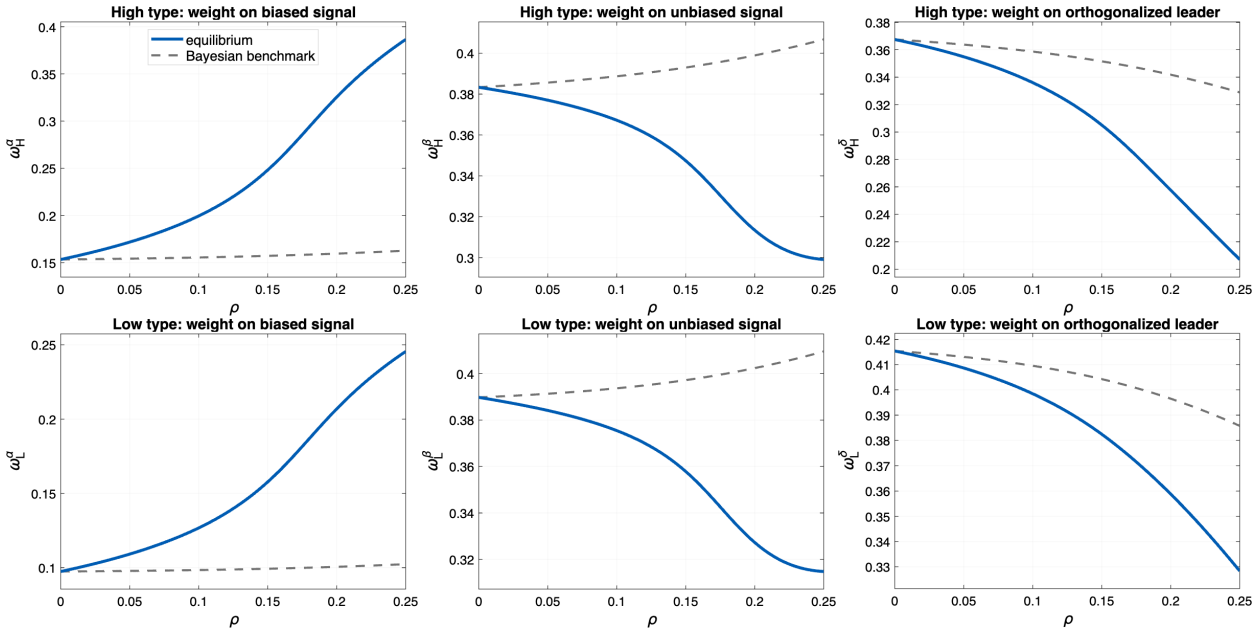


Figure 3: Reputation Concerns and Followers' Equilibrium Weights with a Low-Reputation Leader
Notes: The figure shows how reputation concerns affect followers' equilibrium reporting strategies when the leader has low reputation ($\varphi_j = 0$). The top (bottom) panels correspond to high-type (low-type) followers. The horizontal axis represents the reputation concerns state (ρ). Solid blue curves depict equilibrium reporting weights, while dashed gray curves represent the corresponding Bayesian benchmarks. As reputation concerns increase, both follower types place greater weight on the biased signal and reduce their reliance on the leader, reflecting incentives to differentiate themselves from a low-reputation leader.

In Figure 3, we plot the equilibrium signal weights for both follower types when the leader has low reputation ($\bar{\varphi}_j = 0$) under different reputation concerns states. The dashed gray line is the Bayesian signal weights for a given level of ρ. As reputation concerns increase, the leader relies more on the biased signal[10]. The declining dashed gray lines in the right panels indicate that both types of followers should rely less on the incremental leader signal in a Bayesian framework, once it would be carrying a lot of the bias ξ. In addition, the Bayesian benchmark suggests that followers should place more weight on their unbiased signals as they do not comprise the common bias ξ.

The blue lines in Figure 3 display the followers' optimal behavior when accounting for followers' reputation concerns. Followers heavily underweight the leader as reputation concerns arise, since

[10]Recall that the leader solves the exact same problem as in the baseline model, so Proposition 1 also holds.

they want to distance themselves from the low-type benchmark. Relative to the Bayesian benchmark, reputation concerns amplify the underweighting of the incremental leader signal. Moreover, followers overweight the biased signal and underweight the unbiased signal, making followers behave in the opposite direction from what a Bayesian forecaster would do.

Followers' optimal forecasts become non-monotonic when the leader has a high reputation $\bar{\varphi}_j = 1$, as shown in Figure 4. For low levels of ρ, followers still marginally underweight the incremental leader consensus and the unbiased signal at the expense of overweighting the biased signal. The deviation, nonetheless, is small and the optimal weights are close to the Bayesian benchmark. Followers still marginally underweight the leader when reputation concerns are small because overweighting the leader's forecasts could signal low quality. The intuition is that a high-type analyst would not need to mimic what a high-reputable leader said, since they could infer the same from their own high-type bias processing technology. Therefore, for low levels of reputation concerns, overweighting the incremental information from the orthogonalized leader consensus does not convey quality. Instead, the reputation incentive is to demonstrate independent expertise.

The strategy reverses at higher levels of reputation concerns. Figure 5 helps explain this reversal by comparing the weights that high- and low-type followers place on the leader's forecast. When reputation concerns are low, high-type followers place less weight on the leader than low-type followers, both in the Bayesian benchmark and in equilibrium. However, as reputation concerns increase, the leader increasingly overweights the biased signal, changing the informational content of her forecast and how followers of different types use this information. As shown by the dashed gray line in the right panel, the Bayesian weight of high-type followers eventually exceeds that of low-type followers. Hence, relying more heavily on the leader becomes increasingly characteristic of high-type forecasters. Reputation concerns amplify this incentive when the leader is highly reputable, inducing followers to overweight the leader's forecast relative to their Bayesian benchmark. This generates the sharp reversal in equilibrium reporting strategies observed in Figure 4.

The previous results illustrate how reputation concerns affect followers' reliance on earlier

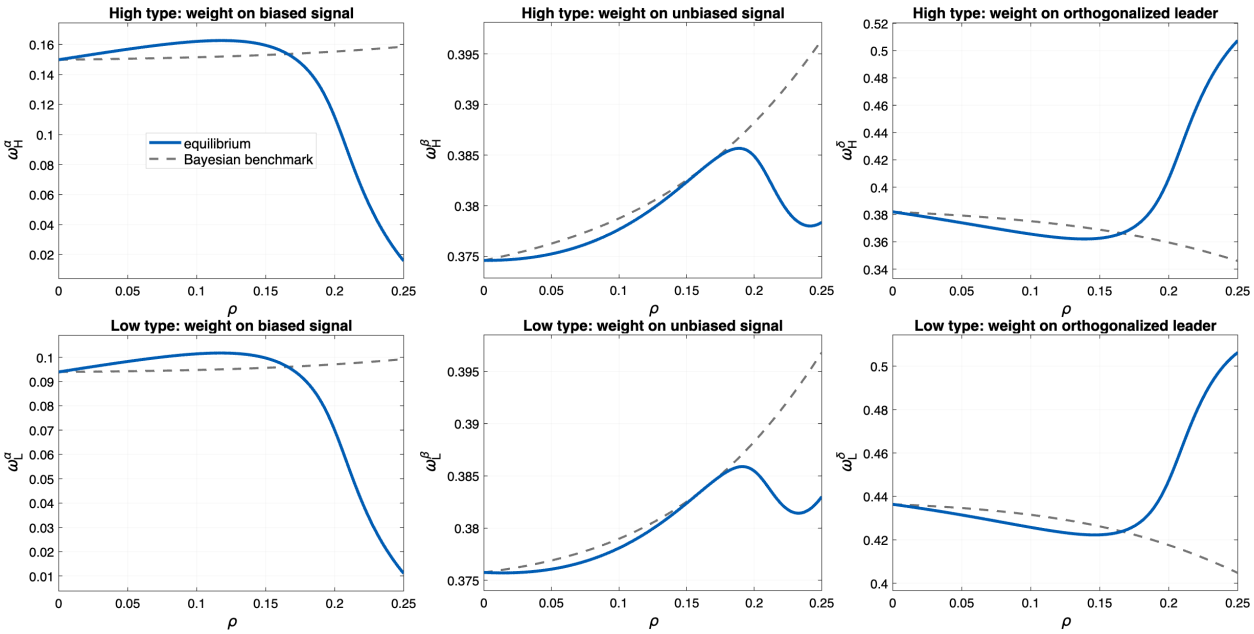


Figure 4: Reputation Concerns and Followers' Equilibrium Weights with a High-Reputation Leader
Notes: The figure shows how reputation concerns affect followers' equilibrium reporting strategies when the leader has a high reputation ($\varphi_j = 1$). The top (bottom) panels correspond to high-type (low-type) followers. The horizontal axis represents the reputation concerns state (ρ). Solid blue curves depict equilibrium reporting weights, while dashed gray curves represent the corresponding Bayesian benchmarks. As reputation concerns increase, both follower types place greater weight on the leader and reduce their reliance on the biased signal, reflecting incentives to align their forecasts with a high-reputable leader.

forecasts, with the direction of these distortions depending on the reputation of the leader. Followers have incentives to differentiate themselves from low-reputation leaders, while sufficiently strong reputation concerns induce them to conform to highly reputable ones. These differences in reporting strategies generate two additional testable implications.

Testable Implication #2: Following Reputable Leaders. *Followers place greater weight on previously submitted forecasts when these come from highly reputable leaders.*

In our model, followers observe both the leader's forecast and her reputation before submitting their own forecasts. Since the reputation of the leader provides information about the leader's forecasting ability, it affects how followers incorporate previously submitted forecasts into their predictions. As illustrated in Figure 6, followers place greater weight on the leader consensus when it has high reputation. The difference arises both from the informational value of highly reputable leaders and from reputational incentives that encourage followers to resemble high-type forecasters while differentiating themselves from low-type forecasters. Consequently, our model

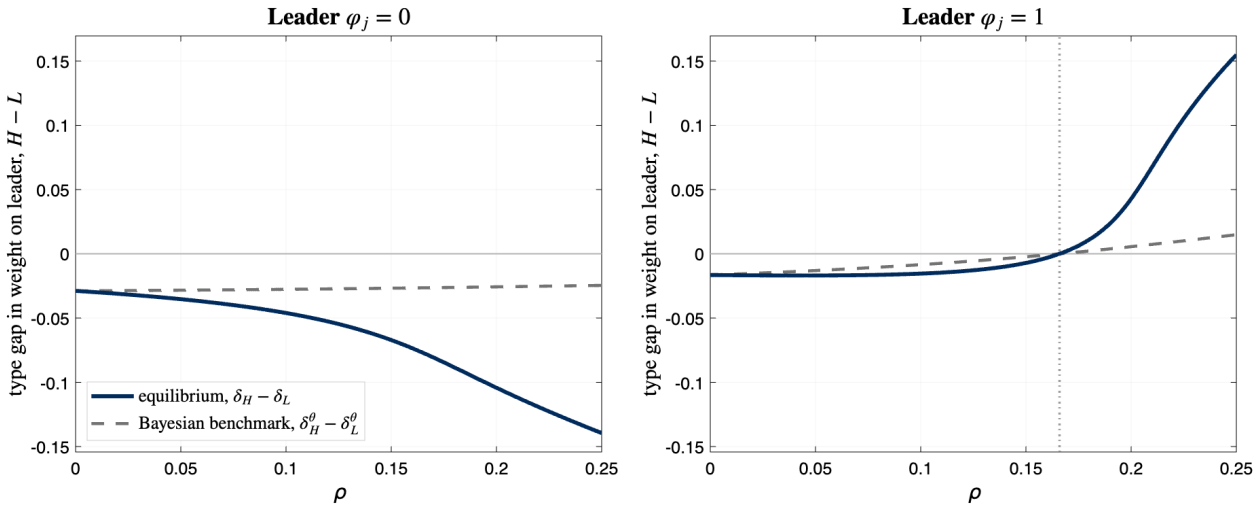


Figure 5: Reputational Concerns and the Reversal in Followers' Reliance on the Leader

Notes: The figure illustrates how reputation concerns affect the difference between high- and low-type followers' weights on the leader's forecast. The left (right) panel corresponds to an economy with a low-reputation (high-reputation) leader. The horizontal axis represents reputation concerns (ρ), while the vertical axis reports the difference in the weights assigned to the leader's forecast, ($\delta_H - \delta_L$). Solid blue curves depict equilibrium reporting weights, while dashed gray curves represent the corresponding Bayesian benchmarks, accounting for the leader's equilibrium reporting strategy at each level of ρ. A positive gap indicates that high-type followers place greater weight on the leader than low-type followers. With a low-reputation leader, the equilibrium gap remains negative and widens as reputation concerns increase. With a high-reputation leader, the gap reverses at sufficiently high levels of reputation concerns, becoming increasingly positive. Although the Bayesian benchmark also exhibits a reversal, reputational incentives substantially amplify the difference between follower types.

predicts that professional forecasters should rely more heavily on previously submitted forecasts when these come from highly reputable forecasters, which is represented by the positive level difference between the blue and orange lines at every ρ in Figure 6.

