Analyst Forecast Biases and the Cost of Wishful Thinking

Snehal Banerjee, Naveen Gondhi, and Xinran Zhao*

August 17, 2026

Abstract

We show that EPS forecasts exhibit underreaction and CAPEX forecasts exhibit overreaction, both at the individual and consensus levels, for the *same analysts and firms*. This evidence is difficult to reconcile in standard models. We propose a wishful thinking model of forecasts in which analysts optimally distort beliefs about signal precision to gain anticipatory utility. A single parameter, the cost of belief distortion, generates both biases: high costs produce underreaction, low costs produce overreaction, and intermediate costs yield individual overreaction but consensus underreaction. The model further predicts that overreaction is stronger at longer horizons and among more experienced analysts, and that low-cost environments feature greater forecast dispersion and consensus volatility. We find evidence consistent with these predictions.

JEL Classification: D62, D82, D83, D84, G14, G24
Keywords: anticipatory utility, wishful thinking, beauty contests, belief heterogeneity, disagreement.

*Banerjee (snehalb@umich.edu) is at the University of Michigan – Ann Arbor. Gondhi (naveen.gondhi@insead.edu) is at INSEAD. Zhao (xinran.zhao@insead.edu) is at INSEAD. We thank Joel Peress and the seminar participants at INSEAD for valuable feedback. Corresponding author: Naveen Gondhi, INSEAD Asia Campus, 1 Ayer Rajah Ave, Singapore 138676. E-mail: naveen.gondhi@insead.edu.

1 Introduction

Analysts play a key role in markets by shaping investor expectations and asset prices through their forecasts. A growing literature has shown that these forecasts are systematically biased and that the nature of the bias varies across settings: while some forecasts exhibit overreaction, other forecasts exhibit underreaction. What drives this variation — and whether a single coherent framework can account for it — remains an open question.

We contribute new evidence that sharpens this puzzle. Using regressions of forecast errors on forecast revisions (as in Coibion and Gorodnichenko (2012)) on I/B/E/S data, we document a striking set of facts: earnings (EPS) forecasts exhibit *underreaction* but capital expenditures (CAPEX) forecasts exhibit *overreaction*. Importantly, this holds at both the individual and consensus levels, and for the same analysts at one-year horizon. Moreover, this predictability varies systematically: overreaction is stronger at longer horizons, in qualitative environments, and among more experienced analysts. This evidence is difficult to reconcile within existing models of belief formation, as we discuss below.

We propose a parsimonious model of belief formation that reconciles not only this new evidence, but also existing stylized facts from the literature. We assume that analysts have incentives to produce accurate forecasts that do not "stand out" too much from the consensus forecast.[1] Importantly, analysts engage in "wishful thinking": they choose subjective beliefs about the precision of their private and public information when forming forecasts, trading off the benefit from anticipatory utility against the cost due to lower accuracy.[2]

We show that variation along a single dimension — the (relative) **costs of belief distortion** — can generate both underreaction and overreaction in forecasts. Specifically, when costs are sufficiently high (e.g., for EPS forecasts), analysts' incentives to herd dominate, leading to underreaction in individual and consensus forecasts. For intermediate costs, individual forecasts exhibit overreaction while consensus forecasts exhibit underreaction.[3] However, when the costs are sufficiently low (e.g., for CAPEX forecasts), analysts significantly overestimate the precision of their information, leading to overreaction in individual and consensus forecasts. Additionally, our model predicts that lower costs of belief distortion lead to greater cross-sectional dispersion in forecasts and higher volatility in consensus forecasts. We find evidence consistent with these predictions.

[1] This is consistent with strategic incentives of forecasters to herd with others — see Croushore (1997), Ottaviani and Sørensen (2006), Marinovic, Ottaviani, and Sørensen (2010), Gemmi and Valchev (2026), and Marinovic, Ottaviani, and Sorensen (2013).

[2] See Bénabou and Tirole (2016) for a recent survey on this topic.

[3] This is consistent with the evidence on macroeconomic forecasts documented by Coibion and Gorodnichenko (2012), Coibion and Gorodnichenko (2015), Bordalo, Gennaioli, Ma, and Shleifer (2020) and others.

Stylized Facts. Using individual analyst forecasts from I/B/E/S over 2007–2025, we estimate Coibion–Gorodnichenko (CG) coefficients of individual and consensus forecasts for EPS and CAPEX as measures of underreaction / overreaction.[4] We document the following stylized facts:

1. EPS forecasts exhibit underreaction and CAPEX forecasts exhibit overreaction at both individual and consensus levels, for the same firms and time periods. This is in contrast to the evidence from macroeconomic forecasts documented in Bordalo et al. (2020), which exhibit overreaction at the individual level but underreaction at the consensus level.

2. The CG coefficient declines with forecast horizon for both EPS and CAPEX, i.e., longer-horizon forecasts exhibit stronger overreaction.

3. The CG coefficient declines with analyst firm-following tenure: EPS underreaction is concentrated among less experienced analysts (higher, more positive CG coefficients), whereas CAPEX overreaction is stronger among more experienced analysts (more negative CG coefficients).

Taken together, these facts are difficult to reconcile within existing frameworks. Standard rational-expectations models with dispersed information (e.g., Coibion and Gorodnichenko, 2015) can generate underreaction at the consensus level but cannot account for overreaction at the individual level. Models in which analysts face strategic incentives to herd (e.g., Gemmi and Valchev, 2026) can generate underreaction at the individual and aggregate levels, but cannot generate overreaction at the consensus level. Recent behavioral models typically focus on simultaneously generating overreaction in individual forecasts but underreaction in consensus forecasts.[5] However, these approaches cannot naturally explain the coexistence of overreaction and underreaction of forecasts by the same analysts about different dimensions of the same firms, or the systematic variation in the biases across different contexts.

Model and Intuition. To reconcile the above empirical evidence, we propose a model of endogenous belief formation in which analysts engage in *wishful thinking* when forming

[4]The individual-level Coibion–Gorodnichenko (CG_i) coefficient is the slope from regressing individual forecast errors on individual forecast revisions. The consensus-level CG coefficient (CG_a) is defined analogously using consensus forecast errors and revisions. Positive values indicate underreaction, while negative values indicate overreaction.

[5]These include models of diagnostic expectations (e.g., Bordalo et al., 2020), extrapolation (e.g., Barberis, Greenwood, Jin, and Shleifer, 2015), overconfidence (e.g., Daniel, Hirshleifer, and Subrahmanyam, 1998), optimally noisy memory (e.g., da Silveira and Woodford, 2020, Sung, 2024), and costly information retrieval (e.g., Afrouzi, Kwon, Landier, Ma, and Thesmar, 2023).

forecasts.[6] To fix ideas, assume that analyst i's payoff u_i is given by:

$$u_i = -(k_i - \theta)^2 - \rho(k_i - K)^2,$$

where θ denotes the target variable, k_i denotes her forecast, and $K \equiv \int k_i di$ denotes the aggregate or consensus forecast. Specifically, the forecaster's payoff depends on (i) the accuracy of her forecast, and (ii) a preference to either deviate from ($\rho < 0$) or conform to ($\rho > 0$) the consensus forecast. Following the empirical literature on strategic incentives, we assume that $\rho > 0$, i.e., forecasters have an incentive to herd due to reputational concerns (see also Ottaviani and Sørensen (2006)).[7] Each analyst optimally chooses her forecast to maximize the above payoff, given noisy public and private information about the target.

When analysts have rational expectations, both individual and consensus forecasts exhibit underreaction. Because analysts have an incentive to herd, they underweight their private information and overweight public information when forming their forecasts. This distortion reflects the coordination motive that arises due to the strategic complementarity in actions (see Morris and Shin (2002)), and leads to underreaction in forecasts (as in Gemmi and Valchev (2026) and Banerjee et al. (2026)).

Next, we consider the case when each analyst faces uncertainty about the precision of her information and engages in "wishful thinking." Specifically, each analyst chooses her subjective beliefs about signal precision to maximize anticipatory utility subject to a **cost of belief distortion** that results from making suboptimal forecasts. Since anticipatory utility increases in the perceived accuracy of one's own forecast and the degree of conformity with the consensus, wishful thinking leads analysts to overestimate the precision of both their private and public information. However, the extent to which they do so depends on the cost of belief distortion.

When this cost is high, beliefs are disciplined and forecasts are close to those under the rational expectations benchmark. As a result, individual and consensus forecasts exhibit underreaction. As the cost of belief distortion decreases, analysts' subjective precision of their information increases, which makes overreaction in forecasts more likely. For intermediate costs, we show that this results in overreaction in individual forecasts, but underreaction in consensus forecasts.[8] Finally, for sufficiently low costs, both individual and consensus forecasts exhibit overreaction.

[6] As we discuss below, we build on the work of Banerjee, Davis, and Gondhi (2026), Banerjee, Davis, and Gondhi (2024), Caplin and Leahy (2019), and Bénabou and Tirole (2016).

[7] To ensure that the equilibrium is well defined, we need to further assume that $\rho \in (-1, 1)$.

[8] Notably, this is consistent with the evidence on macroeconomic forecasts documented by Bordalo et al. (2020).

Implications for EPS and CAPEX forecasts. The contrast between earnings (EPS) and capital expenditures (CAPEX) forecasts provides a sharp illustration of the key mechanism we propose. Belief distortion costs are very high for EPS forecasts, because earnings are closely monitored, quickly realized, and central to analyst evaluation and career concerns. In contrast, CAPEX outcomes are often realized with substantial delay, are harder to verify, and play a less prominent role in performance evaluation for analysts, implying low costs of belief distortion. Given these differences across targets, our model naturally predicts that EPS forecasts exhibit underreaction while CAPEX forecasts exhibit overreaction (at both the individual and consensus levels).

More generally, our model implies that both individual and consensus CG coefficients decrease (i.e., overreaction increases) as the cost of belief distortion decreases, since this leads analysts to further overestimate the precision of their information. This single mechanism provides a unified interpretation of all three stylized facts. Settings with delayed feedback and weaker ex post verification, such as longer forecast horizons, are likely to be associated with lower distortion costs and therefore generate lower CG coefficients. Similarly, qualitative information environments, where signals are harder to verify and performance is less easily evaluated, weaken discipline and lead to stronger overreaction relative to quantitative settings. Finally, analyst characteristics map naturally into this framework. Junior analysts face greater scrutiny and tighter evaluation, implying higher belief distortion costs, which lead to more underreaction, while senior analysts have greater discretion, lower distortion costs, and consequently exhibit weaker underreaction or stronger overreaction. Together, these patterns suggest that variation in expectation biases across forecast targets, horizons, information environments, and analyst characteristics reflects systematic differences in the discipline imposed on belief formation.

Our model generates additional predictions about how the cross-sectional dispersion of forecasts and the time-series volatility of consensus forecasts change with the cost of belief distortion. Since analysts overestimate the precision of their information more when costs are low, this leads to higher dispersion in forecasts (due to over-weighting of private information) and higher volatility of consensus forecasts (due to over-weighting of public news). As the cost of belief distortion increases, analysts' beliefs move closer to the rational benchmark, reducing the weight placed on both private and public information; consequently, forecast dispersion declines and consensus volatility declines.

We take these predictions to the data by exploiting variation in empirical proxies for the cost of belief distortion. Specifically, we use forecast horizon, the qualitative versus quantitative nature of information, and analyst experience as proxies for differences in monitoring, feedback, and accountability. Consistent with the model, we find that settings

associated with lower distortion costs exhibit stronger overreaction, higher cross-sectional dispersion in forecasts, and higher time-series volatility of consensus forecasts, while settings with higher distortion costs exhibit weaker overreaction or underreaction, lower dispersion, and lower consensus volatility.

Overview. The next section discusses the related literature. Section 3 presents the stylized facts we document about EPS and CAPEX forecasts. Section 4 presents the model and derives the main result, characterizing how CG coefficients vary with the cost of belief distortion. Section 5 discusses the additional predictions of the model, and Section 6 concludes. All proofs are in the Appendix.

2 Related Literature

Our paper contributes to three strands of the literature.

Empirical evidence on forecast biases. Building on Coibion and Gorodnichenko (2012, 2015), a growing literature documents systematic predictability in forecast errors. A key finding is that individual forecasts can overreact while consensus forecasts underreact (Bordalo et al., 2020; Afrouzi et al., 2023). In corporate finance, short-horizon earnings forecasts tend to underreact while longer-horizon forecasts overreact (Bouchaud, Krueger, Landier, and Thesmar, 2019; Bordalo, Gennaioli, Porta, and Shleifer, 2019). We add new evidence that sharpens this puzzle: for the *same* analysts, firms, and time periods, EPS forecasts exhibit underreaction while CAPEX forecasts exhibit overreaction, at both the individual and consensus levels. Moreover, these biases vary predictably with forecast horizon, analyst experience, and the qualitative nature of information (Ke, 2024). This coexistence is difficult to reconcile with existing frameworks. Rational expectations models generate consensus underreaction but not individual overreaction (Coibion and Gorodnichenko, 2015). Strategic herding models (Ottaviani and Sørensen, 2006; Gemmi and Valchev, 2026) generate underreaction at both levels but not consensus overreaction. Behavioral models (Bordalo et al., 2020; Bouchaud et al., 2019) generate individual overreaction alongside consensus underreaction, but cannot explain why the same analyst simultaneously underreacts on EPS and overreacts on CAPEX for the same firm at the same time.

Motivated beliefs and wishful thinking. The literature on anticipatory utility assumes that agents choose beliefs to balance accuracy against affective value (e.g., Akerlof and Dickens, 1982; Loewenstein, 1987; Brunnermeier and Parker, 2005; Caplin and Leahy, 2019; see also the survey by Bénabou and Tirole, 2016). Cassella, Golez, Gulen, and Kelly (2023) develops a multi-period model of motivated reasoning to show that optimism in

forecasts can increase with the forecast horizon. The paper finds empirical evidence of this using macroeconomic forecasts, and documents that the optimism bias varies with macroeconomic uncertainty and leads to stronger reaction to good news (vs. bad news).

Our framework builds directly on Banerjee et al. (2026): agents hold a single subjective model and choose the perceived precision of their signals subject to a penalty for departing from accuracy, interpreted as a reduced-form cost of belief distortion rather than knowledge of the true distribution. Relative to Brunnermeier and Parker (2005); Brunnermeier, Gollier, and Parker (2007), which also feature endogenous belief choice with experienced-utility costs, our key contribution is to embed these forces in a coordination game with strategic complementarities. The strategic considerations are important: with herding incentives, both the fundamental uncertainty channel and the aggregate-errors channel push analysts to overestimate signal precision. A single parameter, the cost of belief distortion, then governs whether forecasts overreact, underreact, or display the intermediate pattern (individual overreaction, consensus underreaction) documented for macroeconomic forecasts (Coibion and Gorodnichenko, 2015; Bordalo et al., 2020).

Subjective belief formation. Our paper also contributes to the recent literature on subjective belief formation. Diagnostic expectations (Bordalo, Gennaioli, and Shleifer, 2018; Bordalo et al., 2019, 2020) generate overreaction by making agents overweight representative signals. Memory-based models generate overreaction through distinct channels: fading memory and recency (Nagel and Xu, 2022), capacity-constrained recall (da Silveira and Woodford, 2020, 2022), and associative retrieval of salient past experiences (Enke, Graeber, and Oprea, 2020; Bordalo, Gennaioli, and Shleifer, 2021). Both approaches offer coherent accounts of settings where forecasts overreact. Our framework differs in two respects. First, belief distortions are the *outcome* of an optimization problem rather than a feature of the updating rule. Second, neither class of models naturally generates the coexistence of overreaction and underreaction across forecast targets within the same firms and time periods: both primarily push toward overreaction. In our framework, this coexistence follows directly from variation in the cost of belief distortion across targets (high for EPS and low for CAPEX) without requiring differences in information or analyst composition. We further test the key comparative static of Afrouzi et al. (2023) — that overreaction should increase with process transitoriness — and find that it is inconsistent with the empirical evidence on CAPEX forecasts (see Figure VII).

3 Stylized Facts

3.1 Data and Variable Construction

We obtain individual analyst forecasts of earnings per share (EPS) and capital expenditures (CAPEX) from the I/B/E/S Detail History File.[9] Our sample covers fiscal years 2007 to 2025, with the start date determined by the availability of CAPEX forecast coverage in I/B/E/S. We focus on annual forecasts for one-year-ahead to five-year-ahead horizons, where we identify each horizon using the forecast issuance date and the target fiscal period end date.

We retain only the first forecast issued by each analyst after the most recent annual earnings announcement and before the next quarterly earnings report date. This window is designed so that (i) the analyst has incorporated the latest annual earnings news, and (ii) the forecast is issued before the next quarterly report, limiting the scope for quarterly information to affect the forecast. By retaining only the first forecast in this window, the revision between the two-year-ahead and one-year-ahead forecasts more closely captures the analyst's response to the newly released annual outcomes.[10] To ensure that differences in CG coefficients across items reflect differences in how analysts process information rather than differences in analyst composition or information sets, we retain only analysts who issue EPS and CAPEX forecasts on the same date with both one-year-ahead and two-year-ahead horizons for both items.

Figure I illustrates the timing of forecasts, revisions, and realizations. We use $x_{j,t}$ to denote the target variable being forecasted (EPS or CAPEX for firm j in fiscal year t), and $y_{j,t}$ to denote the realized outcomes reported at the fiscal year t earnings announcement. As shown in the figure, analyst i observes $y_{j,t-1}$ (the realized outcomes for fiscal year $t-1$) before issuing the two-year-ahead forecast $F_{i,t-1}x_{j,t+1}$, and subsequently observes $y_{j,t}$ before issuing the one-year-ahead forecast $F_{i,t}x_{j,t+1}$. For example, after the fiscal year 2014 earnings announcement, analyst i issues forecasts for firm j's EPS and CAPEX for fiscal year 2016 on 17th March 2015. We denote these two-year-ahead forecasts as $F_{i,t-1}x_{j,t+1}$. One year later, after observing the firm's realized outcomes for fiscal year 2015, the analyst revises the forecasts for fiscal year 2016 on 22nd March 2016. We denote these one-year-ahead forecasts

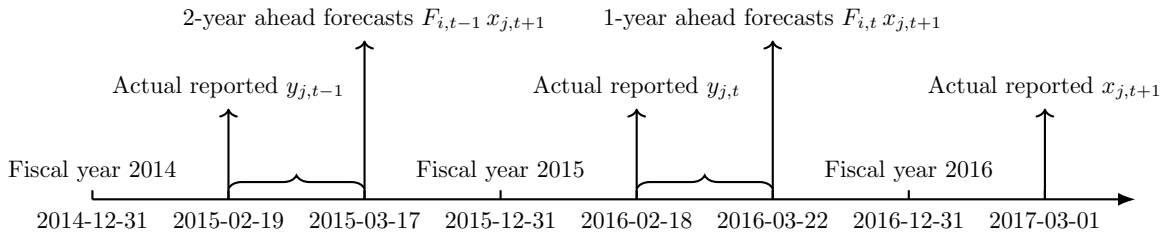
[9] We use the I/B/E/S unadjusted detail history file for EPS forecasts and the unadjusted actuals file for realizations. As documented by Diether, Malloy, and Scherbina (2002), stock splits occurring between the forecast date and the earnings announcement can cause mechanical differences between forecasts and actuals. We adjust all EPS forecasts using cumulative share adjustment factors from CRSP to ensure that forecasts and realizations are on the same per-share basis.

[10] Table I reports the distribution of the number of days between the forecast issuance date and the preceding earnings announcement date. The majority of analysts revise their forecasts on the day following the earnings announcement.

as $F_{i,t}x_{j,t+1}$.

Figure I: Timing of Forecasts, Revisions, and Realizations

This figure illustrates the timing of analyst forecasts and the corresponding realizations for firm j. Here $x_{j,t+1}$ denotes the target variable being forecasted (EPS or CAPEX) for fiscal year $t+1$, and $y_{j,t}$ denotes the realized outcomes reported at the fiscal year t earnings announcement. At each date t, analysts issue one-year-ahead forecasts $F_{i,t}x_{j,t+1}$ after observing the reported outcomes $y_{j,t}$. The two-year-ahead forecasts $F_{i,t-1}x_{j,t+1}$ are formed one year earlier, after observing $y_{j,t-1}$. Realizations of the target variable are reported with a delay relative to the fiscal year end, generating overlapping forecast horizons used in the Coibion–Gorodnichenko regressions.



The forecast revision is defined as the difference between the one-year-ahead and two-year-ahead forecasts, i.e., $F_{i,t}x_{j,t+1} - F_{i,t-1}x_{j,t+1}$ at the individual level and $F_t x_{j,t+1} - F_{t-1}x_{j,t+1}$ at the consensus level. The forecast error is defined as the realized value minus the expectation, i.e., $x_{j,t+1} - F_{i,t}x_{j,t+1}$ at the individual level and $x_{j,t+1} - F_t x_{j,t+1}$ at the consensus level. The consensus forecast $F_t\, x_{j,t+1}$ is the mean forecast across all analysts covering firm j in year t.[11]

Following Bouchaud et al. (2019), we normalize both forecast errors and revisions by a component of $y_{j,t-1}$, the realized outcomes at the end of fiscal year $t-1$. Because $y_{j,t-1}$ is determined before either forecast is issued, the scaling variable does not introduce look-ahead bias. The scaling variable is the stock price for EPS, and net property, plant, and equipment (PP&E) for CAPEX. We estimate the following Coibion–Gorodnichenko regressions at the consensus level (Equation (1)) and individual level (Equation (2)):

$$\frac{FE_{j,t}}{y_{j,t-1}} = \alpha^C + \beta^C \times \frac{REV_{j,t}}{y_{j,t-1}} + u_{j,t}. \tag{1}$$

$$\frac{FE_{i,j,t}}{y_{j,t-1}} = \alpha^I + \beta^I \times \frac{REV_{i,j,t}}{y_{j,t-1}} + u_{i,j,t}. \tag{2}$$

To reduce the impact of outliers, we remove observations for which any variable deviates from its median by more than five times the interquartile range, following Bouchaud et al. (2019) and Bordalo et al. (2020). The final sample contains 16,687 individual analyst forecasts at the analyst-firm-year level and 8,758 consensus forecasts at the firm-year level, per item. Table I reports summary statistics for forecast errors and revisions. We now use these variables to

[11] Bouchaud et al. (2019) compute the consensus forecast as the median across individual forecasts.

Table I: Summary Statistics

This table reports summary statistics for the variables used in the main analysis. The sample is restricted to analysts who, for the same firm and fiscal year, issue both EPS and CAPEX forecasts on the same dates at both the one-year-ahead and two-year-ahead horizons. Forecast errors and revisions are scaled by lagged stock price (EPS) and lagged PP&E (CAPEX). Observations beyond 5 interquartile ranges from the median are removed.