Testable Implication #3: Reputation Anchoring. *The difference in followers' reliance on high- and low-reputation leaders increases with reputation concerns.*

The model further predicts that the importance of the leader reputation for followers' reporting strategies varies with the strength of their reputation concerns. When reputation concerns are low, followers' reliance on the leader consensus is primarily determined by its informational value. As reputation concerns increase, followers face stronger incentives to differentiate themselves from low-reputation leaders or conform to highly reputable ones. As illustrated in Figure 6, these incentives amplify the difference in the weight followers place on the leader consensus across the two reputation environments. Consequently, we expect the influence of highly reputable forecasters

on subsequent forecasts to become more pronounced when reputation concerns are elevated, which is represented by the increasing gap between the blue and orange curves as ρ increases in Figure 6.

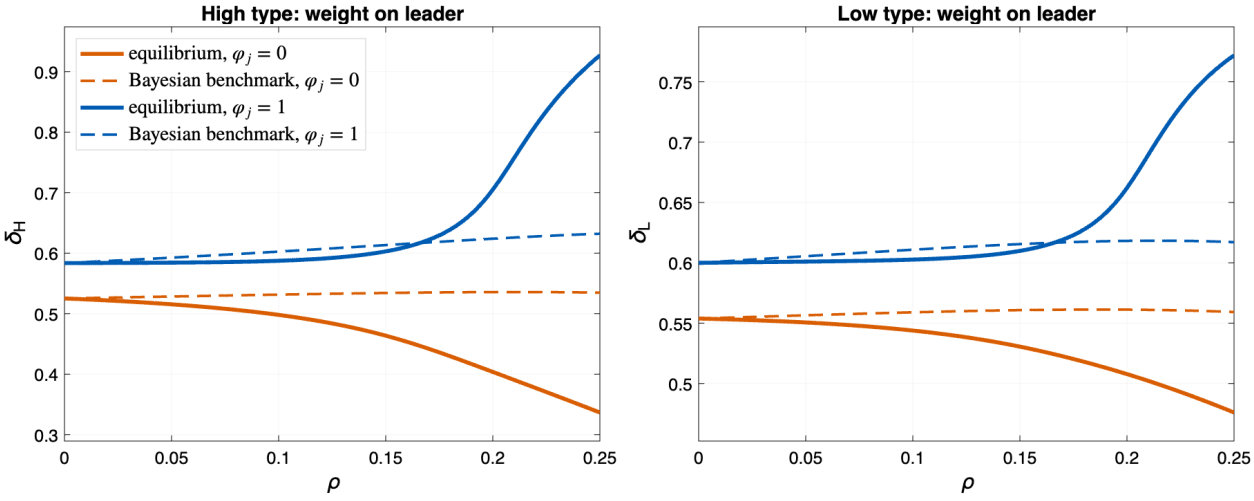


Figure 6: Following Leaders and Reputation Concerns

Notes: The figure illustrates how the reputation of the leader consensus and reputation concerns affect followers' equilibrium weight on the leader consensus. The left (right) panel corresponds to high-type (low-type) followers. The horizontal axis represents the reputation concerns state ρ, while the vertical axis measures the weight assigned to the leader consensus δ_q. Blue (orange) curves correspond to a high-reputation (low-reputation) leader consensus. Solid lines depict equilibrium reporting weights, while dashed lines represent the corresponding Bayesian benchmarks. Followers place greater weight on highly reputable leader forecasts across all levels of reputation concerns. Moreover, as reputation concerns increase, followers increasingly rely on high-reputation leaders while reducing their reliance on low-reputation leaders, widening the difference between the two environments.

3.2 Information Aggregation with Reputation Sequentiality

Having established how reputation concerns affect followers' reporting strategies, we now examine the implications of sequentiality for the quality of information contained in the overall consensus. Relative to the simultaneous economy, sequentiality introduces two competing forces. On the one hand, observing the leader's forecast provides followers with additional information about the fundamental. On the other hand, their reliance on the leader creates a new source of correlated forecasting errors that cannot be diversified away through aggregation.

In the baseline model, idiosyncratic signal errors disappear from the overall consensus because forecasters submit their predictions independently. With sequentiality, however, the leader's idiosyncratic errors influence all followers who incorporate her forecast into their own predictions.

Conditional on the leader's type p, the mean-squared error of the consensus can be written as

$$\mathcal{M}_p = \left[1 - \bar{\alpha} - \bar{\beta} - \bar{\delta}(\alpha_p^j + \beta_p^j)\right]^2 v(U) + (\bar{\alpha} + \bar{\delta}\alpha_p^j)^2 \sigma_\xi^2 + \bar{\delta}^2 V_p^j \tag{35}$$

where $V_p^j = (\alpha_p^j)^2 \sigma_{\eta p}^2 + (\beta_p^j)^2 \sigma_u^2$ is the leader's conditional forecast variance.

The first two terms capture the sensitivity of the consensus to the fundamental and its exposure to the common bias, as in the baseline model. The third term is specific to sequentiality and reflects the propagation of the leader's idiosyncratic signal errors. Although the leader has negligible mass, her errors become aggregate errors through followers' reporting decisions. The effect of sequentiality on information aggregation depends on whether the additional information provided by the leader compensates for the correlated errors introduced by following her forecast.

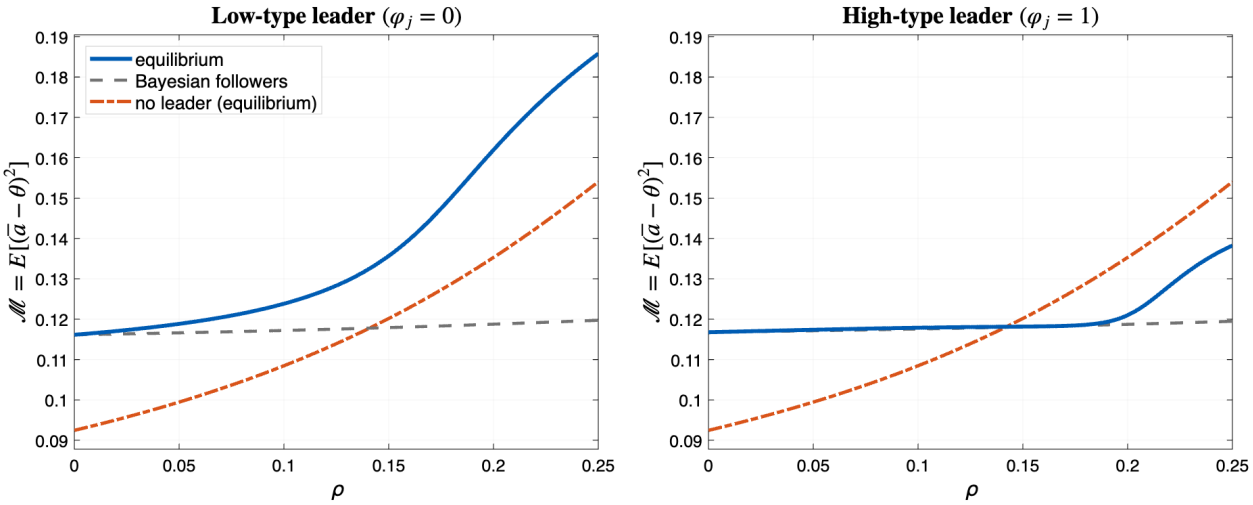


Figure 7: Sequential Forecasting, Reputation Concerns, and Information Aggregation

Notes: The figure illustrates how sequentiality affects the mean-squared error (MSE) of the consensus at different levels of reputation concerns. The left (right) panel corresponds to an economy with a low-reputation (high-reputation) leader. The horizontal axis represents reputation concerns ρ, while the vertical axis measures consensus MSE, with higher values indicating worse information aggregation. Solid blue curves depict the sequential equilibrium, dashed orange curves represent the simultaneous equilibrium without a leader, and dashed gray curves represent the benchmark with Bayesian followers, evaluated using the leader's equilibrium reporting strategy at each level of ρ. With a low-reputation leader, sequentiality increases consensus MSE relative to the simultaneous economy. With a high-reputation leader, sequentiality initially worsens information aggregation but improves it relative to the simultaneous economy when reputation concerns are sufficiently high.

Figure 7 illustrates these competing forces by comparing the sequential equilibrium with the simultaneous economy at different levels of reputation concerns. The dashed orange line represents

the equilibrium without sequentiality, while the blue line corresponds to the sequential equilibrium. We also report the benchmark with Bayesian followers, represented by the dashed gray line, to distinguish the informational value of observing the leader from the additional distortions generated by followers' reputational incentives.

When the leader has low reputation, sequentiality deteriorates information aggregation relative to the simultaneous economy across all levels of reputation concerns. Even when reputation concerns are absent, the leader's idiosyncratic errors propagate through followers and reduce the accuracy of the consensus. As reputation concerns increase, this deterioration becomes more pronounced. As shown in Figure 3, followers increasingly overweight their biased signal while reducing their reliance on the unbiased private signal. This increases the exposure of the consensus to the common bias, while the remaining dependence on the leader allows her idiosyncratic errors to persist in the aggregate.

When the leader has high reputation, the consequences of sequentiality depend on the level of reputation concerns. For low values of ρ, the sequential economy has a higher consensus MSE than the simultaneous economy, as the additional information from the leader does not compensate for the propagation of her idiosyncratic errors. However, this relationship reverses when reputation concerns become sufficiently high. As shown in Figure 4, followers increasingly rely on the highly reputable leader rather than overweighting their own biased signal. Since the leader has a greater ability to process the common bias, her forecast provides an alternative source of information that limits the deterioration in aggregate accuracy observed in the simultaneous economy. Consequently, sequentiality improves information aggregation relative to the simultaneous equilibrium when reputation concerns are sufficiently strong.

Importantly, this improvement does not imply that reputational distortions among followers are beneficial. In both reputation environments, equilibrium consensus MSE exceeds the benchmark with Bayesian followers. Rather, sequentiality can mitigate the deterioration generated by reputation concerns in the simultaneous economy, particularly when followers observe a highly reputable leader.

These results highlight that the consequences of sequentiality for information aggregation are not unambiguous. Although observing previous forecasts introduces additional information, it also allows individual forecasting errors to become aggregate errors. Whether sequentiality improves or deteriorates consensus accuracy depends on the reputation of earlier forecasters and the strength of the reputational incentives affecting subsequent reporting decisions.

4 Data

Our empirical analysis combines individual inflation forecasts from the monthly Bloomberg survey of professional forecasters with the inflation risk premium estimated by the Federal Reserve Bank of Cleveland. The Bloomberg survey allows us to observe how professional forecasters respond to previously submitted forecasts and how this behavior depends on the reputation of other forecasters. We use the inflation risk premium to study how forecasting behavior changes with the importance of inflation risk or, in other words, the salience of inflation to policymakers and market participants.

Bloomberg Inflation Forecasts. We use individual forecasts of the monthly change in the U.S. Consumer Price Index (CPI) collected by Bloomberg before each inflation announcement. The survey contains forecasts from professional economists working at financial institutions, including banks, broker-dealers, and research firms. Our sample covers 260 CPI announcements between January 2005 and October 2026, comprising 19,583 individual forecasts. On average, approximately 75 forecasters participate in each announcement.

There are three reasons for using the Bloomberg survey relative to other professional forecaster surveys. First, macroeconomic announcements surprises that move financial variables are calculated using the Bloomberg survey as the benchmark expectation (Elenev et al., 2024). Second, the Bloomberg survey allows forecasters to observe the forecasts of others before submitting their own. Finally, the survey also rewards forecasters by assigning a time-varying ranking for the Top 10 forecasters.

For each announcement t, we observe the forecast submitted by forecaster i, denoted by $\hat{\pi}_{i,t}$, the realized inflation rate π, and the date at which the forecast was submitted. The forecasts correspond to the seasonally adjusted month-over-month percentage change in headline CPI. Both forecasts and realized inflation are reported in increments of 0.1 percentage points. We identify individual forecasters by their names and institutions, allowing us to follow their forecasting behavior over time. Since forecasters enter and exit the survey, the resulting panel is unbalanced.