	N	**Mean**	**SD**	**25%**	**Median**	**75%**
$x_{j,t+1} - F_{i,t}x_{j,t+1}$						
EPS	16,687	-0.0003	0.02	-0.0061	0.0007	0.0064
CAPEX	16,687	-0.0074	0.07	-0.0347	-0.0070	0.0173
$F_{i,t}x_{j,t+1} - F_{i,t-1}x_{j,t+1}$						
EPS	16,687	-0.0063	0.03	-0.0141	-0.0024	0.0042
CAPEX	16,687	0.0081	0.08	-0.0246	0.0035	0.0363
$x_{j,t+1} - F_{t}x_{j,t+1}$						
EPS	8,758	-0.0006	0.02	-0.0067	0.0006	0.0064
CAPEX	8,758	-0.0082	0.07	-0.0368	-0.0080	0.0174
$F_{t}x_{j,t+1} - F_{t-1}x_{j,t+1}$						
EPS	8,758	-0.0064	0.03	-0.0147	-0.0024	0.0043
CAPEX	8,758	0.0090	0.08	-0.0241	0.0040	0.0373
Revision days						
H1	16,687	9.9	18.2	1	1	7
H2	16,687	13.8	21.7	1	2	17

estimate the degree of under- and overreaction in analyst forecasts across items.

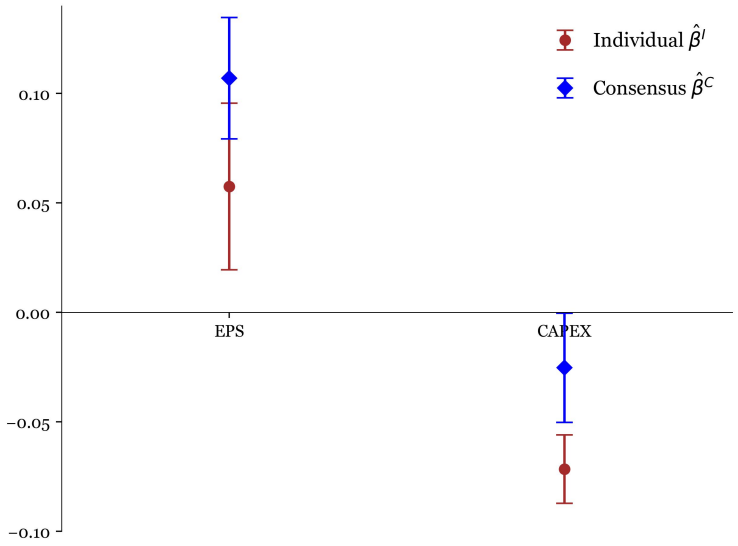
3.2 Unconditional Estimates

We first analyze the predictability of forecast errors from forecast revisions at both the consensus and individual levels for EPS and CAPEX by estimating Equations (1) and (2).

Figure II plots the estimates, and the results for both regressions are reported in Table II. Panel A in Table II shows the results using consensus forecasts. For EPS forecasts, the coefficient is positive and statistically significant, with $\hat{\beta}^C = 0.107$ in column (1). This indicates underreaction in the consensus forecast for short-term EPS, consistent with the previous literature. The estimated value for the EPS consensus forecast is close to what Bouchaud et al. (2019) find using firm-level EPS consensus forecasts for fiscal years ending between 1989 and 2015 (they find $\beta \approx 0.16$). Our smaller estimated β is mainly because our sample period spans 2007 to 2025, during which analyst forecasts exhibit a smaller magnitude of underreaction relative to earlier periods, as suggested by Table B.1 in the Appendix. In contrast, the coefficient for the CAPEX sample is -0.025, suggesting overreaction in CAPEX

Figure II: CG Coefficient Estimates

This figure plots the estimated coefficients $\hat{\beta}^C$ (consensus) and $\hat{\beta}^I$ (individual) from Equations (1) and (2) for EPS and CAPEX, corresponding to Panel A column (1) and Panel B column (1) in Table II, respectively. Horizontal bars show 90% confidence intervals.



consensus forecasts.

Panel B in Table II shows the results using individual forecasts. Although the magnitude of the coefficient decreases, the coefficient remains positive and significant for EPS forecasts at the individual level, with $\hat{\beta}^I = 0.058$. In contrast, the coefficient for individual CAPEX forecasts is -0.072 and significant at the 1% level.

The regression results show that we have β^C and β^I with the same sign but different magnitudes for EPS and CAPEX forecasts. It is important to note that for macroeconomic variables, as documented in Bordalo et al. (2020), the evidence is $\beta^C > 0$ while $\beta^I < 0$.

3.3 Forecast Horizon

Next, we examine the term structure of CG coefficients by estimating CG coefficients separately for each forecast horizon, from one-year-ahead through four-year-ahead,[12] i.e., we compute the CG coefficient for forecasts of x_{t+1}, x_{t+2}, and longer-horizon targets (x_{t+3}, x_{t+4}). Survey evidence indicates that overreaction is stronger at longer horizons. Bordalo et al. (2019) find overreaction in long-term earnings growth forecasts, while Bouchaud et al. (2019) document underreaction for short-term earnings. The same pattern holds at the consensus level for earnings (Ke, 2024). Similarly, for interest rates, the term structure is downward sloping (d'Arienzo, 2020). The same patterns hold in a lab setting: Afrouzi et al.

[12]Since computing the CG revision at horizon k requires the $(k+1)$-year-ahead forecast as the prior, five-year-ahead forecasts enter only as inputs to the four-year-ahead revision.

Table II: Forecast Errors on Forecast Revision

This table reports regression estimates of the predictability of forecast errors from forecast revisions at both the consensus and individual analyst levels, corresponding to Equations (1) and (2). Results are presented separately for EPS and CAPEX. Forecast errors and revisions are scaled by lagged price (for EPS) and lagged PP&E (for CAPEX). Panel A reports consensus-level results: columns (1) and (3) use pooled OLS; columns (2) and (4) include firm fixed effects. Standard errors are clustered by firm and year. Panel B reports individual-level results: columns (1) and (3) include analyst fixed effects; columns (2) and (4) add year fixed effects. Standard errors are clustered by analyst and year. The sample is restricted to analysts who, for the same firm and fiscal year, issue both EPS and CAPEX forecasts on the same dates at both the one-year-ahead and two-year-ahead horizons. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

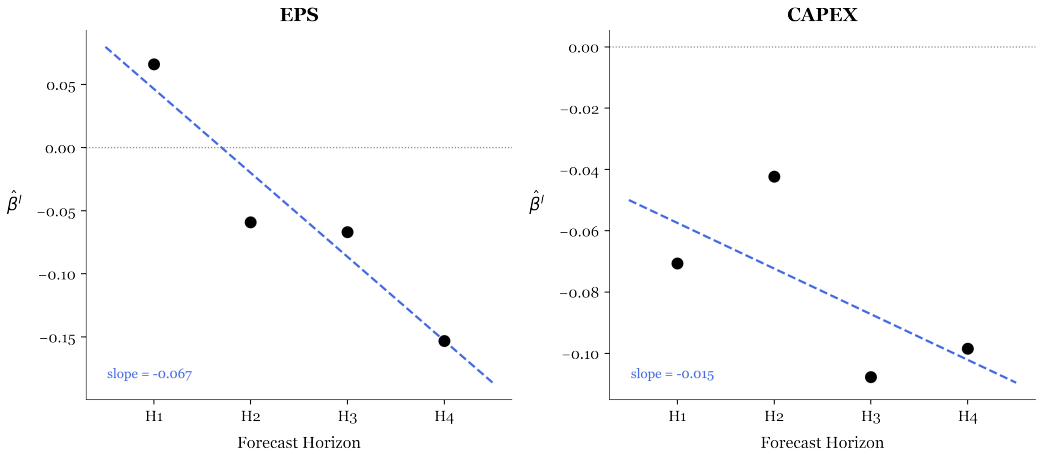
	EPS		**CAPEX**	
	(1)	(2)	(3)	(4)
Panel A: Consensus				
Forecast Revision	0.107***	0.047**	−0.025*	−0.064***
	(0.017)	(0.020)	(0.015)	(0.021)
Observations	8,758	8,758	8,758	8,758
R-squared	0.019	0.004	0.001	0.006
Firm FE		✓		✓
Cluster	Firm-Year	Firm-Year	Firm-Year	Firm-Year
Panel B: Individual				
Forecast Revision	0.058**	0.051**	−0.072***	−0.083***
	(0.023)	(0.024)	(0.010)	(0.011)
Observations	16,687	16,687	16,687	16,687
R-squared	0.005	0.004	0.007	0.009
Analyst FE	✓	✓	✓	✓
Year FE		✓		✓
Cluster	Analyst-Year	Analyst-Year	Analyst-Year	Analyst-Year

(2023), for example, find that overreaction is stronger for longer-horizon forecasts.

We document that the same pattern holds for analysts' CAPEX forecasts as well. Figure III plots the individual-level CG coefficient at each horizon and shows a downward-sloping pattern consistent with stronger overreaction at longer horizons. Table III confirms this formally: at the individual level, the CG coefficient is −0.069 for short-term forecasts (H1–H2) and −0.105 for long-term forecasts (H3+). A similar pattern holds at the consensus level, where the coefficient decreases from −0.087 for short-term forecasts to −0.329 for long-term forecasts. These results imply that overreaction in CAPEX forecasts strengthens with the forecast horizon, consistent with a downward-sloping term structure of CG coefficients. Overall, our findings for CAPEX closely mirror the patterns documented in the survey and experimental literature.

Figure III: Forecast Horizon and the Individual CG Coefficient

Each panel plots the individual-level CG coefficient $\hat{\beta}^I$ from Equation (2), estimated separately at each forecast horizon H_k ($k = 1, \ldots, 4$). The left panel shows EPS (scaled by lagged price) and the right panel shows CAPEX (scaled by lagged PP&E). All regressions include analyst fixed effects with standard errors clustered by analyst and year. The dashed line is the OLS best fit across the four horizon-specific point estimates. Each item uses its full sample from 2007 to 2025 to maximize coverage.



3.4 Experience Effect

Next, we examine how analyst experience shapes expectation formation. Table IV does so by interacting forecast revisions with an indicator for junior analysts. For EPS forecasts, the baseline CG coefficient is positive, consistent with underreaction. The interaction term with the junior indicator is also positive and statistically significant at the 10% level, indicating that junior analysts exhibit stronger underreaction than their more experienced peers. This pattern is consistent with Bouchaud et al. (2019) and suggests that junior analysts update more conservatively.

A natural interpretation of this evidence is that experience makes analysts more rational, moving the CG coefficient closer to zero. However, the results for CAPEX forecasts challenge this interpretation. The baseline CG coefficient is negative, indicating overreaction, and the interaction term is positive, implying that junior analysts overreact significantly less than senior analysts. In other words, more experienced analysts exhibit more extreme overreaction. These results are illustrated in Figure IV.

Importantly, these findings suggest that experience does not uniformly move analysts toward rational expectations. Instead, it amplifies the prevailing bias: experienced analysts underreact less in settings characterized by underreaction (EPS), but overreact more in settings characterized by overreaction (CAPEX).

Table III: Forecast Horizon and the CG Coefficient

This table reports CG regression estimates at short and long forecast horizons for EPS and CAPEX. Individual-level regressions estimate Equation (2) with analyst FE and horizon FE, and standard errors clustered by analyst and year. Consensus-level regressions estimate Equation (1) with firm FE and horizon FE, and standard errors clustered by firm and year. The short-term pools one- and two-year horizons; the long-term pools horizons from three years ahead. Each item uses its full sample from 2007 to 2025 to maximize coverage. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Short-term (H1–H2)		**Long-term (H3+)**	
	Individual	Consensus	Individual	Consensus
Panel A: EPS				
Forecast Revision	0.042^{***}	0.028^{*}	-0.110^{*}	-0.294^{***}
	(0.012)	(0.017)	(0.060)	(0.061)
Observations	201,360	49,632	13,380	8,884
Analyst FE	✓		✓	
Firm FE		✓		✓
Horizon FE	✓	✓	✓	✓
Cluster	Analyst-Year	Firm-Year	Analyst-Year	Firm-Year
Panel B: CAPEX				
Forecast Revision	-0.069^{***}	-0.087^{***}	-0.105^{**}	-0.329^{***}
	(0.011)	(0.015)	(0.052)	(0.056)
Observations	36,804	17,376	4,993	3,768
Analyst FE	✓		✓	
Firm FE		✓		✓
Horizon FE	✓	✓	✓	✓
Cluster	Analyst-Year	Firm-Year	Analyst-Year	Firm-Year

4 Model

In this section, we propose a model of endogenous belief formation in which analysts face strategic considerations and engage in wishful thinking when producing forecasts. The model captures two important features. First, an analyst's payoffs depend on both the accuracy of her forecast and how closely it tracks the consensus forecast. Second, an analyst faces uncertainty about the precision of her information and engages in wishful thinking when choosing the perceived precision of this information. We show how this stylized model of forecasts is able to reconcile the existing empirical evidence discussed above, and provide a novel set of predictions which we discuss further in Section 4.3.

Specifically, we assume that analyst i chooses to submit a forecast k_i of a target random variable θ that maximizes her payoff:

$$u_i(k_i, K, \theta) \equiv -(k_i - \theta)^2 - \rho(k_i - K)^2, \tag{3}$$

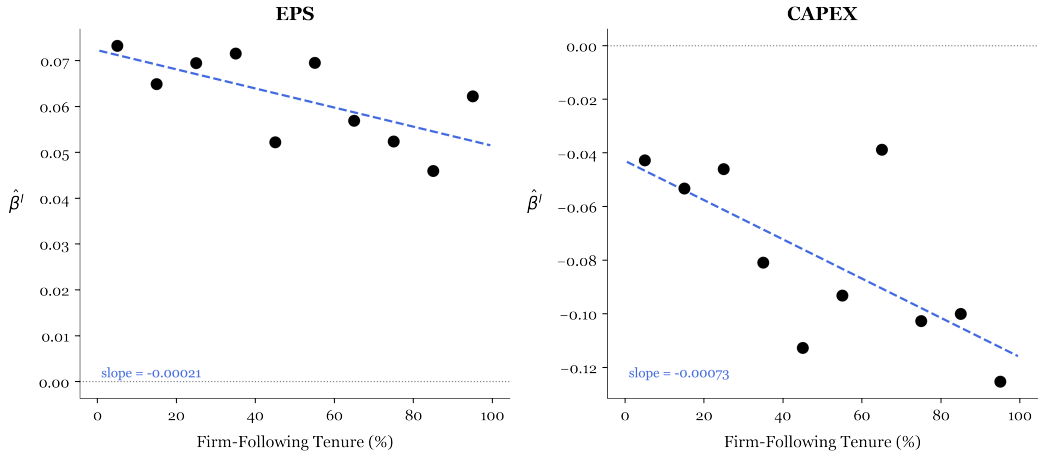
Table IV: Analyst Experience and the CG Coefficient

This table reports CG regression estimates with an interaction between the forecast revision and a Junior analyst indicator. Junior equals one if the analyst's firm-following tenure falls in the bottom 20th percentile within each fiscal year, where tenure is measured as the number of years since the analyst's first forecast for the firm. Individual-level regressions include analyst fixed effects with standard errors clustered by analyst and year. Consensus-level regressions aggregate individual variables to the firm-year level. Standard errors are clustered by firm and year. Each item uses its full sample from 2007 to 2025 to maximize coverage. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	EPS		**CAPEX**	
	Individual	Consensus	Individual	Consensus
REV	0.0627^{***}	0.1363^{***}	-0.0778^{***}	-0.0322^{**}
	(0.0135)	(0.0167)	(0.0115)	(0.0135)
REV × Junior	0.0172^{*}	0.0328^{**}	0.0341^{***}	0.0422^{*}
	(0.0096)	(0.0135)	(0.0132)	(0.0238)
Observations	159,671	34,160	26,955	11,909
Analyst FE	✓		✓	
Cluster	Analyst-Year	Firm-Year	Analyst-Year	Firm-Year

Figure IV: Analyst Experience and the Individual CG Coefficient

This figure plots the individual-level CG coefficient $\hat{\beta}^I$ from Equation (2), estimated within ten equal-sized bins of firm-following tenure (defined as years since the analyst's first forecast for the firm). All regressions include analyst fixed effects with standard errors clustered by analyst and year. The left panel shows EPS (scaled by lagged price) and the right panel shows CAPEX (scaled by lagged PP&E). The horizontal axis reports the percentile midpoint of each tenure bin; the dashed line is the OLS best fit across the ten bin-level estimates. Each item uses its full sample from 2007 to 2025 to maximize coverage.



where $K = \int_i k_i di$ is the consensus forecast. The first term captures the incentive to produce an accurate forecast. In the second term, the parameter $\rho \in (-1, 1)$ reflects the strategic considerations that the analyst faces: when $\rho < 0$, analysts want to stand out (i.e., forecasts are strategic substitutes), while when $\rho > 0$, analysts want to herd (i.e., forecasts are strategic complements). This is in line with strategic considerations that analysts face in practice (e.g., see Croushore (1997), Ottaviani and Sørensen (2006), Marinovic et al. (2010), Marinovic et al. (2013), Gemmi and Valchev (2026), and Banerjee et al. (2026)).

Moreover, assume that the target $\theta \sim N(0, 1/\tau)$ and analyst i observes both a private signal s_i and a public signal s, where

$$s_i = \theta + \varepsilon_i, \quad \text{and} \quad s = \theta + \eta, \tag{4}$$

where the common error has distribution $\eta \sim N(0, 1/\tau_\eta)$, the individual errors have distribution $\varepsilon_i \sim N(0, 1/\tau_e)$, and all random variables $\{\theta, \eta, \{\varepsilon_i\}_i\}$ are mutually independent.

4.1 Rational Expectations

First, consider a benchmark in which analysts know the precision of their information (i.e., τ_η and τ_e) and have rational expectations. Given analyst i's payoffs, beliefs and the realizations of s_i and s, her optimal action, k_i is

$$k_i^*(s_i, s) = r\mathbb{E}[K|s_i, s] + (1 - r)\mathbb{E}[\theta|s_i, s], \tag{5}$$

where $r \equiv \frac{\rho}{1+\rho}$ and $\mathbb{E}[\cdot]$ denotes the expectation under the objective joint distribution of fundamentals and signals. Bayesian updating implies that analyst i's conditional beliefs about θ are given by

$$\mathbb{E}[\theta|s_i, s] = As_i + Bs, \text{ and } \operatorname{var}[\theta|s_i, s] = \frac{1 - A - B}{\tau}, \tag{6}$$

where analyst i's weights on the private and public signals are given by:

$$A \equiv \frac{\tau_e}{\tau + \tau_e + \tau_\eta}, \text{ and } B \equiv \frac{\tau_\eta}{\tau + \tau_e + \tau_\eta}. \tag{7}$$

The following lemma characterizes equilibrium actions.

Lemma 1. *A linear equilibrium exists, is unique, and is characterized by*

$$k_i(s_i, s) = \alpha s_i + \beta s, \quad \textit{and} \quad K = \alpha\theta + \beta s, \tag{8}$$

where $\alpha = \frac{1-r}{1-rA}A$ *and* $\beta = \frac{B}{1-rA}$.

Note that since we have normalized the unconditional mean of the target as zero (i.e., $\mathbb{E}[\theta] = 0$), one can interpret analyst i's forecast $k_i(s_i, s)$ as the forecast revision (FR) after she conditions on public and private information, i.e., $FR_i = k_i(s_i, s)$. Similarly, we can

define the forecast error (FE) for analyst i as

$$FE_i \equiv \theta - k_i(s_i, s). \tag{9}$$

We can define the consensus forecast revision $\bar{FR}$ and consensus forecast error $\bar{FE}$ by taking expectations across analysts for the corresponding expressions.

This implies that the Coibion–Gorodnichenko regression coefficient for individual forecasts, CG_i, and for consensus forecasts, CG_a, are given by:

$$CG_i = \frac{\text{cov}(FE_i, FR_i)}{\text{var}(FR_i)}, \quad \text{and} \quad CG_a = \frac{\text{cov}(\bar{FE}, \bar{FR})}{\text{var}(\bar{FR})}. \tag{10}$$

These correspond to the coefficients $\hat{\beta}^I$ and $\hat{\beta}^C$ in Section 3. Under rational expectations, this implies the following result.

Proposition 1. *When forecasters exhibit rational expectations, then*

(i) consensus forecasts always exhibit underreaction ($CG_a > 0$), and

(ii) individual forecasts exhibit overreaction ($CG_i < 0$) if and only if $\rho < 0$.

Note that irrespective of strategic considerations, consensus forecasts always exhibit underreaction. The weight that an individual forecaster puts on their private signal reflects the idiosyncratic error (ε_i). However, the consensus forecast reflects an aggregate of these private signals, which is more informative than any individual signal because the errors are independent. As a result, the consensus forecast appears to underreact to the (aggregate) information available to analysts.[13]

Under rational expectations, individual forecasts exhibit a bias if and only if analysts face strategic considerations. When $\rho = 0$, analysts want to maximize accuracy of forecasts and under rational expectations, put the statistically optimal weight on their private and public information. In turn, this implies that individual forecast errors are not predictable (i.e., $CG_i = 0$). In contrast, when $\rho > 0$, analysts have an additional incentive to herd with the consensus forecast - this leads them to put less weight on the private signal but more weight on the public signal, which in turn, leads to underreaction in individual forecasts (i.e., $CG_i > 0$). Similarly, when $\rho < 0$, analysts have an incentive to stand out, which leads them to put relatively more weight on their private signals and, consequently, leads to overreaction (i.e., $CG_i < 0$).

The above results imply that strategic considerations alone are insufficient to account for the facts documented in Section 3.2.[14] In particular, under rational expectations, this

[13] This is the mechanism that Coibion and Gorodnichenko (2015) propose as an explanation for consensus underreaction (as in the noisy information model of Woodford (2003)).

[14] Note that Gemmi and Valchev (2026) use the case of $\rho < 0$ to reconcile individual overreaction and consensus underreaction documented in macroeconomic forecasts (e.g., Bordalo et al. (2020)).

setting cannot simultaneously generate both overreaction (at both individual and consensus level for CAPEX) and underreaction (for EPS) within the same analysts and firms, nor can it explain why the direction and magnitude of these biases vary systematically with forecast horizon, information type, and analyst characteristics.

4.2 Wishful Thinking

Our model is a special case of the generalized setting formalized in Banerjee et al. (2026). We assume that analysts face uncertainty about the quality of their private and public information and allow them to entertain subjective beliefs about the precision of such information.

Specifically, we assume that analyst i perceives the error in her private and public signals to be

$$\varepsilon_i \sim_i N\left(0, \frac{1}{\delta_{e,i}\tau_e}\right) \qquad \text{and} \qquad \eta \sim_i N\left(0, \frac{1}{\delta_{\eta,i}\tau_\eta}\right), \tag{11}$$

respectively, where she chooses $\delta_{e,i}$ and $\delta_{\eta,i}$ as described below. If $\delta_{e,i} = \delta_{\eta,i} = 1$, analyst i's beliefs coincide with the objective distribution, and so she exhibits rational expectations. However, when $\delta_{e,i}$ ($\delta_{\eta,i}$) is greater than one, analyst i overestimates the precision of the private (public) signal when forming expectations; when $\delta_{e,i}$ ($\delta_{\eta,i}$) is less than one, she underestimates the precision of the information.