An important feature of the Bloomberg survey is that we can observe when each forecast was submitted. Most forecasts are submitted during the week preceding the CPI announcement, with a median submission date of five days before the release[11]. This allows us to construct the information available from previously submitted forecasts when a forecaster makes her prediction.

For each forecaster i and announcement t, we construct a leader consensus $\tilde{\pi}_{i,t}$, as the median of all forecasts submitted strictly before forecaster i. We require at least three previous forecasts to construct the leader consensus. Since submission dates are observed at a daily frequency, we do not include forecasts submitted on the same day[12]. The measure $\bar{\pi}_{i,t}$ is the empirical counterpart of the theoretical leader's forecast a_j in the model, capturing the information from earlier forecasters that subsequent forecasters can observe. We use the absolute distance between individual forecasts and the leader consensus to measure the extent to which forecasters conform to previously submitted forecasts.

Bloomberg also provides a publicly available ranking of the best forecasters for each economic indicator. We use the Bloomberg Top-10 ranking as a measure of forecaster reputation. Since the ranking reported for a given announcement incorporates the accuracy of the forecast after inflation is realized, we use the ranking from the forecaster's previous appearance to measure her reputation at the time of submission. This allows us to identify whether highly reputable forecasters have already submitted their predictions when subsequent forecasters make their decisions. This would be the empirical counterpart of φ_j in the model.

[11] From anecdotal evidence, we take the submission timing as exogenous, with firms setting deadlines for analysts to submit their forecasts.

[12] Our understanding is that Bloomberg collects all forecasts submitted throughout a day and makes them publicly available once a day.

Finally, we construct measures of forecast accuracy and dispersion. We denote by $\bar{\pi}_t$ the overall Bloomberg consensus, measured by the median of individual forecasts for announcement t. Individual forecast accuracy is measured by $|\hat{\pi}_t - \pi_t|$, while consensus accuracy is measured by $|\bar{\pi} - \pi_t|$. We also calculate the cross-sectional standard deviation of forecasts within each announcement.

Table 1 reports summary statistics for the main variables. The average absolute individual forecast error is 0.108 percentage points, compared with 0.089 percentage points for the overall consensus. Cross-sectional forecast dispersion averages 0.088 percentage points. Approximately 48% of individual forecasts coincide exactly with the empirical leader consensus, showing that sequentiality and matching are important characteristics of the survey.

Table 1: Summary Statistics

Variable	Mean	SD	Median	*N*
Panel A. Announcement-level variables				
Realized inflation, π_t	0.213	0.342	0.200	260
Absolute consensus error, $\lvert\bar{\pi}_t - \pi_t\rvert$	0.089	0.094	0.100	260
Cross-sectional forecast dispersion	0.088	0.032	0.083	260
Inflation risk premium, IRP_{t-1}	0.397	0.086	0.413	260
Number of forecasters	75.3	–	75.0	260
Panel B. Forecaster-level variables				
Individual forecast, $\hat{\pi}_{i,t}$	0.207	0.284	0.200	19,583
Absolute forecast error, $\lvert\hat{\pi}_{i,t} - \pi_t\rvert$	0.108	0.112	0.100	19,583
Distance to empirical leader consensus, $\lvert\hat{\pi}_{i,t} - \tilde{\pi}_{i,t}\rvert$	0.062	0.075	0.050	18,026
Exact match with empirical leader consensus	0.483	–	–	18,026
Lagged Bloomberg Top-10 indicator	0.120	–	–	19,583
Days before CPI release	5.7	3.5	5.0	19,583

Notes: The table reports summary statistics for monthly U.S. CPI inflation forecasts collected by Bloomberg between January 2005 and October 2026. Panel A reports announcement-level variables, while Panel B reports individual forecaster-announcement observations. The Bloomberg consensus is the median forecast for each announcement. The empirical leader consensus is the median of forecasts submitted strictly before an individual forecaster, requiring at least three previous forecasts. The exact-match indicator equals one when an individual forecast coincides with the empirical leader consensus. The lagged Bloomberg Top-10 indicator is based on the ranking available from each forecaster's previous appearance. The inflation risk premium is the ten-year measure estimated by the Federal Reserve Bank of Cleveland, lagged to ensure that it was available before the CPI announcement. Inflation forecasts, errors, dispersion, and the IRP are expressed in percentage points.

Inflation Risk Premium. To measure the salience of inflation risk, we use the ten-year Inflation Risk Premium (IRP) estimated by the Federal Reserve Bank of Cleveland. The Cleveland Fed estimates the IRP through a term-structure model that combines nominal Treasury yields, inflation swaps, realized inflation, and survey-based inflation expectations, following Haubrich, Pennacchi and Ritchken (2012)[13]. The IRP measures the compensation investors require for bearing the risk that realized inflation differs from their expectations.

We use IRP as an empirical proxy for the model's reputation concerns ρ. The idea is that, when inflation risk becomes more important for financial market participants, the accuracy of professional inflation forecasts may also become more important for their clients. Consequently, forecasters may face stronger incentives to protect their reputation. In the model, these incentives are represented by the state variable ρ, which determines the importance forecasters assign to their future reputation. We interpret the IRP as an observable indicator of these conditions, rather than a direct measure of reputation concerns.

The IRP is available at a monthly frequency and is expressed in percentage points[14]. In our sample, the average ten-year inflation risk premium is 0.4 percent, with a standard deviation of 0.09 percentage points. We also classify announcements into five groups according to the distribution of IRP_{t-1}. We refer to the bottom two quintiles as low-IRP periods and the top two quintiles as high-IRP periods[15].

[13]The Bloomberg survey is not used to estimate IRP, which prevents any mechanical relationship between our empirical variables.

[14]The Cleveland Fed updates its estimates on the day of the CPI announcement, after the inflation data is released. Hence, we use the lagged IRP as our baseline time-consistent measure.

[15]The choice of quintiles is motivated by the distribution of IRP in our sample. The distribution is highly concentrated around its median, with 88% of observations between 0.25 and 0.5 percent. Consequently, a median split would compare announcements with relatively similar inflation risk premium. Excluding the middle quintile increases the separation between two regimes, while using the bottom and top two quintiles preserves equally sized groups and a larger sample than a comparison of the extreme terciles alone.

5 Empirical Results

In this section, we first document a few empirical facts that helps us to justify some of the features in our model and to argue that high Inflation Risk Premium periods are associated with high reputation concerns states. Then we assess the testable implications from the model to validate our reputation mechanisms.

5.1 Empirical Facts

We now introduce three empirical facts from the Bloomberg survey data that motivates a reputational friction to explain the behavior of professional forecasters. We argue that periods of high inflation risk premium are associated to times of high reputation concerns. The intuition comes from investors caring more about the forecasts when inflation risk premium is high, which turns future reputation more valuable for professional forecasters.

Fact 1 – Forecast dispersion decreases when inflation risk is high. There is a clear negative statistical relationship between the cross-sectional dispersion in the Bloomberg survey and inflation risk premium. As we can see from Figure 8, forecasters tend to agree more in times of high inflation risk premium. If IRP only reflected uncertainty about the fundamental, one would expect larger dispersion rather than herding and that mistakes would be concentrated at high IRP levels. Thus, inflation risk premium seems to comprise more than merely uncertainty about inflation.

Fact 2 – Mistakes are more costly when inflation risk is high. Forecasters drop more in the Bloomberg ranking when making a mistake during a high inflation risk period. From Table 2, a 10-bp mistake is associated with a 1.1 ranking position drop during low inflation risk states, while a mistake of the same magnitude drops the rank by 1.7 ranking positions when IRP is high. The difference is statistically significant. Mistakes by the consensus increase a forecaster's reputation when keeping their accuracy fixed, particularly during high inflation risk times. Fact 2 could be naturally explained when taking Fact 1 into account: forecasts are more concentrated when inflation

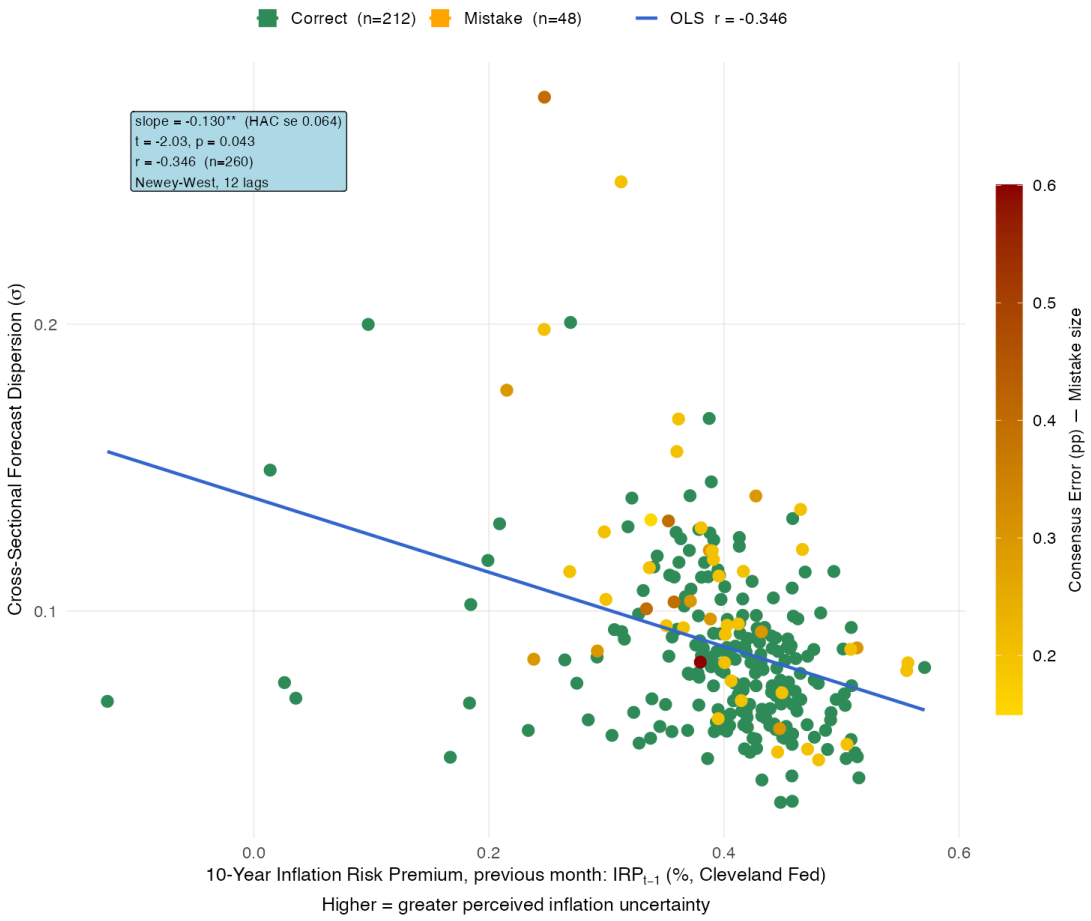


Figure 8: Inflation Risk Premium and Cross-Sectional Forecast Dispersion

Notes: The figure illustrates the relationship between the Cleveland Fed's 10-year Inflation Risk Premium (IRP) and the cross-sectional dispersion of professional inflation forecasts from the Bloomberg survey. The horizontal axis reports the previous month's IRP, while the vertical axis measures forecast dispersion as the standard deviation of individual forecasts within each announcement. Green marks indicate announcements without consensus mistakes, while yellow-to-red markers indicate announcements with consensus mistakes. Darker colors represent larger absolute forecast errors. The solid blue line represents the OLS linear fit. The correlation is negative and statistically significant $r = -0.346$*,* $p < 0.001$*, indicating that professional forecasters tend to be less dispersed during period of elevated inflation risk.*

risk is high, causing deviations from the norm to drastically change the ranking. However, rather than claiming causality, we aim to use the facts to explore the association between inflation risk premium and reputation concerns.

Fact 3 – More herding, without greater accuracy. Although forecasters exhibit greater convergence during periods of elevated inflation risk, this behavior does not appear to be accompanied by

Table 2: Reputational Cost of Forecast Errors by IRP Regime

	Bloomberg ranking, $\Delta\varphi_{i,t+1}$			Pseudo-ranking, $\Delta\hat{\varphi}_{i,t+1}$		
	Low IRP	All	High IRP	Low IRP	All	High IRP
	(1)	(2)	(3)	(4)	(5)	(6)
Own error	-0.107^{***}	-0.139^{***}	-0.169^{***}	-0.172^{***}	-0.195^{***}	-0.219^{***}
	(0.012)	(0.013)	(0.020)	(0.016)	(0.015)	(0.020)
Consensus error	0.076^{***}	0.093^{***}	0.105^{***}	0.136^{***}	0.144^{***}	0.146^{***}
	(0.010)	(0.011)	(0.018)	(0.014)	(0.011)	(0.015)
High − Low IRP, own error		-0.062^{***}			-0.047^{**}	
		(0.022)			(0.021)	
		[$p = 0.005$]			[$p = 0.027$]	
Observations	544	1,356	531	2,790	6,908	2,665
Forecasters	64	79	71	160	166	151
R^2	0.155	0.154	0.152	0.200	0.185	0.169

Notes: OLS regressions of the change in a forecaster's reputation score between consecutive survey appearances, $\Delta\varphi_{i,t+1} = \varphi_{i,t+1} - \varphi_{i,t}$, on the forecaster's own absolute error, $|\hat{\pi}_{i,t+1} - \pi_{t+1}|$, and the absolute error of the Bloomberg (median) consensus, $|\bar{\pi}_{t+1} - \pi_{t+1}|$, at announcement $t + 1$, both in basis points. $\varphi = 11 -$ rank for forecasters in Bloomberg's Top-10 and zero otherwise (columns 1–3); the pseudo-ranking ranks all forecasters with at least 18 forecasts by mean absolute error over the trailing 24 announcements, $\hat{\varphi} = 11 -$ rank (columns 4–6). The sample includes forecaster-announcements with a complete 24-announcement history and a forecast at $t + 1$; columns 1–3 further require the forecaster to be ranked at t or $t + 1$. IRP regimes classify the forecast for announcement $t + 1$ by the quintile (across announcements) of IRP_t, the inflation risk premium released with announcement t and hence known, together with $\varphi_{i,t}$, before the forecast is made: Low = Q1–Q2, High = Q4–Q5; columns 2 and 5 use all announcements. The High − Low row reports the interaction of the own error with a High-IRP indicator in a pooled regression on the Low and High samples, with both regressors interacted. Standard errors clustered by forecaster in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

a greater accuracy advantage from following the empirical leader consensus[16]. To investigate this relationship, we examine how individual forecast errors vary with the distance between individual forecasts the previously submitted consensus. We find that deviations from the leader consensus are generally associated with larger forecasting errors, but this relationship weakens significantly when IRP is high. This suggests that the increased conformity documented in Fact 1 is not explained by a corresponding improvement in the relative accuracy of the leader consensus. Together with the larger reputational penalties documented in Fact 2, this pattern is consistent with reputational incentives contributing to forecasting convergence, beyond purely informational considerations.

[16]See table with results in Appendix C.1.

Fact 4 – Consensus Accuracy Explains Ranking Variation. The empirical counterpart of reputation in the model is the Bloomberg ranking. We investigate whether changes in ranking are solely associated with a forecaster's own accuracy or whether the behavior of others, in particular of the consensus, also plays a role. Table 2 provides initial evidence that reputation updates depend on relative forecasting performance. Conditional on a forecaster's own absolute error, a larger consensus forecast error is associated with a more favorable change in her ranking. In other words, making a forecasting mistake appears to be less costly to reputation when other forecasters also make mistakes. This finding is consistent with our model, in which the evaluator uses information contained in the aggregate consensus when assessing individual forecasting ability.

To further assess the importance of relative performance, we examine the incremental explanatory power of consensus accuracy for changes in reputation. As shown in Table 3, the adjusted R^2 increases from 17% when considering only a forecaster's own accuracy to 31% when also including consensus accuracy. When generating a pseudo-ranking for the entire sample of forecasters[17], we got a similar result with the adjusted R^2 increasing from 27.3% to 45.3%. Thus, the accuracy of the aggregate consensus provides substantial additional explanatory power for reputation changes, beyond individual performance. Although this relationship partly reflects the relative nature of the Bloomberg ranking, it is consistent with the central mechanism of our model and with Testable Implication 1: reputation evaluations depend not only on individual forecast accuracy but also on the performance of other forecasters.

5.2 Empirically Assessing Our Testable Implications

Testable Implication #1: Reputation Concerns and Lower Dispersion. We proxy the model's reputation concerns state with the Cleveland Fed's inflation risk premium. As documented in Fact 1, periods of elevated inflation risk are associated with lower cross-sectional dispersion of professional forecasts. To assess this relationship more directly, Figure 9 plots the cross-sectional variance of individual Bloomberg forecasts against the previous month's IRP. This measure corresponds more

[17] We follow the broad guidelines provided by Bloomberg on how they calculate the ranking.

Table 3: Variance Decomposition of Reputation Changes

	Bloomberg ranking (Top-10 sample)		Pseudo-ranking (all forecasters)	
	R^2	Adj. R^2	R^2	Adj. R^2
Own accuracy only	18.5	17.0	27.6	27.3
Consensus accuracy only	0.5	−1.4	0.5	0.2
Both	33.5	31.0	45.7	45.3
Incremental own (Both − Consensus only)	33.0		45.2	
Incremental consensus (Both − Own only)	15.0		18.2	
Observations	1,356		6,908	
Forecasters	79		166	

Notes: Share of the variance of the reputation change $\Delta\varphi_{i,t} = \varphi_{i,t+1} - \varphi_{i,t}$ explained by OLS regressions, in percent. "Own accuracy" regressors are forecaster i's absolute errors $|\hat{\pi}_{i,s} - \pi_s|$ and "consensus accuracy" regressors are the absolute errors of the Bloomberg survey median, $|\bar{\pi}_s - \pi_s|$, each entered separately for the 25 announcements $s = t-23, \ldots, t+1$ (25 regressors per block, 50 in "Both"). Rankings are end-of-period: $\varphi_{i,t}$ reflects accuracy over announcements $t - 23$ to t, so $\Delta\varphi_{i,t}$ reflects the error at $t + 1$ entering and the error at $t - 23$ leaving the window. Bloomberg ranking: $\varphi_{i,t} = 11 -$ rank for the Bloomberg Top 10 and 0 otherwise; the sample keeps forecasters ranked at t or $t + 1$. Pseudo-ranking: $\varphi_{i,t} = 11 -$ rank over all forecasters with at least 18 forecasts in the window, ranked by mean absolute error over announcements $t - 23$ to t. Both samples require forecasts for all 25 consecutive announcements. Incremental rows are differences in unadjusted R^2.

closely to the expected cross-sectional variance characterized in our theoretical model.

We find a negative relationship between forecast variance and the inflation risk premium. Importantly, this relationship does not establish that reputation concerns cause forecast dispersion to decline. The inflation risk premium may reflect changes in both fundamental uncertainty and the importance of forecasting performance to market participants. Moreover, while the model's prediction concerns expected dispersion conditional on the uncertainty state (U), our empirical measure captures realized dispersion across announcements. Nevertheless, the negative association provides suggestive evidence consistent with reputational incentives affecting the aggregation of information in professional forecasts.

Testable Implication #2: Following Reputable Leaders. We investigate whether professional forecasters place greater weight on previously submitted forecasts when these come from highly reputable forecasters. To proxy for the reputation of the leader consensus, we construct an indicator equal to one when at least one forecaster previously ranked among Bloomberg's Top 10 has

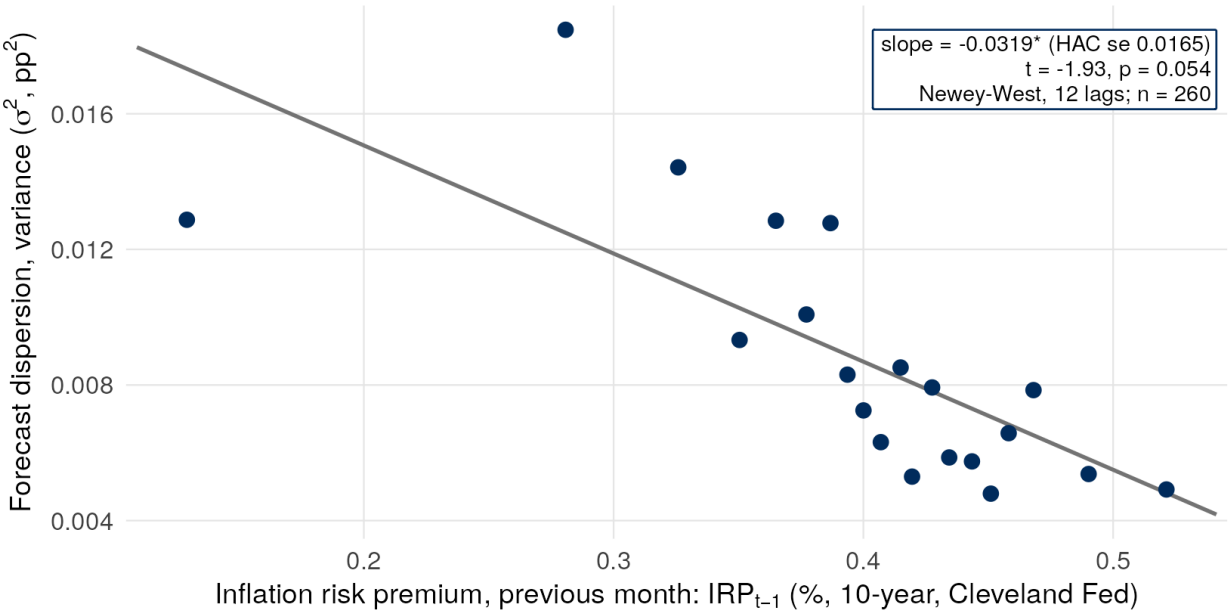


Figure 9: Inflation Risk Premium and Cross-Sectional Forecast Variance

The figure presents a binscatter illustrating the relationship between the Cleveland Fed's 10-year Inflation Risk Premium (IRP) and the cross-sectional variance of professional inflation forecasts from the Bloomberg survey. The horizontal axis reports the previous month's IRP, while the vertical axis measures forecast dispersion as the cross-sectional variance of individual forecasts within each CPI announcement, expressed in squared percentage points. Each dot represents the average IRP and forecast variance within a bin of announcements grouped by IRP. The solid gray line represents the fitted linear regression using the underlying 260 monthly observations rather than the binned averages.

submitted a forecast strictly before analyst i. We then examine whether the presence of a highly reputable forecaster in the preceding pool is associated with a smaller absolute distance between the analyst's prediction and the empirical leader consensus.

Table 4 reports the results controlling for analyst and announcement fixed effects, the number of preceding forecasts, and the time remaining until the CPI announcement. We find that the presence of a Top-10 forecaster in the preceding pool of submissions is associated with a reduction of approximately 0.9 basis points in the distance to the leader consensus. This relationship is statistically significant and is primarily driven by the extensive margin: forecasters become approximately 7.8 percentage points more likely to submit a forecast that exactly matches the previously observed consensus.

An alternative explanation is that the presence of highly reputable forecasters is correlated with the existence of more accurate information in the empirical leader consensus. To investigate

Table 4: Following Reputable Leaders

	Absolute deviation $d_{i,t}$		Probability of deviation $\mathbf{1}\{d_{i,t} > 0\}$	
	(1)	(2)	(3)	(4)
Top-10 in preceding pool	−0.0099***	−0.0091***	−0.0777***	−0.0778***
	(0.0028)	(0.0029)	(0.0189)	(0.0202)
Own Top-10 status	0.0025	0.0026	−0.0022	−0.0022
	(0.0020)	(0.0021)	(0.0160)	(0.0160)
Historical pool MAE	–	0.1308	–	−0.0246
		(0.1398)		(0.8421)
Historical pool MAE control		✓		✓
Other controls	✓	✓	✓	✓
Analyst fixed effects	✓	✓	✓	✓
Announcement fixed effects	✓	✓	✓	✓
Mean of dependent variable	0.0614	0.0614	0.5105	0.5105
Observations	16,881	16,881	16,881	16,881
Forecasters	373	373	373	373
Within R^2	0.0015	0.0016	0.0022	0.0022

Notes: OLS regressions of individual deviations from the empirical leader consensus on the presence of highly reputable forecasters in the preceding pool; columns (3)–(4) are linear probability models. In columns (1)–(2) the dependent variable is $d_{i,t} = |\hat{\pi}_{i,t} - \tilde{\pi}_{i,t}|$, the absolute difference (in percentage points) between analyst i's forecast and the empirical leader consensus $\tilde{\pi}_{i,t}$, the median of forecasts submitted on dates strictly before analyst i's (at least three required). In columns (3)–(4) it is an indicator equal to one when the forecast differs from $\tilde{\pi}_{i,t}$. Top-10 in preceding pool equals one when at least one forecaster ranked in the Bloomberg Top 10 as of the previous announcement submitted a forecast before analyst i; own Top-10 status is analyst i's ranking as of the previous announcement. Historical pool MAE is the average, over the preceding forecasters, of each forecaster's mean absolute forecast error over the 24 announcements up to and including that forecaster's previous survey appearance (at least 18 forecasts required), so it uses only information available before the current announcement. Other controls are the number of preceding forecasts and the number of days between the forecast and the CPI release. The sample is the 16,881 analyst-announcement observations for which historical pool MAE is defined, and is identical across columns (forecasters observed only once are dropped by the fixed effects). Standard errors clustered by analyst in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

this possibility, we control for the historical mean absolute error of the preceding submission pool, calculated using information available before the current announcement. The estimated relationship remains virtually unchanged, suggesting that the influence of highly reputable forecasters is not fully explained by this measure of historical forecasting ability. These findings are consistent with Testable Implication 2, whereby followers place greater weight on the leader consensus when it contains highly reputable forecasters.