Forecasts given subjective beliefs. Given these subjective beliefs, analyst i's forecast $k_i^*(\delta_{e,i}, \delta_{\eta,i})$ maximizes her expected payoff, i.e.,

$$k_i^*(s_i, s; \delta_{e,i}, \delta_{\eta,i}) \equiv \arg\max_{k_i} \mathbb{E}_i\left[u_i(k_i, K, \theta) \,\middle|\, s_i, s\right] \tag{12}$$

$$= r\mathbb{E}_i[K|s_i, s] + (1-r)\,\mathbb{E}_i[\theta|s_i, s], \tag{13}$$

where $r = \frac{\rho}{1+\rho}$ and $\mathbb{E}_i[\cdot]$ denotes the analyst's *subjective* expectation, given her choice of $\delta_{e,i}$ and $\delta_{\eta,i}$.

Since the joint distribution of fundamentals and signals is normal, Bayesian updating implies that analyst i's conditional beliefs about θ are given by

$$\mathbb{E}_i[\theta|s_i, s] = A_i s_i + B_i s, \text{ and } \mathrm{var}_i[\theta|s_i, s] = \frac{1 - A_i - B_i}{\tau}, \tag{14}$$

where analyst i's weights on the private and public signals are given by:

$$A_i \equiv \frac{\delta_{e,i}\tau_e}{\tau + \delta_{e,i}\tau_e + \delta_{\eta,i}\tau_\eta}, \text{ and } B_i \equiv \frac{\delta_{\eta,i}\tau_\eta}{\tau + \delta_{e,i}\tau_e + \delta_{\eta,i}\tau_\eta}. \tag{15}$$

We define the aggregate weight on private information, $A \equiv \int_i A_i di$, and the public information, $B \equiv \int_i B_i di$, analogously. The following lemma characterizes equilibrium forecasts, given subjective beliefs.

Lemma 2. *Given a choice of subjective beliefs* $\{\delta_{e,i}, \delta_{\eta,i}\}_i$ *for each analyst* i*, a linear equilibrium exists, is unique, and is characterized by*

$$k_i\left(s_i, s; \delta_{e,i}, \delta_{\eta,i}\right) = \alpha_i s_i + \beta_i s \qquad \text{and} \qquad K = \alpha\theta + \beta s,$$

where $\alpha_i = \frac{1-r}{1-rA} A_i$, $\beta_i = \frac{(1-r)B_i + Br}{1-rA}$, $\alpha = \frac{A(1-r)}{1-rA}$, *and* $\beta = \frac{B}{1-rA}$.

Analyst i's forecast puts weight α_i and β_i on her private and public information, respectively. Importantly, these weights depend not only on strategic considerations (as in the rational expectations benchmark), but also on the choice of subjective beliefs.

Subjective belief choice. We next characterize the key trade-off that an analyst faces when choosing her subjective beliefs. Denote analyst i's **anticipatory utility** by:

$$AU_i\left(\delta_{e,i}, \delta_{\eta,i}\right) \equiv \mathbb{E}_i\left[u\left(k_i^*\left(s_i, s; \delta_{e,i}, \delta_{\eta,i}\right), K, \theta\right)\right], \tag{16}$$

and her expected **experienced utility**, given her chosen action, by

$$EU_i\left(\delta_{e,i}, \delta_{\eta,i}\right) = \mathbb{E}\left[u\left(k_i^*\left(s_i, s; \delta_{e,i}, \delta_{\eta,i}\right), K, \theta\right)\right]. \tag{17}$$

Specifically, her anticipatory utility reflects the contemporaneous utility flow she receives from the anticipation of her (subjective) expected payoff given her choice of beliefs $\{\delta_{e,i}, \delta_{\eta,i}\}$ and forecast $k_i^*(\cdot)$. Similarly, experienced utility reflects the expected payoff under the objective distribution that results from her (distorted) forecast $k_i^*(\cdot)$.

Each analyst chooses her subjective beliefs $\{\delta_{e,i}, \delta_{\eta,i}\}$ to maximize a weighted average of her anticipatory utility and her experienced utility. Formally, analyst i's objective function is

$$\max_{\delta_{e,i}, \delta_{\eta,i}} \frac{1}{1+\psi}\left\{AU_i(\delta_{e,i}, \delta_{\eta,i}) + \psi EU_i(\delta_{e,i}, \delta_{\eta,i})\right\} \equiv TU_i(\delta_{e,i}, \delta_{\eta,i}), \tag{18}$$

where $\psi \geq 0$ scales the relative importance of experienced utility versus anticipatory utility. The objective in (18) reflects the tradeoff between forming "desirable" beliefs (which increase anticipatory utility) versus "accurate" beliefs (which increase experienced utility).

When ψ is low, the analyst puts a lot of weight on anticipatory utility, which leads her to distort her subjective beliefs more (i.e., choose $\delta_{e,i}$ and $\delta_{\eta,i}$ further from 1) in an effort to improve her subjective perception of the future expected payoffs. To understand how she distorts her beliefs, note that anticipatory utility is proportional to the sum of two subjective conditional variances:

$$AU_i\left(\delta_{e,i}, \delta_{\eta,i}\right) \propto \underbrace{-\mathrm{var}_i\left[\theta | s_i, s\right]}_{\text{fundamental uncertainty channel}} - \underbrace{\rho\, \mathrm{var}_i\left[K|\theta\right]}_{\text{aggregate-errors channel}}, \tag{19}$$

The above implies that an analyst's subjective beliefs affect her anticipatory utility through two channels. First, anticipatory utility is higher when subjective forecast accuracy is higher

(i.e., when $\text{var}_i[\theta|s_i, s]$ is lower) — we refer to this as the **fundamental uncertainty channel**. This channel creates an incentive to overestimate the precision of both private and public information.

Second, the analyst's anticipatory utility is affected by her beliefs about the excess volatility in the consensus forecast K, which is driven by $\rho\text{var}_i[K|\theta]$. We refer to this as the **aggregate-errors channel**, because it captures the impact of "errors" in the aggregate action (relative to the full information case). Specifically, when analysts have an incentive to stand out (i.e., $\rho < 0$), larger errors in the consensus forecasts improve anticipatory utility, which leads towards a tendency to underestimate the precision of public information. However, when analysts have an incentive to herd (i.e., $\rho > 0$), larger errors reduce anticipated payoffs, and so there is an incentive to overestimate the precision of public information.

As ψ increases, the analyst puts relatively more weight on experienced utility, which penalizes subjective beliefs that are distorted away from rational expectations. In fact, when $\psi \to \infty$, this implicit cost of belief distortion becomes prohibitive, and the analyst's subjective beliefs converge to the rational expectations benchmark, i.e., $\delta_{\eta,i} = \delta_{e,i} = 1$.

Given these observations, the following proposition characterizes subjective belief choice when analysts have an incentive to herd (i.e., $\rho > 0$).

Proposition 2. *When $\rho > 0$, a unique equilibrium exists, and*

(i) If $\psi \leq \frac{(\tau_e+(1+\rho)\tau_\eta)}{2\tau}$, then all analysts choose to ignore their prior and choose $A_i = \frac{\tau_e}{\tau_e+\tau_\eta}$ and $B_i = \frac{\tau_\eta}{\tau_e+\tau_\eta}$.

(ii) If $\psi > \frac{(\tau_e+(1+\rho)\tau_\eta)}{2\tau}$, then all analysts choose

$$\delta_{e,i} = \frac{2\tau\psi\left(2\psi+1\right)+\left(\rho+1\right)\rho\tau_\eta}{2\psi\left(2\tau\psi-\left(\rho+1\right)\tau_\eta-\tau_e\right)} > 1, \quad \textit{and} \quad \delta_{\eta,i} = \delta_{e,i} + \frac{\rho}{2\psi} > 1 \tag{20}$$

Moreover, δ_e and δ_η weakly decrease in ψ.

The intuition for this result follows from the earlier discussion. Since $\rho > 0$, both the fundamental uncertainty channel and the aggregate-errors channel lead the analysts to overestimate the precision of public and private information. The extent to which these beliefs are distorted depends on the relative cost ψ. When the cost of belief distortion is sufficiently low (i.e., when $\psi \leq \frac{\tau_e+(1+\rho)\tau_\eta}{2\tau}$), analysts overestimate the signal precisions maximally: their subjective beliefs completely ignore the prior beliefs about fundamentals, i.e., they choose $A_i = \frac{\tau_e}{\tau_e+\tau_\eta}$ and $B_i = \frac{\tau_\eta}{\tau_e+\tau_\eta}$. When costs are higher (i.e., $2\psi\tau > (\tau_e + (1+\rho)\tau_\eta)$), the analysts still choose to overestimate the precision of their information (i.e., choose $\delta_{e,i}, \delta_{\eta,i} > 1$), but the distortion is not as extreme and decreases

with the cost parameter ψ. Furthermore, since $\rho > 0$, the aggregate-errors channel induces an incremental marginal benefit of increasing perceived public precision, and consequently, $\delta_{\eta,i} > \delta_{e,i}$.

The following result characterizes how the above subjective belief choice affects the CG regression coefficients for both individual and consensus forecasts.

Proposition 3. *When forecasters exhibit wishful thinking and have an incentive to herd (i.e., $\rho > 0$), then:*

(i) If ψ is sufficiently low, both consensus and individual forecasts overreact to information i.e., $cov\left(\bar{FE}, \bar{FR}\right) < 0$ and $cov(FE_i, FR_i) < 0$.

(ii) For intermediate ψ, consensus forecasts underreact and individual forecasts overreact to information i.e., $cov\left(\bar{FE}, \bar{FR}\right) > 0$ and $cov(FE_i, FR_i) < 0$.

(iii) If ψ is sufficiently high, consensus forecasts and individual forecasts underreact to information i.e., $cov\left(\bar{FE}, \bar{FR}\right) > 0$ and $cov(FE_i, FR_i) > 0$.

As ψ increases, both CG_i and CG_a increase.

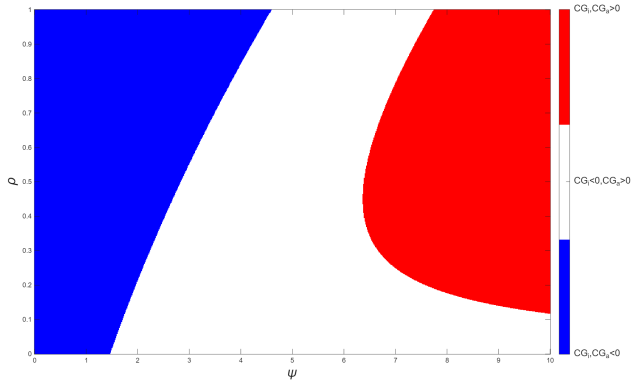
Proposition 3 presents a key takeaway of our analysis. In our model, whether individual and consensus forecasts feature overreaction or underreaction depends on a single parameter: the relative cost of belief distortion ψ. Figure V provides an illustration by plotting the implied CG coefficients for individual and consensus forecasts for different values of ρ and ψ. When the analysts have incentives to herd (i.e., $\rho > 0$), a low cost of belief distortion (low ψ) leads to overreaction in both consensus and individual forecasts, while a high cost leads to underreaction in both forecasts. It is worth noting that when the cost of belief distortion is intermediate, our model is also able to generate overreaction in individual forecasts and underreaction in consensus forecasts (in line with the evidence on macroeconomic forecasts in Coibion and Gorodnichenko (2015) and Bordalo et al. (2020)).

The contrast between EPS and CAPEX forecasts provides a sharp illustration of this mechanism. EPS is the central object of analyst evaluation: earnings are realized within months, errors are immediately observable, and forecast accuracy is directly tied to analyst reputation and career outcomes. This rapid, high-stakes feedback imposes strong discipline on belief formation, corresponding to high ψ. Our model therefore predicts underreaction in EPS forecasts, consistent with the evidence in Panel A of Table II.

CAPEX outcomes are fundamentally different. Capital expenditure plans are realized over multiple years, subject to revision, and tend to play a smaller role in the criteria by which analysts are evaluated. The feedback loop is slower, noisier, and more disconnected

Figure V: CG Coefficients as a Function of ρ and ψ

This figure plots the Coibion–Gorodnichenko (CG) coefficients at both the individual level and consensus level. The horizontal axis corresponds to the cost of belief distortion, ψ and the y-axis corresponds to ρ. There are three regions: blue corresponds to overreaction in both individual and consensus forecasts; red corresponds to underreaction in both forecasts; and white corresponds to overreaction in individual forecasts and underreaction in consensus forecasts. Other parameters are $\tau = \tau_\eta = \tau_e = 1$ and $\bar{\delta} = 1.8$.



from analyst accountability. This corresponds to low ψ: the cost of distorting beliefs about CAPEX is small, giving analysts wide scope for wishful thinking. Our model therefore predicts overreaction in CAPEX forecasts, again consistent with Table II.

Crucially, both sets of forecasts are produced by the same analysts about the same firms at the same time. The coexistence of underreaction and overreaction is therefore difficult to attribute to differences in information or firm characteristics — instead, it reflects the differential discipline imposed on belief formation across forecast targets.

Figure VI illustrates how individual and consensus CG coefficients vary with the cost of belief distortion ψ. As ψ increases, the model transitions from joint overreaction ($CG_i < 0$ and $CG_a < 0$), to an intermediate region ($CG_i < 0 < CG_a$) with individual overreaction and consensus underreaction, and ultimately to joint underreaction ($CG_i > 0$ and $CG_a > 0$). Importantly, this transition is monotonic: both CG_i and CG_a increase with ψ. This monotonic relationship provides a sharp and testable implication of the model.

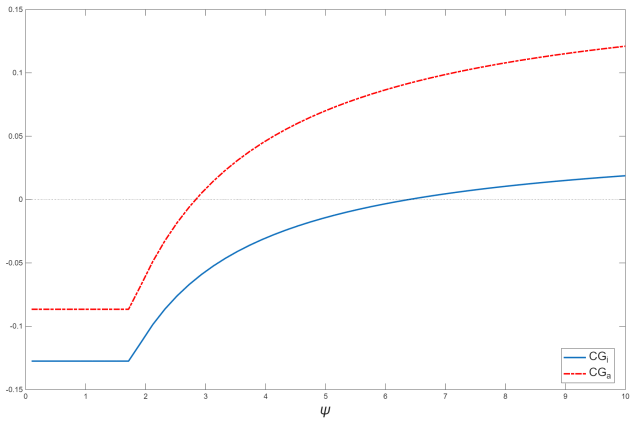
4.3 Additional Testable Predictions

Our model predicts that the cost of belief distortion should be a critical determinant of the cross-sectional dispersion in analysts' forecasts and volatility of consensus forecasts due to its impact on the analysts' subjective beliefs. This generates novel predictions of our model on the moments of the forecast distribution across analysts and over time.

Specifically, let $\sigma_k^2 \equiv \int (k_i - K)^2 \, di$ denote the (cross-sectional) dispersion in analysts' forecasts and let $\nu^2 \equiv \mathbb{E}\left[(K)^2\right]$ denote the variance in consensus forecasts. The following

Figure VI: CG Coefficients and the Cost of Belief Distortion

This figure plots the Coibion–Gorodnichenko (CG) coefficients at both the individual level and consensus level. The horizontal axis corresponds to the cost of belief distortion, ψ. Other parameters are $\tau = \tau_\eta = \tau_e = 1$, $\rho = 0.5$, and $\bar{\delta} = 1.8$.



result characterizes how these moments vary with the cost of belief distortion ψ.

Proposition 4. *The following results hold:*

- *Dispersion in analysts' forecasts decreases with the cost of belief distortion ψ.*
- *Variance in consensus forecasts decreases with the cost of belief distortion ψ.*

To gain some intuition for this result, note that one can express dispersion and consensus volatility as:

$$\sigma_k^2 = \frac{\alpha^2}{\tau_e} \qquad \text{and} \qquad \nu^2 = \frac{(\alpha+\beta)^2}{\tau} + \frac{\beta^2}{\tau_\eta}, \tag{21}$$

respectively. When ψ is low, the penalty for distorting beliefs is low and analysts choose to overestimate the precision of public and private signals (i.e., choose high $\delta_{e,i}$ and high $\delta_{\eta,i}$). This results in a large cross-sectional dispersion in analyst forecasts, because the weight α on private signal errors is high. Moreover, it results in high consensus volatility, because the weight β on the public signal is also high.

As ψ increases, deviating from rational expectations becomes more costly, so both subjective beliefs shrink and move toward one. In equilibrium, this reduces both loadings, with α falling (hence lower dispersion across agents) and β falling (hence lower sensitivity of consensus to common public news). This implies that both cross-sectional dispersion and consensus variance decline as the cost of belief distortion increases.

Table V summarizes the model's comparative statics with respect to an increase in the cost of belief distortion. In the next section, we discuss how empirical evidence aligns with these and other predictions of our model.

Table V: Comparative Statics with Cost of Belief Distortion (ψ)

Moment	**Change when ψ increases**
Coibion–Gorodnichenko coefficients (CG_i, CG_a)	$\uparrow$
Cross-sectional dispersion in forecasts (σ_k^2)	$\downarrow$
Variance of the consensus forecast (ν^2)	$\downarrow$

5 Testing the Model

Our framework predicts that the cost of belief distortion, ψ, governs how analysts process information and drives the degree of over-/under-reaction in forecasts. Lower values of ψ lead analysts to overestimate the precision of public and private information, which generates overreaction, while higher values of ψ induce more disciplined forecasts and generate underreaction. As we discuss above, our model's predictions are consistent with the stylized facts we document in Section 3 and with the evidence on macro-economic forecasts in the literature.

A key challenge in testing our model's predictions more directly is identifying proxies for the cost of belief distortion. In this section, we propose three variables — forecast horizon, analyst experience, and the share of qualitative information — that arguably generate variation in this cost. We discuss how the degree of over-/under-reaction in forecasts varies with these measures, and how they are related to forecast dispersionand consensus forecast volatility. We conclude the section with a discussion of alternate belief formation mechanisms that have been proposed in the literature.

5.1 Forecast Horizon

One source of variation in the effective cost of belief distortion is the horizon of the forecast. Short-horizon forecasts can be evaluated more quickly (since the targets are realized sooner), which makes forecast errors tightly attributable to the analysts' interpretation of information. This rapid feedback imposes strong ex ante discipline on beliefs and, consequently, increases the cost of distorting subjective beliefs, i.e., it is associated with higher effective values of ψ. By contrast, long-horizon forecasts are only evaluated with a larger lag and associated with more intervening shocks, so forecast errors are harder to attribute to belief distortions and easier to rationalize. This weakens the discipline on subjective beliefs and lowers the effective ψ.

Given this relation, our model then predicts that variation in forecast horizons should be related to different degrees of overreaction. Shorter horizons correspond to higher effective ψ and hence more disciplined updating, and so such forecasts should exhibit less overreaction;

longer horizons correspond to lower ψ and so long-horizon forecasts should exhibit less underreaction or stronger overreaction. This is consistent with the evidence we document in Table III in Section 3.3.

More broadly, evidence from survey expectations and controlled experiments shows that overreaction becomes stronger as the forecast horizon increases, consistent with the mechanism we propose. Using equity analyst data, Bouchaud et al. (2019) show that short-term earnings forecasts tend to exhibit underreaction, whereas Bordalo et al. (2019) find that expectations about long-run firm performance overreact to news. Afrouzi et al. (2023) document that the degree of overreaction rises monotonically with the horizon, both in experimental settings and in survey data, and that models calibrated to short-horizon forecasts systematically underpredict the degree of overreaction at longer horizons. This pattern is also present in macroeconomic expectations: long-term interest rate forecasts exhibit significantly more overreaction than short-term forecasts (d'Arienzo, 2020).

Taken together, these findings point to a robust "term structure" of expectation biases: forecasts at short horizons tend to be more disciplined and closer to rational expectations (as in Proposition 1), while forecasts at longer horizons display greater overreaction. Through the lens of our model, this pattern arises naturally from variation in the effective cost of belief distortion ψ, with shorter horizons corresponding to higher costs and longer horizons corresponding to lower costs.

5.2 Experience and Seniority

Another source of variation in belief distortion costs is analyst experience and seniority. Career concerns and reputational incentives impose discipline on belief formation, effectively increasing ψ. Specifically, junior analysts, or those with shorter tenure, face stronger incentives to establish credibility and are subject to closer scrutiny. Their forecasts are more tightly evaluated, making distorted beliefs more costly and corresponding to higher effective values of ψ. In contrast, more experienced analysts may face weaker marginal career concerns and have greater discretion in interpreting information, corresponding to lower effective values of ψ.

The model therefore predicts that overreaction should be weaker among junior analysts and stronger among senior analysts. We proxy for experience using firm-following tenure, defined as the number of years since the analyst's first forecast for the firm, capturing variation in the strength of career concerns and hence in the effective cost of belief distortion. As shown in Table IV, we find that as experience increases (ψ decreases), the CG coefficient decreases for both EPS and CAPEX, consistent with our model's predictions.

5.3 Qualitative vs Quantitative Information

Recent evidence using textual data from analyst reports sheds light on how analysts react to different types of information. A particularly striking pattern emerges when relating over-/under-reaction to the nature of information. Using textual data from analyst reports, Ke (2024) documents a strong negative relationship between the CG coefficient and the share of qualitative information across topics. In particular, topics with a higher fraction of qualitative, intangible content — such as business strategy, product development, and R&D — exhibit significantly lower (and often negative) CG coefficients, while topics grounded in quantitative, accounting-based information display higher (and often positive) CG coefficients. This downward-sloping relationship is especially pronounced for long-horizon forecasts, where overreaction is concentrated in qualitative dimensions of information. Taken together, these findings point to systematic differences in how analysts process qualitative versus quantitative information, and highlight the importance of the nature of information in shaping expectation biases.

Our analysis provides a natural explanation for these results. Rather than attributing differences in these types of information solely to information precision, we emphasize the role of belief distortion costs in shaping how analysts process qualitative versus quantitative information. Intuitively, more qualitative information environments—that is, environments in which signals are harder to verify and evaluate ex post—should be associated with lower effective values of ψ. The subjective interpretation of such information is less tightly constrained, which gives analysts greater scope to deviate from rational expectations. By contrast, when information is quantitative, analysts' interpretations can be more easily benchmarked, making departures from rational expectations more costly. The model therefore predicts that overreaction should be more likely in response to qualitative information than to quantitative information.