We also examine whether the influence of reputable leaders depends on the follower's own type.

In our model, a forecaster's prior reputation enters additively into posterior reputation log-odds and therefore does not directly affect her optimal reporting strategy. Moreover, Figure 12 shows that high- and low-type followers exhibit similar responses to leader reputation, particularly when reputation concerns are relatively low. Consistent with these predictions, we find no statistically significance difference in their following behavior for ranked and non-ranked followers. This finding suggests that the influence of reputable leaders extends broadly across forecasters and is consistent with the model results.

Testable Implication #3: Reputation Anchoring. Our third testable implication predicts that the influence of leader reputation on subsequent forecasts becomes stronger as reputation concerns increase. As illustrated in Figure 12, this amplification is nonlinear, becoming particularly pronounced at sufficiently high levels of reputation concerns. To investigate this prediction, we examine whether the relationship between the presence of highly reputable forecasters in the preceding pool and subsequent forecast conformity varies with the inflation risk premium, which we use as a proxy for reputation concerns.

We first estimated a pooled regression that interacts the presence of a Top-10 forecaster in the leader consensus with an indicator for high-IRP periods, as shown in Figure 10. Although the estimated effect on the probability of deviating from the leader consensus is stronger during high-IRP periods, the difference is not statistically significant. However, the model suggests that the amplification of reputational incentives becomes particularly pronounced only at sufficiently high levels of reputation concerns. Motivated by this non-linearity, we investigate how the influence of reputable forecasters varies across the IRP distribution.

Figure 10 presents the results. Panel (a) shows that the presence of a Top-10 forecaster in the leader consensus is associated with stronger conformity in the highest IRP quintile (Q5) than in most intermediate IRP regimes. Importantly, Panel (b) shows that the estimated reputation-anchor effect is stronger and statistically significant than intermediate regions of the distribution (Q2-Q3 and Q2-Q4 pooled groups). These findings are consistent with our model's prediction

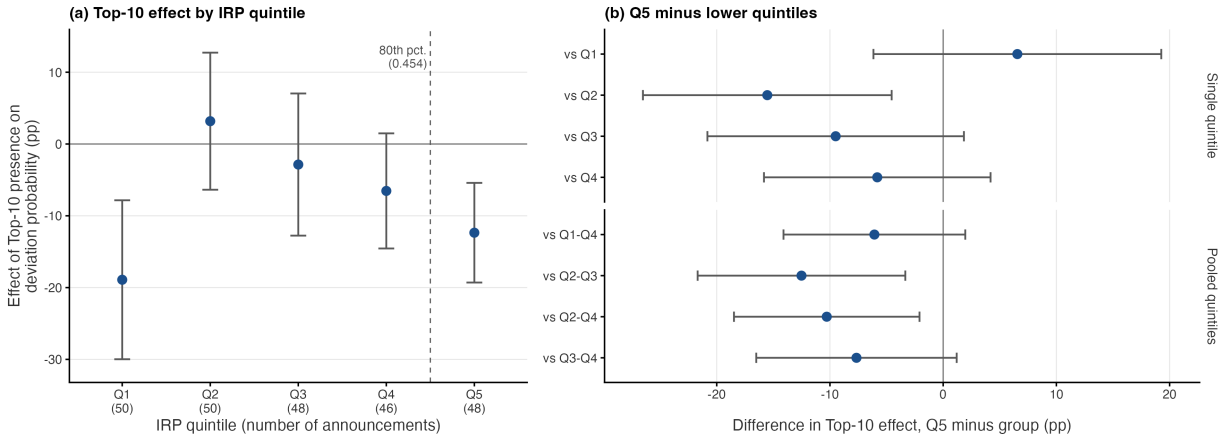


Figure 10: Inflation Risk and Reputation Anchoring

Notes: The figure illustrates how the influence of highly reputable forecasters on subsequent predictions varies with the inflation risk premium. Panel (a) reports the estimated effect of having at least one previously ranked Bloomberg Top-10 forecaster in the preceding pool on the probability of deviating from the empirical leader consensus, separately by IRP quintile. Panel (b) reports the differences between the estimated Top-10 effect in the highest IRP quintile (Q5) and the effects in lower quintiles, both individually and pooled. Negative coefficients indicate greater conformity to the preceding consensus when highly reputable forecasters are present. All specifications control for forecaster and announcement fixed effects, historical forecasting accuracy, accuracy of the preceding pool, submission timing, and the number of preceding forecasts. Vertical and horizontal bars represent 95% confidence intervals, with standard errors clustered by forecaster and announcement.

that reputation concerns amplify the influence of highly reputable forecasters at elevated levels of reputation concerns.

Nevertheless, the relationship is not monotonic throughout the IRP distribution, as substantial reputation anchoring is also observed in the lowest quintile. Further investigation suggests that this result is partly associated with episodes of abrupt declines and subsequent recoveries in the Clevelan Fed's IRP series. Moreover, the upper tail of IRP is concentrated in recent years (2022–2025), making it difficult to fully distinguish elevated inflation risk from other changes in the forecasting environment. We therefore interpret these findings as suggestive evidence of the nonlinear amplification predicted by our model, rather than that inflation risk captures reputation concerns across all macroeconomic environments.

6 Conclusion

A forecast consensus is valuable because it can bring together information that no single forecaster possesses and aggregate dispersed information and models used by experts. Reputational incentives can erode that advantage. When forecasters lean more heavily on a shared source of information to demonstrate their skill, they also give greater weight to its errors. Those errors survive aggregation. A panel of forecasters whose reports cluster more tightly can therefore give investors and policymakers a less informative view of the economy.

The strength of this distortion can vary over the cycle. When a forecast miss carries a larger reputational penalty, forecasters have more reason to favor reports that protect their standing. The consensus can then become less informative even when the information available to forecasters is just as good. This changes how markets can read a release.

The same mechanism can operate wherever expert judgments serve both as inputs to decisions and as records of professional ability. Earnings forecasts help investors value firms and help analysts build a reputation with clients. Credit assessments can similarly shape lending decisions while influencing how evaluators judge the analysts who produce them. In each setting, the way experts are evaluated can affect how much independent information reaches the people who rely on their collective judgment.

References

Baley, Isaac and Javier Turen. 2024. Lumpy forecasts. Economics Working Papers 1898 Department of Economics and Business, Universitat Pompeu Fabra.
URL: *https://ideas.repec.org/p/upf/upfgen/1898.html*

Bikhchandani, Sushil, David Hirshleifer and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100(5):992–1026.
URL: *http://www.jstor.org/stable/2138632*

Bostanci, Gorkem and Guillermo Ordoñez. 2026. "Business, Liquidity, and Information Cycles." *The Quarterly Journal of Economics* 141(3):2313–2361.
URL: *https://doi.org/10.1093/qje/qjag029*

Chousakos, Kyriakos, Gary Gorton and Guillermo Ordoñez. 2023. "Information Dynamics and Macro Fluctuations." *American Economic Journal: Macroeconomics* 15(4):372–400.
URL: *https://www.aeaweb.org/articles?id=10.1257/mac.20210101*

Elenev, Vadim, Tzuo-Hann Law, Dongho Song and Amir Yaron. 2024. "Fearing the Fed: How wall street reads main street." *Journal of Financial Economics* 153:103790.
URL: *https://www.sciencedirect.com/science/article/pii/S0304405X24000138*

Gemmi, Luca and Rosen Valchev. 2026. "Biased surveys." *Journal of Monetary Economics* 157(C):None.
URL: *https://ideas.repec.org/a/eee/moneco/v157y2026ics0304393225001394.html*

Gorton, Gary and Guillermo Ordoñez. 2014. "Collateral Crises." *American Economic Review* 104(2):343–78.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.104.2.343*

Graham, John R. 1999. "Herding among Investment Newsletters: Theory and Evidence." *The Journal of Finance* 54(1):237–268.
URL: *http://www.jstor.org/stable/222415*

Haubrich, Joseph, George Pennacchi and Peter Ritchken. 2012. "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps." *The Review of Financial Studies* 25(5):1588–1629.
URL: *https://ideas.repec.org/a/oup/rfinst/v25y2012i5p1588-1629.html*

Jurado, Kyle, Sydney C. Ludvigson and Serena Ng. 2015. "Measuring Uncertainty." *American Economic Review* 105(3):1177–1216.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.20131193*

Laster, David, Paul Bennett and In Sun Geoum. 1999. "Rational Bias in Macroeconomic Forecasts." *The Quarterly Journal of Economics* 114(1):293–318.
URL: *http://www.jstor.org/stable/2586954*

Lichtendahl, Kenneth C., Yael Grushka-Cockayne and Phillip E. Pfeifer. 2013. "The Wisdom of Competitive Crowds." *Operations Research* 61(6):1383–1398.
URL: *https://ideas.repec.org/a/inm/oropre/v61y2013i6p1383-1398.html*

Ordoñez, Guillermo. 2013. "The Asymmetric Effects of Financial Frictions." *Journal of Political Economy* 121(5):844–895.
URL: *https://www.jstor.org/stable/10.1086/673886*

Ottaviani, Marco and Peter Norman Sørensen. 2006*a*. "Professional advice." *Journal of Economic Theory* 126(1):120–142.
URL: *https://ideas.repec.org/a/eee/jetheo/v126y2006i1p120-142.html*

Ottaviani, Marco and Peter Norman Sørensen. 2006*b*. "The strategy of professional forecasting." *Journal of Financial Economics* 81(2):441–466.
URL: *https://www.sciencedirect.com/science/article/pii/S0304405X06000286*

Qiu, Yajie, Bruno Deschamps and Xiaoquan Liu. 2024. "Uncertainty and macroeconomic forecasts: Evidence from survey data." *Journal of Economic Behavior & Organization* 224(C):463–480.
URL: *https://ideas.repec.org/a/eee/jeborg/v224y2024icp463-480.html*

Scharfstein, David S and Jeremy C Stein. 1990. "Herd Behavior and Investment." *American Economic Review* 80(3):465–479.
URL: *https://ideas.repec.org/a/aea/aecrev/v80y1990i3p465-79.html*

Trueman, Brett. 1994. "Analyst Forecasts and Herding Behavior." *The Review of Financial Studies* 7(1):97–124.
URL: *https://doi.org/10.1093/rfs/7.1.97*

Van Nieuwerburgh, Stijn and Laura Veldkamp. 2006. "Learning asymmetries in real business cycles." *Journal of Monetary Economics* 53(4):753–772.
URL: *https://ideas.repec.org/a/eee/moneco/v53y2006i4p753-772.html*

Veldkamp, Laura L. 2005. "Slow boom, sudden crash." *Journal of Economic Theory* 124(2):230–257. Learning and Bounded Rationality.
URL: *https://www.sciencedirect.com/science/article/pii/S002205310400064X*

A Model

A.1 Equilibrium Illustration

In Figure 11, we plot the equilibrium private signal weights for both types under different reputation concerns states. In an economy with no reputation concerns (blue line), the optimal forecast is just their posterior belief. The flat blue curves illustrate that the forecasters' optimal forecast does not depend on the external evaluator's conjecture. Taking the Bayesian economy as a benchmark, we then show how reputation concerns distort the optimal forecast towards the biased signal for both types. The orange curve represents the optimal private signal weight given the evaluator's conjecture. In the left panels, we can see that both types set a higher weight on the biased signal relative to the Bayesian benchmark. This comes at the expense of setting less weight on the unbiased private signal (right panels).