5.4 Dispersion of Beliefs

Our model implies novel predictions on the cross-sectional dispersion in beliefs and the volatility of consensus forecasts. Specifically, Proposition 4 in Section 4.3 implies that forecast dispersion decreases with the cost of belief distortion, ψ. When ψ is low, analysts have greater discretion to deviate from disciplined updating, leading to more dispersed forecasts. When ψ is high, beliefs are more tightly disciplined, and individual forecasts cluster more closely around the rational benchmark, which results in less dispersion.

To test this prediction, we use forecast horizon and analyst experience as sources of variation for ψ. Longer-horizon forecasts are subject to weaker ex post discipline, corresponding to lower ψ, so the model predicts higher dispersion at longer horizons.

Similarly, more experienced analysts face weaker career concerns and greater discretion, corresponding to lower effective ψ and hence greater forecast dispersion.

Consistent with these predictions, Table VI shows that forecast dispersion is systematically higher at longer horizons and among more experienced analysts. These patterns provide supporting evidence for the role of the cost of belief distortion in shaping not only the direction of expectation biases, but also the cross-sectional distribution of beliefs.

Table VI: Forecast Dispersion and the Cost of Belief Distortion

This table reports regressions of forecast dispersion on horizon and analyst experience composition. The dependent variable is $\log(\sigma_k)$, where σ_k is the cross-sectional standard deviation of analyst forecasts within each firm-year-horizon cell. Long Horizon is an indicator for $h \geq 3$. Share Junior is the fraction of analysts in the cell whose firm-following tenure falls in the bottom 20th percentile. Columns (1) and (3) pool all horizons with firm fixed effects; columns (2) and (4) use H1 sample. Standard errors are clustered by firm and year. Each item uses its full sample from 2007 to 2025 to maximize coverage. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	EPS		**CAPEX**	
	(1)	(2)	(3)	(4)
Long Horizon	0.676***		0.427***	
	(0.049)		(0.038)	
Share Junior		−0.168**		−0.896***
		(0.085)		(0.227)
Observations	56,530	25,939	12,696	5,973
R^2	0.071	0.001	0.023	0.007
Firm FE	✓		✓	
Cluster	Firm-Year	Firm-Year	Firm-Year	Firm-Year

5.5 Variance of Consensus Forecast

Our model also has implications for how the volatility of consensus beliefs depends on the cost of belief distortion. In particular, Proposition 4 implies that the variance of the consensus forecast, ν^2, decreases with the cost of belief distortion, ψ. Intuitively, when ψ is low, analysts place greater weight on distorted common narratives and the consensus forecast moves more strongly with common shocks, increasing its volatility. When ψ is high, beliefs are more disciplined and the consensus forecast is less volatile.

To test this prediction, we again use forecast horizon and analyst experience as proxies for ψ. The model predicts higher consensus volatility at longer horizons and for more experienced analysts. Table VII provides evidence consistent with these predictions. For both EPS and CAPEX, the Long Horizon coefficient is positive and the Share Junior coefficient is negative; both are statistically significant, indicating that lower costs of belief distortion are associated

with greater volatility in consensus beliefs, in line with the model.

Table VII: Consensus Volatility and the Cost of Belief Distortion

This table reports regressions of consensus volatility on horizon and analyst experience composition. The dependent variable is $\log(\nu)$, where ν is the time-series standard deviation of the consensus forecast across fiscal years within each firm-horizon pair. Long Horizon is an indicator for $h \geq 3$. Share Junior is the average fraction of junior analysts across years. Columns (1) and (3) cluster standard errors by firm; columns (2) and (4) use H1 sample and robust standard errors. Each item uses its full sample from 2007 to 2025 to maximize coverage. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	EPS		**CAPEX**	
	(1)	(2)	(3)	(4)
Long Horizon	0.443***		0.916***	
	(0.042)		(0.146)	
Share Junior		−1.631***		−2.979***
		(0.275)		(0.981)
Observations	4,271	1,957	1,031	487
R^2	0.023	0.016	0.038	0.018
Cluster	Firm		Firm	

5.6 Alternate Mechanisms

In this section, we discuss how the stylized facts documented in Section 3 help distinguish our framework from existing theories of expectation formation. Our analysis documents a number of stylized facts about EPS and CAPEX forecasts:

1. EPS forecasts underreact whereas CAPEX forecasts overreact, both at the individual and consensus levels, within the same firms and time periods (Table II).
2. Overreaction becomes stronger at longer forecast horizons (Table III).
3. CG coefficients decline with analyst experience, with CAPEX overreaction strongest among more experienced analysts (Table IV).

These facts jointly place strong restrictions on any theory of expectation formation: a successful framework must account not only for overreaction and underreaction in isolation, but for their coexistence across forecast targets, horizons, information environments, and analyst characteristics, *within the same firms and time periods*. Our model ties all three facts to a single parameter, ψ, the cost of belief distortion. Existing approaches can match subsets of this evidence, but none accounts for the full pattern.

Rational inattention and noisy-information models. Models of noisy rational expectations or rational inattention (e.g., Woodford (2003), Maćkowiak and Wiederholt (2009), Myatt and Wallace (2012), Kohlhas and Walther (2021)) can generate consensus

underreaction, making them natural candidates for the EPS side of fact 1. However, benchmark versions typically assume that agents hold correct beliefs about the information they observe and update optimally given those signals. As such, these models do not naturally generate persistent overreaction at the individual and consensus level simultaneously, and so cannot match the evidence on CAPEX forecasts.

Strategic forecasters. Models with reputational herding incentives (e.g., Gemmi and Valchev (2026), Marinovic et al. (2010)) can generate underreaction by inducing analysts to shade toward the consensus, consistent with the EPS side of fact 1. But the same force dampens revisions uniformly across environments and cannot generate the consensus overreaction we document for CAPEX in the same firms and periods.

Diagnostic expectations. Diagnostic-expectations models (e.g., Bordalo et al. (2018), Bordalo et al. (2019), Bordalo et al. (2020)) generate overreaction by making agents overweight representative signals, offering a natural account of CAPEX overreaction and possibly the horizon pattern in fact 2 if longer-horizon forecasts rely more on narrative or representative information. However, benchmark diagnostic models primarily push in the direction of overreaction, and therefore do not naturally explain why EPS forecasts underreact in the same firms and periods.

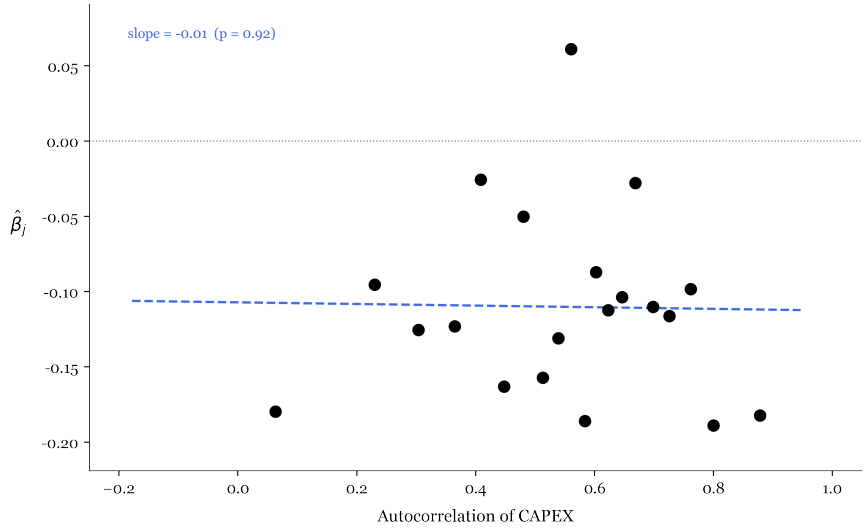
Sticky expectations and partial adjustment. Sticky-information models (e.g., Mankiw and Reis (2002), Coibion and Gorodnichenko (2012)) rationalize underreaction through sluggish updating, fitting the EPS evidence well. However, they do not naturally give rise to the overreaction needed to explain the CAPEX evidence. They also predict more sluggish adjustment at longer horizons, because forecasts are effectively staler and updated less frequently. This implies that underreaction should become stronger with horizon, which is inconsistent with fact 2.

Memory-based and information-processing explanations. Models based on imperfect recall, recency, or salience (e.g., Nagel and Xu (2022), da Silveira and Woodford (2020), da Silveira and Woodford (2022), Enke et al. (2020), Bordalo et al. (2021), Afrouzi et al. (2023)) generate overreaction through excess sensitivity to recent information. In particular, Afrouzi et al. (2023) predict that overreaction should be strongest when the underlying process is more transitory. We test this directly: Figure VII plots CG coefficients against process persistence (autocorrelation) and finds no monotonic relationship. The key comparative static does not appear to be a first-order driver of CAPEX overreaction in our data.

In summary, these approaches do not naturally account for why analysts underreact in EPS while simultaneously overreacting in CAPEX for the same firms and time periods. Moreover, the key state variables in these models (e.g., recency, salience, process persistence,

Figure VII: Process Persistence and the CG Coefficient

This figure plots firm-level CG coefficients against CAPEX process persistence. For each firm j, we estimate $\hat{\beta}_j$ by running Equation (2) at the firm level. The x-axis is the AR(1) coefficient of firm j's CAPEX. We group firms into 20 equal-sized bins based on the AR(1) and plot bin averages; the dashed line is the OLS best fit. Firms with fewer than 12 observations are excluded. Both actual and projected CAPEX are normalized by lagged PP&E.



or signal-processing frictions) do not map directly into the dimensions that organize the empirical evidence, including how the degree of overreaction should vary across horizons (fact 2), analyst experience (fact 3), and the qualitative nature of the information environment.

6 Conclusion

This paper documents a new set of facts about analyst expectations. Using matched forecasts from I/B/E/S, we show that analyst earnings (EPS) forecasts exhibit underreaction in the Coibion–Gorodnichenko regression, while capital expenditures (CAPEX) forecasts exhibit overreaction for the same firms over the same short horizons. This evidence poses a puzzle not only for rational expectations, but also for a number of alternate belief formation models.

We propose a tractable model of analyst expectations based on wishful thinking that reconciles this evidence. Analysts observe private and public information but endogenously choose subjective beliefs about the precision of this information to trade off anticipatory utility against ex post accuracy. A central object is the cost of belief distortion, ψ: when ψ is high, distorted beliefs are expensive and forecasts are disciplined, leading to underreaction; when ψ is low, analysts overweight common or narrative-driven information, generating overreaction. The framework links the sign and magnitude of reaction biases to features of the forecasting environment that shape accountability and monitoring.

The model delivers additional testable predictions about the joint behavior of forecast biases and cross-sectional moments. In low-ψ environments, forecasts load more on both sources of information, leading to higher cross-sectional dispersion and higher volatility of the consensus forecast; as ψ rises, both dispersion and aggregate volatility fall. Consistent with these implications, we show that empirically plausible proxies for ψ—forecast horizon, analyst experience, and the qualitative versus quantitative nature of information—are associated with systematic variation in expectation biases in the directions predicted by the model.

More broadly, our results suggest that understanding expectation formation requires distinguishing between environments where realized outcomes impose strong discipline and those where forecasts are harder to evaluate ex post. This perspective helps unify disparate findings on under- and overreaction and highlights belief-distortion costs as a useful organizing parameter for cross-item and cross-context differences in analyst behavior.

References

H. Afrouzi, S. Y. Kwon, A. Landier, Y. Ma, and D. Thesmar. Overreaction in expectations: Evidence and theory. *The Quarterly Journal of Economics*, 138(3):1713–1764, 2023.

G. A. Akerlof and W. T. Dickens. The economic consequences of cognitive dissonance. *The American economic review*, 72(3):307–319, 1982.