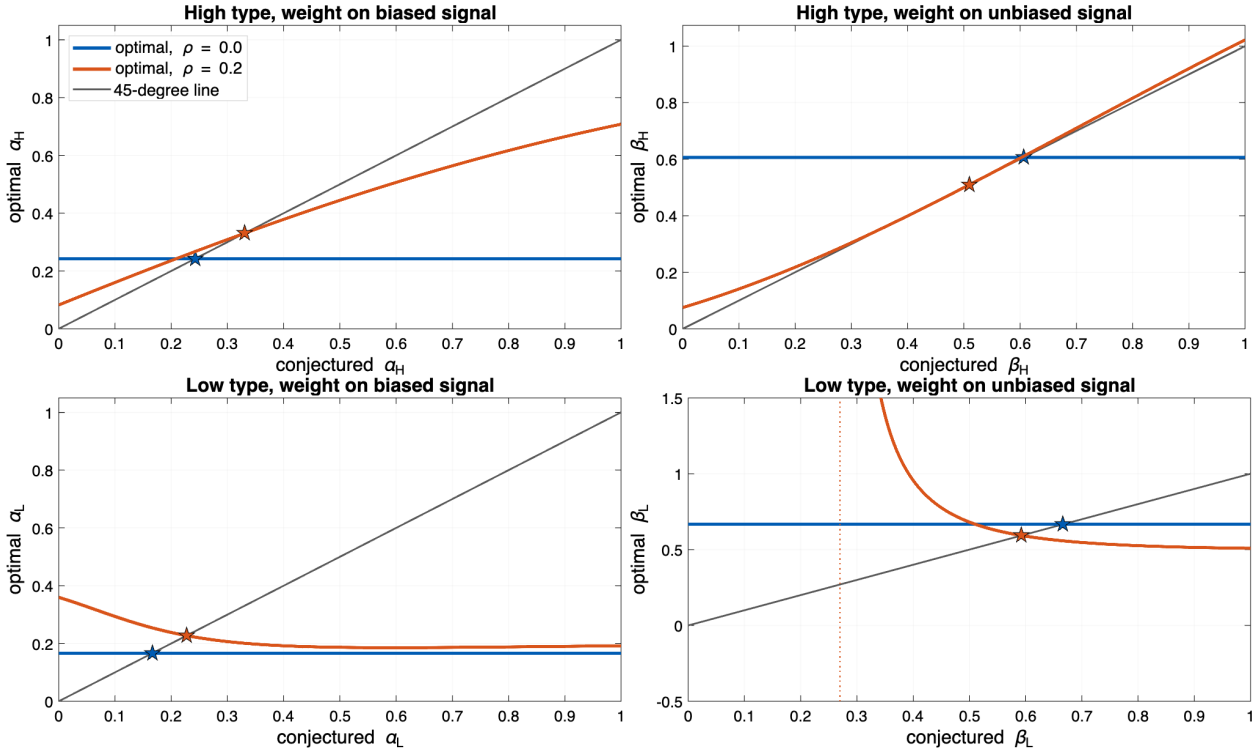


Figure 11: Reputation Concerns and Equilibrium Private Signal Weights.

Notes: The figure illustrates how reputation concerns affect equilibrium reporting strategies. The top (bottom) panels correspond to high-type (low-type) leaders, while the left (right) panel shows the weight assigned to the biased (unbiased) private signal. The horizontal axis represents the evaluator's conjecture weight, and the vertical axis represents forecasters' optimal weight. Blue lines depict optimal weights without reputation concerns ($\rho = 0$), while orange curves represent optimal responses when reputation concerns are present ($\rho = 0.2$). The 45-degree lines indicate equilibrium consistency, and stars mark equilibrium weights. Reputation concerns induce both forecaster types to place greater weight on the biased signal relative to their Bayesian benchmarks.

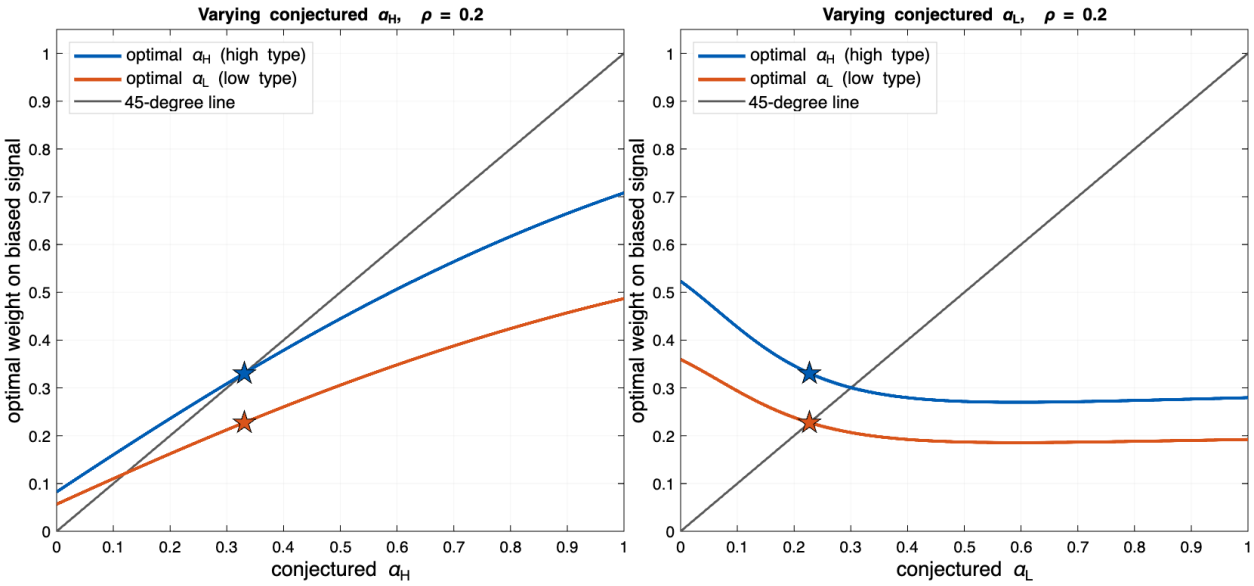


Figure 12: External Evaluator Beliefs and Optimal Leader Forecast

Notes: The figure illustrates how forecasters' optimal weights on the biased signal respond to changes in the external evaluator's conjectured reporting strategies when reputation concerns are present ($\rho = 0.2$). The left (right) panel varies the conjectured weight of high-type (low-type) leaders, holding the other type's conjectured weight fixed at the equilibrium value. Blue (orange) curves represent the optimal weights of the high-type (low-type) forecasters. The 45-degree line indicates equilibrium consistency for the corresponding type, and stars mark equilibrium weights. An increase in the conjectured high-type weight raises the optimal weights of both types, reflecting incentives to resemble high-type forecasters. Conversely, an increase in the conjectured low-type weight reduces the optimal weights of both types locally around the equilibrium, reflecting incentives to differentiate themselves from low-type forecasters.

The slopes of the orange curves illustrate the reputation chasing mechanism, which is better depicted in Figure 12. When the evaluator conjectures that high types place more weight in the private signals, forecasters optimally set more weight on the private signals. Meanwhile, the slope is locally negative around the equilibrium due to the exact same reputation channel. When the evaluator expects a low-type to set more weight on the private signals, forecasters optimally set a lower weight in these signals to avoid being perceived as a low-type. The differences in the extent of weight set by both types come from the accuracy cost in forecasters' preferences. It is more costly for low-type forecasters to set a higher weight on the biased signal, since it brings them farther from their posterior belief. This assumption not only prevents a cheap-talk game, but also the existence of a pooling equilibrium.

A.2 Proof of Proposition 1

Fix the uncertainty state U. At $\rho = 0$, both types report their Bayesian posterior beliefs, so that $\alpha_q = \alpha_q^B$ and $\beta_q = \beta_q^B$.

The best-response mapping is smooth around the Bayesian equilibrium, where $V_s^B > 0$ and $D = \gamma > 0$. Moreover, at $\rho = 0$, best responses do not depend on the evaluator's conjectured reporting coefficients. The implicit function theorem therefore guarantees a unique differentiable equilibrium branch locally around $\rho = 0$.

From Equations (17) and (19), we have

$$\mu^*_{s,q} - m^\theta_q = -(1 - \alpha_s - \beta_s)(m^\theta_q - \mu_0) + \alpha_s m^\xi_q.$$

Differentiating the optimal reporting rule with respect to ρ at the Bayesian equilibrium yields

$$\left.\frac{\partial a_{i,q}}{\partial \rho}\right|_{\rho=0} = \frac{1}{\gamma}\left[\left(\frac{1-\alpha_L^B-\beta_L^B}{V_L^B} - \frac{1-\alpha_H^B-\beta_H^B}{V_H^B}\right)(m^\theta_q - \mu_0) + \left(\frac{\alpha_H^B}{V_H^B} - \frac{\alpha_L^B}{V_L^B}\right) m^\xi_q\right],$$

where

$$V_s^B = (\alpha_s^B)^2 \sigma^2_{\eta s} + (\beta_s^B)^2 \sigma^2_u.$$

Recall from Equation (5) that

$$m^\xi_q = \frac{\sigma^2_\xi}{\sigma^2_\xi + \sigma^2_{\eta q}}\left[(1-\alpha_q^B)(y_i - \mu_0) - \beta_q^B (x_i - \mu_0)\right].$$

Substituting this expression and matching the coefficients on the two private signals gives

$$\left.\frac{\partial}{\partial \rho}\left(\frac{\alpha_q}{\beta_q}\right)\right|_{\rho=0^+} = \frac{\sigma^2_\xi}{\gamma \beta_q^B (\sigma^2_\xi + \sigma^2_{\eta q})}\left(\frac{\alpha_H^B}{V_H^B} - \frac{\alpha_L^B}{V_L^B}\right).$$

The first term is strictly positive under the assumptions of the proposition. To establish that the second term is also positive, notice that the Bayesian weights imply

$$\frac{\alpha_s^B}{V_s^B} = \frac{1 + (\tau_0 + \tau_u)(\sigma^2_\xi + \sigma^2_{\eta s})}{\sigma^2_{\eta s} + \tau_u (\sigma^2_\xi + \sigma^2_{\eta s})^2}.$$

This expression is strictly decreasing in $\sigma^2_{\eta s}$, as can be verified by differentiation. Since high-type

forecasters have lower processing noise, $\sigma^2_{\eta H} < \sigma^2_{\eta L}$, it follows that

$$\frac{\alpha^B_H}{V^B_H} > \frac{\alpha^B_L}{V^B_L}.$$

Therefore,

$$\frac{\partial}{\partial \rho}\left(\frac{\alpha_q}{\beta_q}\right)\bigg|_{\rho=0^+} > 0, \qquad q \in \{H, L\}.$$

Hence, an increase in reputation concerns from the Bayesian equilibrium raises the relative weight placed on the biased signal for both forecaster types.

A.3 Proof of Proposition 2

For any realization of the fundamental and private signals, the cross-sectional decomposition of squared forecasting errors implies:

$$\int_0^1 (a_i - \theta)^2 \, di = (\bar{a} - \theta)^2 + \int_0^1 (a_i - \bar{a})^2 \, di$$

Taking expectations conditional on (U, ρ) gives:

$$\mathbb{E}\left[\int_0^1 (a_i - \theta)^2 \, di \,\middle|\, U, \rho\right] = \mathcal{M}(U, \rho) + \mathcal{D}(U, \rho)$$

For each forecaster, let $m^{\theta}_{q_i,i} = \mathbb{E}[\theta \mid y_i, x_i, q_i]$ denote her Bayesian posterior belief. Since reported forecasts are measurable with respect to the forecaster's information set, the orthogonality property of conditional expectations implies:

$$\mathbb{E}[(a_i - \theta)^2 \mid U, \rho] = \mathbb{E}[(m^{\theta}_{q_i,i} - \theta)^2 \mid U] + \mathbb{E}[(a_i - m^{\theta}_{q_i,i})^2 \mid U, \rho]$$

Integrating across forecasters:

$$\mathcal{M}(U, \rho) + \mathcal{D}(U, \rho) = \mathcal{B}(U) + \mathcal{R}(U, \rho)$$

where $\mathcal{B}(U)$ denotes the average Bayesian posterior mean-squared error.

Since forecasters report their Bayesian posterior beliefs when $\rho = 0$, we have $\mathcal{R}(U, 0) = 0$, implying:

$$\mathcal{B}(U) = \mathcal{M}(U, 0) + \mathcal{D}(U, 0)$$

Subtracting the two expressions establishes:

$$\mathcal{M}(U, \rho) - \mathcal{M}(U, 0) = -\Big[\mathcal{D}(U, \rho) - \mathcal{D}(U, 0)\Big] + \mathcal{R}(U, \rho) \tag{36}$$

A.4 Proof of Lemma 1

Consider a leader who submits a forecast a_j. Followers observe a_j and use the affine reporting strategies

$$a_{i,q} = \mu_0 + \alpha_q(y_i - \mu_0) + \beta_q(x_i - \mu_0) + \delta_q(a_j - \mu_0),$$

where the coefficients depend on the follower's type but not on the leader's realized type.

Since there is a continuum of followers and the leader has negligible mass, idiosyncratic signal errors disappear through aggregation. The overall consensus is therefore given by

$$\bar{a} - \mu_0 = (\bar{\alpha} + \bar{\beta})(\theta - \mu_0) + \bar{\alpha}\xi + \bar{\delta}(a_j - \mu_0),$$

where $\bar{\alpha}$, $\bar{\beta}$, and $\bar{\delta}$ denote the population-average follower coefficients.

The evaluator observes $(\theta, a_j, \bar{a})$. Provided that $\bar{\alpha} \neq 0$, she can recover the common bias as

$$\widehat{\xi} = \frac{\bar{a} - \mu_0 - (\bar{\alpha} + \bar{\beta})(\theta - \mu_0) - \bar{\delta}(a_j - \mu_0)}{\bar{\alpha}}.$$

Importantly, $\widehat{\xi} = \xi$ for any report submitted by the leader, including deviations from her equilibrium strategy. Any change in the consensus induced by followers' responses to a_j is exactly offset in the evaluator's inference.

Under the hypothesis that the leader is type $s \in \{H, L\}$, her conjectured forecast distribution conditional on (θ, ξ) is normal, with mean

$$\mu_s^j = \mu_0 + (\alpha_s^j + \beta_s^j)(\theta - \mu_0) + \alpha_s^j \xi$$

and variance

$$V_s^j = (\alpha_s^j)^2 \sigma_{\eta s}^2 + (\beta_s^j)^2 \sigma_u^2 > 0.$$

Let $g_s(a_j \mid \theta, \xi)$ denote the corresponding conditional density. The joint density of (θ, ξ, a_j) under type s is

$$p_\theta(\theta) p_\xi(\xi) g_s(a_j \mid \theta, \xi).$$

The mapping from (θ, ξ, a_j) to $(\theta, a_j, \bar{a})$ is invertible, with Jacobian determinant $\bar{\alpha}$. By a change of variables, the evaluator's joint likelihood satisfies

$$f_s(\theta, a_j, \bar{a}) = \frac{p_\theta(\theta) p_\xi(\widehat{\xi}) g_s(a_j \mid \theta, \widehat{\xi})}{|\bar{\alpha}|}.$$

Since the prior densities and the Jacobian are identical under both type hypotheses, they cancel in the likelihood ratio:

$$\frac{f_H(\theta, a_j, \bar{a})}{f_L(\theta, a_j, \bar{a})} = \frac{g_H(a_j \mid \theta, \widehat{\xi})}{g_L(a_j \mid \theta, \widehat{\xi})}.$$

Because $\widehat{\xi} = \xi$ for every possible leader report, the evaluator updates the leader's reputation according to

$$\ell_j' = \ell_j + \log \frac{g_H(a_j \mid \theta, \xi)}{g_L(a_j \mid \theta, \xi)}.$$

Consequently, a type-p leader who anticipates followers' responses solves

$$\max_{a_j} E\left[\rho\left(\ell_j + \log \frac{g_H(a_j \mid \theta, \xi)}{g_L(a_j \mid \theta, \xi)} \right) - \frac{\gamma}{2}(a_j - \theta)^2 \,\middle|\, x_j, y_j, p \right].$$

This is exactly the optimization problem faced by a forecaster in the simultaneous economy. In particular, the objective does not depend on the followers' reporting coefficients. Moreover, since $V_s^j > 0$, both Gaussian densities have full support, so the likelihood ratio is well-defined for any leader report, including deviations from equilibrium.

Therefore, anticipating followers' responses does not alter the leader's optimal reporting problem. The leader's equilibrium reporting coefficients satisfy the same fixed-point conditions as in the baseline economy and can be determined independently of the followers' equilibrium.

A.5 Derivation of $\Omega_{q,p}$

Recall that $\Omega_{q,p} \equiv \text{Var}(e^{\perp}_{a,p})$.

$$
\begin{aligned}
\Omega_{q,p} &= \text{Var}(e_{a,p}) + \lambda^2_{q,p}\text{Var}(e_{y,q}) - 2\lambda_{q,p}\text{Cov}(e_{a,p}, e_{y,q}) \\
&= \text{Var}(e_{a,p}) + \frac{\text{Cov}(e_{a,p}, e_{y,q})^2}{\text{Var}(e_{y,q})^2}\text{Var}(e_{y,q}) - 2\frac{\text{Cov}(e_{a,p}, e_{y,q})}{\text{Var}(e_{y,q})}\text{Cov}(e_{a,p}, e_{y,q}) \\
&= \text{Var}(e_{a,p}) - \frac{\text{Cov}(e_{a,p}, e_{y,q})^2}{\text{Var}(e_{y,q})} \\
&= \text{Var}(e_{a,p}) - \frac{(\alpha^j_p)^2\sigma^4_\xi}{\sigma^2_\xi + \sigma^2_{\eta,q}} \\
&= (\alpha^j_p)^2(\sigma^2_\xi + \sigma^2_{\eta,p}) + (\beta^j_p)^2\sigma^2_u - \frac{(\alpha^j_p)^2\sigma^4_\xi}{\sigma^2_\xi + \sigma^2_{\eta,q}} \\
&= (\alpha^j_p)^2\left[\sigma^2_{\eta,p} + \frac{\sigma^2_\xi\sigma^2_{\eta,q}}{\sigma^2_\xi + \sigma^2_{\eta,q}}\right] + (\beta^j_p)^2\sigma^2_u
\end{aligned}
$$

A.6 Derivation of Bayesian Weights in Terms of $\bar{a}_j$

Recall the definition of $z_{j,qp}$ from Equation 27:

$$z_{j,qp} = \frac{Z_{q,p}}{h_{q,p}} + \mu_0$$

Using the definition of $Z_{q,p}$ from Equation 25:

$$z_{j,qp} - \mu_0 = \frac{(\bar{a}_j - \mu_0) - \lambda_{q,p}(y_i - \mu_0)}{h_{q,p}}$$

Plugging the previous equation into 28:

$$m^{\theta}_{q,p} = \mu_0 + \omega^{\alpha,\theta}_{q,p}(y_i - \mu_0) + \omega^{\beta,\theta}_{q,p}(x_i - \mu_0) + \omega^{\delta,\theta}_{q,p}\left(\frac{\bar{a}_j - \mu_0}{h_{q,p}} - \frac{\lambda_{q,p}(y_i - \mu_0)}{h_{q,p}}\right)$$

Matching the coefficients in the previous equations:

$$m^{\theta}_{q,p} = \mu_0 + \underbrace{\left(\omega^{\alpha,\theta}_{q,p} - \frac{\lambda_{q,p}}{h_{q,p}}\right)}_{\alpha^{\theta}_{q,p}}(y_i - \mu_0) + \underbrace{\omega^{\beta,\theta}_{q,p}}_{\beta^{\theta}_{q,p}}(x_i - \mu_0) + \underbrace{\left(\frac{\omega^{\delta,\theta}_{q,p}}{h_{q,p}}\right)}_{\delta^{\theta}_{q,p}}(\bar{a}_j - \mu_0)$$

A.7 Derivation of the Common-Bias Posterior Coefficients

Recall that the follower observes two signals with orthogonal noise components:

$$y_i - \theta = e_{y,q} = \xi + \eta_i, \tag{37}$$

$$z_{j,qp} - \theta = v_{j,qp} = \frac{e^{\perp}_{a,p}}{h_{q,p}}. \tag{38}$$

By construction, $\mathrm{Cov}(e_{y,q}, e^{\perp}_{a,p}) = 0$. Under Gaussianity, the two noise components are independent.

Define the Gaussian projection coefficients:

$$\kappa_{y,q} \equiv \frac{\mathrm{Cov}(\xi, e_{y,q})}{\mathrm{Var}(e_{y,q})} = \frac{\sigma^2_{\xi}}{\sigma^2_{\xi} + \sigma^2_{\eta q}}, \tag{39}$$

$$\kappa_{z,qp} \equiv \frac{\mathrm{Cov}(\xi, v_{j,qp})}{\mathrm{Var}(v_{j,qp})} = \frac{(\alpha^{j}_{p} - \lambda_{q,p})\sigma^2_{\xi} h_{q,p}}{\Omega_{q,p}}. \tag{40}$$

Conditional on the realization of the fundamental, Gaussian signal extraction gives:

$$\mathbb{E}[\xi \mid \theta, y_i, z_{j,qp}, q, p] = \kappa_{y,q}(y_i - \theta) + \kappa_{z,qp}(z_{j,qp} - \theta). \tag{41}$$

Since the follower does not observe θ, integrating over the posterior belief $m^{\theta}_{q,p}$ yields:

$$m^{\xi}_{q,p} = \kappa_{y,q}(y_i - m^{\theta}_{q,p}) + \kappa_{z,qp}(z_{j,qp} - m^{\theta}_{q,p}). \tag{42}$$

Substituting the affine fundamental posterior,

$$\begin{aligned} m^{\theta}_{q,p} = \mu_0 &+ \alpha^{\theta}_{q,p}(y_i - \mu_0) + \beta^{\theta}_{q,p}(x_i - \mu_0) \\ &+ \delta^{\theta}_{q,p}(z_{j,qp} - \mu_0), \end{aligned} \tag{43}$$

and collecting coefficients gives:

$$\begin{aligned}\alpha^{\xi}_{q,p} &= \kappa_{y,q} - (\kappa_{y,q} + \kappa_{z,qp})\alpha^{\theta}_{q,p},\\ \beta^{\xi}_{q,p} &= -(\kappa_{y,q} + \kappa_{z,qp})\beta^{\theta}_{q,p},\\ \delta^{\xi}_{q,p} &= \kappa_{z,qp} - (\kappa_{y,q} + \kappa_{z,qp})\delta^{\theta}_{q,p}.\end{aligned} \tag{44}$$

A.8 Follower Affine Bayesian Weights

Since 28 and 32 are affine, we can re-write 33 and 34 as affine expressions.

$$m^{\theta}_{q} = \mu_0 + \alpha^{\theta}_{q}(y_i - \mu_0) + \beta^{\theta}_{q}(x_i - \mu_0) + \bar{\varphi}_j\delta^{\theta}_{qH}(z_{j,qH} - \mu_0) + (1-\bar{\varphi}_j)\delta^{\theta}_{qL}(z_{j,qL} - \mu_0)$$

$$m^{\xi}_{q} = \alpha^{\xi}_{q}(y_i - \mu_0) + \beta^{\xi}_{q}(x_i - \mu_0) + \bar{\varphi}_j\delta^{\xi}_{q,H}(z_{j,qH} - \mu_0) + (1-\bar{\varphi}_j)\delta^{\xi}_{q,L}(z_{j,qL} - \mu_0)$$

where

$$\alpha^{\theta}_{q} = \bar{\varphi}_j\alpha^{\theta}_{q,H} + (1-\bar{\varphi}_j)\alpha^{\theta}_{q,L}, \quad \beta^{\theta}_{q} = \bar{\varphi}_j\beta^{\theta}_{q,H} + (1-\bar{\varphi}_j)\beta^{\theta}_{q,L}$$

$$\alpha^{\xi}_{q} = \bar{\varphi}_j\alpha^{\xi}_{q,H} + (1-\bar{\varphi}_j)\alpha^{\xi}_{q,L}, \quad \beta^{\xi}_{q} = \bar{\varphi}_j\beta^{\xi}_{q,H} + (1-\bar{\varphi}_j)\beta^{\xi}_{q,L}$$

A.9 External Evaluator and Followers' Problem

Evaluator Using the Average Forecast. The evaluator conjectures that a type-s forecaster follows an affine reporting strategy on the observed signals.

$$a_{i,s} = \mu_0 + \alpha_s(y_i - \mu_0) + \beta_s(x_i - \mu_0) + \delta_s(a_j - \mu_0) \tag{45}$$

where α_s, β_s, and δ_s are strategic equilibrium objects for $s \in \{H, L\}$.

Let χ denote the actual population share of high-type followers. Hence, we can write the average follower weights as:

$$\bar{\alpha} = \chi\alpha_H + (1-\chi)\alpha_L, \quad \bar{\beta} = \chi\beta_H + (1-\chi)\beta_L, \quad \bar{\delta} = \chi\delta_H + (1-\chi)\delta_L$$

Idiosyncratic errors from 45 average out:

$$\bar{a} - \mu_0 = (\bar{\alpha} + \bar{\beta})(\theta - \mu_0) + \bar{\alpha}\xi + \bar{\delta}(a_j - \mu_0) \tag{46}$$

Provided that $\bar{\alpha} \neq 0$, the evaluator can use the consensus $\bar{a}$ to invert 46 and recover the common bias.

$$\xi = \frac{\bar{a} - \mu_0 - (\bar{\alpha} + \bar{\beta})(\theta - \mu_0) - \bar{\delta}(a_j - \mu_0)}{\bar{\alpha}} \tag{47}$$

Evaluator Assessing the Followers. The evaluator updates the followers' reputations in the same fashion as it does with the forecasters in the baseline model. Conditional on the realized $(\theta, \xi, \bar{a}_j)$, the evaluator benchmarks a type-s forecast using:

$$\mu_s = (1 - \alpha_s - \beta_s - \delta_s)\mu_0 + (\alpha_s + \beta_s)\theta + \alpha_s\xi + \delta_s a_j \tag{48}$$

with conditional variance[18],

$$V_s = \alpha_s^2\sigma_{\eta s}^2 + \beta_s^2\sigma_u^2 \tag{49}$$

The external evaluator then updates forecaster i's reputation through Bayes rule.

$$\ell_i' = \ell_i + \frac{1}{2}\log\left(\frac{V_L}{V_H}\right) - \frac{(a_i - \mu_H)^2}{2V_H} + \frac{(a_i - \mu_L)^2}{2V_L} \tag{50}$$

The Follower's Problem. Each follower chooses a forecast a_i to maximize the following problem:

$$\max_{a_i} \mathbb{E}\left[\rho\ell_i' - \frac{\gamma}{2}(a_i - \theta)^2 \mid x_i, y_i, a_j, q_i = q, \varphi_j\right] \tag{51}$$

The followers know that the evaluator updates future reputation using 50. Similar to the forecasters' problem in Section 2, the followers must form expectations about the external evaluator's benchmark. Let $\mu_{s,q}^* \equiv \mathbb{E}[\mu_s \mid x_i, y_i, q_i = q, a_j, \varphi_j]$ be the type-q follower's expectation of how the the evaluator benchmarks

[18] Notice that a_j affects the conditional mean but not the conditional variance because the leader's forecast is publicly observed.

a type-s forecaster. Using the followers' posterior beliefs m_q^θ and m_q^ξ:

$$\mu_{s,q}^* = (1 - \alpha_s - \beta_s - \delta_s)\mu_0 + (\alpha_s + \beta_s)m_q^\theta + \alpha_s m_q^\xi + \delta_s a_j \tag{52}$$

Followers' Optimal Forecast. The first-order condition for a type-q follower solving 51 is:

$$\gamma(a_i - m_q^\theta) = \rho\left[-\frac{a_i - \mu_{Hq}^*}{V_H} + \frac{a_i - \mu_{Lq}^*}{V_L}\right] \tag{53}$$

Defining $D \equiv \gamma + \rho\left(\frac{1}{V_H} - \frac{1}{V_L}\right)$ and re-arranging 53, we can write the optimal forecast for a type-q follower as $a_{i,q}^*$:

$$a_{i,q}^* = m_q^\theta + \frac{\rho}{D}\left[\frac{\mu_{Hq}^* - m_q^\theta}{V_H} - \frac{\mu_{Lq}^* - m_q^\theta}{V_L}\right] \tag{54}$$

B Solution Method

In this appendix, we describe the numerical procedure used to solve the affine rational expectations equilibria characterized in Sections 2 and 3. In both environments, an equilibrium is a fixed point in the reporting coefficients, such that the evaluator's conjectured strategies coincide with forecasters' optimal reporting strategies.

B.1 Recursive Equilibrium

As established in Lemma 1 (Appendix A.4), the leader's optimal reporting problem is independent of the followers' equilibrium strategies. Consequently, we can solve the sequential economy recursively. We first determine the leader's reporting coefficients, which solve the same problem as in the simultaneous economy, and then solve the followers' reporting coefficients conditional on the leader's equilibrium strategy.

Leader's Problem. For given public states (U, ρ), the leader's equilibrium is characterized by four coefficients: $(\alpha_H^j, \beta_H^j, \alpha_L^j, \beta_L^j)$.

Starting from conjectured reporting coefficients, we compute the evaluator's conditional means and variances for each type, as defined in Equations 13 and 14. Given these conjectures, we obtain the optimal

reporting strategy from Equation 19. Since the optimal forecast is affine in the private signals, it can be written as

$$a^*_{j,p} = \mu_0 + \widehat{\alpha}^j_p(y_j - \mu_0) + \widehat{\beta}^j_p(x_j - \mu_0),$$

where $p \in \{H, L\}$. The optimal coefficients are obtained by matching the weights on the two private signals in Equation 19.

An equilibrium requires that the optimal coefficients coincide with the evaluator's conjectures:

$$\widehat{\alpha}^j_p = \alpha^j_p, \qquad \widehat{\beta}^j_p = \beta^j_p, \qquad p \in \{H, L\}.$$

We solve these four nonlinear equations jointly. The second-order condition is

$$\gamma + \rho\left(\frac{1}{V^j_H} - \frac{1}{V^j_L}\right) > 0,$$

which ensures that the optimal forecast maximizes the forecaster's expected objective. The resulting coefficients also characterize the equilibrium of the simultaneous economy in Section 2.

Followers' Problem. Given the leader's equilibrium strategy and public reputation φ_j, we solve for the six follower reporting coefficients $(\alpha_H, \beta_H, \delta_H, \alpha_L, \beta_L, \delta_L)$.

For a conjectured set of follower coefficients, we first compute their Bayesian posterior beliefs about the fundamental and common bias, conditional on each possible leader type. These beliefs are characterized in Equations 28 and 31, and combined using the reputation-weighted inference rules in Equations 33 and 34.

Conditional on the leader's type p, we express the followers' problem using the orthogonalized leader signal $z_{j,qp}$, defined in Equation 27. This transformation removes the correlation between the noise in the leader's forecast and the noise in the follower's biased private signal. It allows us to calculate the Bayesian posterior weights using the precisions of conditionally independent signals.

We then use Equation 54 to compute the optimal reporting coefficients in this transformed signal space. Conditional on leader type p, the corresponding optimal forecast can be written as

$$a^*_{i,q,p} = \mu_0 + w^y_{q,p}(y_i - \mu_0) + w^x_{q,p}(x_i - \mu_0) + w^j_{q,p}(z_{j,qp} - \mu_0).$$

To recover the coefficients on the signals actually observed by followers, we use the transformation implied by Equations 25–27:

$$z_{j,qp} - \mu_0 = \frac{(a_j - \mu_0) - \lambda_{q,p}(y_i - \mu_0)}{h_{q,p}},$$

where $\lambda_{q,p}$ and $h_{q,p}$ are defined in Equations 24 and 26, respectively. The optimal coefficients in the original reporting strategy are therefore

$$\widehat{\alpha}_{q,p} = w^y_{q,p} - \frac{\lambda_{q,p}}{h_{q,p}} w^j_{q,p},$$

$$\widehat{\beta}_{q,p} = w^x_{q,p},$$

$$\widehat{\delta}_{q,p} = \frac{w^j_{q,p}}{h_{q,p}}.$$

Since followers do not observe the leader's type, we combine these optimal coefficients using the publicly observed leader reputation, consistently with the inference rules in Equations 33 and 34. Thus,

$$\widehat{\alpha}_q = \varphi_j \widehat{\alpha}_{q,H} + (1 - \varphi_j)\widehat{\alpha}_{q,L},$$

and analogously for $\widehat{\beta}_q$ and $\widehat{\delta}_q$. This aggregation is valid because the optimal reporting rule is linear in posterior beliefs and the evaluator's conditional variances, defined in Equation 49, do not depend on the leader's type.

Finally, the follower equilibrium requires

$$\widehat{\alpha}_q = \alpha_q, \qquad \widehat{\beta}_q = \beta_q, \qquad \widehat{\delta}_q = \delta_q, \qquad q \in \{H, L\}.$$

We solve these six nonlinear equations jointly, conditional on the leader's equilibrium coefficients. As in the leader's problem, we verify that the second-order condition is satisfied.

B.2 Numerical Solution and Equilibrium Selection

Bayesian Equilibrium. We initialize the numerical solution at the economy without reputation concerns, $\rho = 0$. In this case, the reputational terms in Equations 19 and 54 disappear, and all forecasters optimally

report their Bayesian posterior beliefs.

The equilibrium is therefore known analytically. The leader's reporting coefficients are the Bayesian weights in Equation 4, while the followers' coefficients are obtained from Equations 29 and 30, combined using the reputation-weighted inference rule in Equation 33 and evaluated at the Bayesian leader's strategy.

Importantly, at $\rho = 0$, optimal reporting coefficients do not depend on the evaluator's conjectures. Hence, this equilibrium provides a unique starting point for the numerical solution.

Continuation in Reputation Concerns. Starting from the Bayesian equilibrium, we increase ρ gradually and solve the fixed-point problem at each new value. We use the previous equilibrium coefficients to construct the initial guess for the next solution, with secant extrapolation when two previous solutions are available. At each value of ρ, we first solve the leader's four-equation system. Given the resulting coefficients, we update followers' posterior beliefs and solve their six-equation system, separately for each leader-reputation environment.

The nonlinear systems are solved using a damped Gauss–Newton method, with numerical derivatives and a line search. We use a grid spacing of $\Delta\rho = 0.001$ and reduce the step size when the solver fails to converge to a sufficiently close solution. A solution is accepted only if the maximum absolute fixed-point residual is below 10^{-10}, the evaluator's conditional variances are strictly positive, and the second-order conditions are satisfied. We also verify that successive solutions remain close to the extrapolated equilibrium path.

This continuation procedure selects the equilibrium branch connected to the Bayesian equilibrium at $\rho = 0$. All numerical results in the paper are computed along this branch. The procedure does not establish that other equilibria do not exist for $\rho > 0$.

Numerical Illustration. For the numerical results in Sections 2 and 3, we use the following illustrative parameter values:

$$\gamma = 1, \qquad v(U) = 2, \qquad \sigma_\xi^2 = 1, \qquad \sigma_u^2 = 0.5,$$

$$\sigma_{\eta H}^2 = 0.25, \qquad \sigma_{\eta L}^2 = 1, \qquad \chi = 0.5.$$

C Empirical Facts

C.1 Table Supporting Empirical Fact 3

Table 5: Forecast Accuracy and Deviations from the Empirical Leader Consensus

	$\lvert\hat{\pi}_{i,t} - \pi_t\rvert$
IRP_{t-1}	−0.003
	(0.015)
$d_{i,t}$	0.887^{***}
	(0.090)
$IRP_{t-1} \times d_{i,t}$	-0.947^{***}
	(0.204)
$n_{i,t}^{\text{prior}}$	−0.0001
	(0.0001)
Forecaster fixed effects	Yes
Observations	17,978
Forecasters	386
Within R^2	0.139

Notes: OLS regression of the absolute forecast error, $|\hat{\pi}_{i,t} - \pi_t|$, on the inflation risk premium IRP_{t-1}, the distance to the empirical leader consensus, $d_{i,t} = |\hat{\pi}_{i,t} - \tilde{\pi}_{i,t}|$, their interaction, and the number of forecasts submitted before forecaster i, $n_{i,t}^{\text{prior}}$, with forecaster fixed effects. Errors, distances and IRP are in percentage points. The empirical leader consensus $\tilde{\pi}_{i,t}$ is the median of the forecasts submitted on dates strictly before forecaster i's, and is defined when there are at least three. IRP_{t-1} is the Cleveland Fed ten-year inflation risk premium of the month preceding the announcement. Standard errors clustered by forecaster in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.