S. Banerjee, J. Davis, and N. Gondhi. Choosing to disagree: Endogenous dismissiveness and overconfidence in financial markets. *The Journal of Finance*, 79(2):1635–1695, 2024.

S. Banerjee, J. Davis, and N. Gondhi. Strategic wishful thinking: Implications for forecasts. *Journal of Political Economy*, 2026.

N. Barberis, R. Greenwood, L. Jin, and A. Shleifer. X-CAPM: An extrapolative capital asset pricing model. *Journal of financial economics*, 115(1):1–24, 2015.

R. Bénabou and J. Tirole. Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–64, 2016.

P. Bordalo, N. Gennaioli, and A. Shleifer. Diagnostic expectations and credit cycles. *The Journal of Finance*, 73(1):199–227, 2018.

P. Bordalo, N. Gennaioli, R. L. Porta, and A. Shleifer. Diagnostic expectations and stock returns. *The Journal of Finance*, 74(6):2839–2874, 2019.

P. Bordalo, N. Gennaioli, Y. Ma, and A. Shleifer. Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–2782, 2020.

P. Bordalo, N. Gennaioli, and A. Shleifer. Memory, attention, and choice. *Quarterly Journal of Economics*, 136(3):1969–2011, 2021.

J.-P. Bouchaud, P. Krueger, A. Landier, and D. Thesmar. Sticky expectations and the profitability anomaly. *The Journal of Finance*, 74(2):639–674, 2019.

M. K. Brunnermeier and J. A. Parker. Optimal expectations. *American Economic Review*, 95(4):1092–1118, 2005.

M. K. Brunnermeier, C. Gollier, and J. A. Parker. Optimal beliefs, asset prices, and the preference for skewed returns. *American Economic Review*, 97(2):159–165, 2007.

A. Caplin and J. Leahy. Wishful thinking. 2019.

S. Cassella, B. Golez, H. Gulen, and P. Kelly. Motivated beliefs in macroeconomic expectations. *Available at SSRN 3759035*, 2023.

O. Coibion and Y. Gorodnichenko. What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159, 2012.

O. Coibion and Y. Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. *The American Economic Review*, pages 2644–2678, 2015.

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

B. S. da Silveira and M. Woodford. Optimally imprecise memory and choice. *Econometrica*, 88(4):1397–1445, 2020.

B. S. da Silveira and M. Woodford. Noisy memory and overreaction to news. *American Economic Review*, 112(12):3967–3998, 2022.

K. Daniel, D. Hirshleifer, and A. Subrahmanyam. Investor psychology and security market under-and overreactions. *the Journal of Finance*, 53(6):1839–1885, 1998.

D. d'Arienzo. Maturity increasing overreaction and bond market puzzles. *Available at SSRN 3733056*, 2020.

K. Diether, C. Malloy, and A. Scherbina. Differences of opinion and the cross section of stock returns. *Journal of Finance*, 57(5):2113–2141, 2002.

B. Enke, T. Graeber, and R. Oprea. Cognitive uncertainty and the structure of beliefs. *American Economic Review*, 110(10):3105–3130, 2020.

L. Gemmi and R. Valchev. Biased surveys. *Journal of Monetary Economics*, 157, 2026.

S. Ke. Analysts' belief formation in their own words. *Available at SSRN 5025830*, 2024.

A. N. Kohlhas and A. Walther. Asymmetric attention. *American Economic Review*, 111(9): 2879–2925, 2021.

G. Loewenstein. Anticipation and the valuation of delayed consumption. *The Economic Journal*, 97(387):666–684, 1987.

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

N. G. Mankiw and R. Reis. Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, 117(4):1295–1328, 2002.

I. Marinovic, M. Ottaviani, and P. N. Sørensen. Modeling idea markets: Between beauty contests and prediction markets. 2010.

I. Marinovic, M. Ottaviani, and P. Sorensen. Forecasters' objectives and strategies. In *Handbook of economic forecasting*, volume 2, pages 690–720. Elsevier, 2013.

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

D. P. Myatt and C. Wallace. Endogenous information acquisition in coordination games. *The Review of Economic Studies*, 79(1):340–374, 2012.

S. Nagel and Z. Xu. Fading memory and asset pricing. *The Review of Financial Studies*, 35 (5):2190–2245, 2022.

M. Ottaviani and P. N. Sørensen. The strategy of professional forecasting. *Journal of Financial Economics*, 81(2):441–466, 2006.

Y. Sung. Macroeconomic expectations and cognitive noise. Federal Reserve Bank of San Francisco, 2024.

M. Woodford. Imperfect common knowledge and the effects of monetary policy. In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Princeton University Press, 2003.

A Proofs

A.1 Proof of Lemma 1

Under rational expectations, $A_i = A$ and $B_i = B$ for all i, where

$$A = \frac{\tau_e}{\tau + \tau_e + \tau_\eta}, \qquad B = \frac{\tau_\eta}{\tau + \tau_e + \tau_\eta}.$$

Given information (s_i, s), player i solves

$$\max_{k_i} \mathbb{E}_i\big[-(k_i - \theta)^2 - \rho(k_i - K)^2 \mid s_i, s\big].$$

The first-order condition is

$$(1+\rho)k_i = \mathbb{E}_i[\theta \mid s_i, s] + \rho\, \mathbb{E}_i[K \mid s_i, s],$$

or, with $r = \rho/(1+\rho)$,

$$k_i = (1-r)\mathbb{E}_i[\theta \mid s_i, s] + r\, \mathbb{E}_i[K \mid s_i, s].$$

Conjecture a symmetric linear equilibrium

$$k_i = \alpha s_i + \beta s, \qquad K = \alpha\theta + \beta s.$$

Since $\mathbb{E}_i[\theta \mid s_i, s] = As_i + Bs$, the conjectured aggregate action $K = \alpha\theta + \beta s$ implies

$$\mathbb{E}_i[K \mid s_i, s] = \alpha\, \mathbb{E}_i[\theta \mid s_i, s] + \beta s = \alpha(As_i + Bs) + \beta s.$$

Hence player i's best response is

$$\begin{aligned} k_i &= (1-r)(As_i + Bs) + r\big[\alpha(As_i + Bs) + \beta s\big] \\ &= \big(1 - r + r\alpha\big)As_i + \big[\big(1 - r + r\alpha\big)B + r\beta\big]s. \end{aligned}$$

Matching coefficients on (s_i, s) with $k_i = \alpha s_i + \beta s$ gives

$$\alpha = A(1 - r + r\alpha), \qquad \beta = B(1 - r + r\alpha) + r\beta.$$

From the first equation:

$$\alpha(1 - rA) = A(1 - r) \implies \alpha = \frac{A(1-r)}{1 - rA}.$$

Using this in the second equation:

$$\beta(1-r) = B(1 - r + r\alpha) = B\frac{1-r}{1-rA} \implies \beta = \frac{B}{1-rA}.$$

Therefore a linear equilibrium exists with

$$k_i = \alpha s_i + \beta s, \qquad K = \alpha\theta + \beta s,$$

and the coefficients in the lemma.

Uniqueness follows because the two equilibrium conditions are linear and $1 - rA > 0$ (since $A \in (0,1)$ and $r < 1$ for $\rho \in (-1, 1)$), so the solution for (α, β) is unique. □

A.2 Proof of Proposition 1

Using Lemma 1,

$$k_i = \alpha s_i + \beta s, \qquad K = \alpha\theta + \beta s, \qquad s_i = \theta + \varepsilon_i,\ s = \theta + \eta,$$

with $\varepsilon_i \perp \eta \perp \theta$, $\mathrm{var}(\theta) = 1/\tau$, $\mathrm{var}(\varepsilon_i) = 1/\tau_e$, and $\mathrm{var}(\eta) = 1/\tau_\eta$. Define

$$FR_i = k_i, \qquad FE_i = \theta - k_i, \qquad \bar{FR} = K, \qquad \bar{FE} = \theta - K.$$

Then

$$FE_i = (1-\alpha-\beta)\theta - \alpha\varepsilon_i - \beta\eta, \quad FR_i = (\alpha+\beta)\theta + \alpha\varepsilon_i + \beta\eta.$$

So

$$\operatorname{cov}(FE_i, FR_i) = \frac{(1-\alpha-\beta)(\alpha+\beta)}{\tau} - \frac{\alpha^2}{\tau_e} - \frac{\beta^2}{\tau_\eta},$$

$$\operatorname{var}(FR_i) = \frac{(\alpha+\beta)^2}{\tau} + \frac{\alpha^2}{\tau_e} + \frac{\beta^2}{\tau_\eta} > 0.$$

Similarly,

$$\operatorname{cov}(\bar{FE}, \bar{FR}) = \frac{(1-\alpha-\beta)(\alpha+\beta)}{\tau} - \frac{\beta^2}{\tau_\eta},$$

$$\operatorname{var}(\bar{FR}) = \frac{(\alpha+\beta)^2}{\tau} + \frac{\beta^2}{\tau_\eta} > 0.$$

Under rational expectations, set $D := \tau + \tau_e + \tau_\eta$, $A = \tau_e/D$, $B = \tau_\eta/D$. With $\alpha = \frac{(1-r)A}{1-rA}$ and $\beta = \frac{B}{1-rA}$, note that

$$\alpha + \beta = \frac{(1-r)\tau_e + \tau_\eta}{D(1-rA)}, \qquad 1-\alpha-\beta = \frac{\tau}{D(1-rA)}.$$

Substituting into $\operatorname{cov}(FE_i, FR_i)$:

$$\begin{aligned}
\operatorname{cov}(FE_i, FR_i) &= \frac{1}{(1-rA)^2}\left[\frac{\tau\big[(1-r)\tau_e + \tau_\eta\big]}{D^2\tau} - \frac{(1-r)^2\tau_e^2}{D^2\tau_e} - \frac{\tau_\eta^2}{D^2\tau_\eta}\right] \\
&= \frac{1}{(1-rA)^2 D^2}\Big[(1-r)\tau_e + \tau_\eta - (1-r)^2\tau_e - \tau_\eta\Big] \\
&= \frac{r(1-r)\tau_e}{(1-rA)^2 D^2}.
\end{aligned}$$

Because $\tau_e > 0$, $(1-r) > 0$, and $(1-rA)^2D^2 > 0$, the sign of $\operatorname{cov}(FE_i, FR_i)$ equals the sign of $r = \rho/(1+\rho)$, hence the sign of ρ. Therefore

$$CG_i < 0 \iff \rho < 0.$$

For consensus,

$$\operatorname{cov}(\bar{FE}, \bar{FR}) = \frac{1}{(1-rA)^2 D^2}\Big[(1-r)\tau_e + \tau_\eta - \tau_\eta\Big] = \frac{(1-r)\tau_e}{(1-rA)^2 D^2} > 0,$$

so $CG_a > 0$. This proves Proposition 1. □

A.3 Proof of Lemma 2

Fix any profile $\{\delta_{e,i}, \delta_{\eta,i}\}_i$, and thus fixed

$$A_i = \frac{\delta_{e,i}\tau_e}{\tau + \delta_{e,i}\tau_e + \delta_{\eta,i}\tau_\eta}, \qquad B_i = \frac{\delta_{\eta,i}\tau_\eta}{\tau + \delta_{e,i}\tau_e + \delta_{\eta,i}\tau_\eta},$$

with aggregates $A = \int_i A_i\, di$ and $B = \int_i B_i\, di$. Given (s_i, s), the best response satisfies

$$k_i = (1-r)\mathbb{E}_i[\theta \mid s_i, s] + r\,\mathbb{E}_i[K \mid s_i, s], \qquad r = \frac{\rho}{1+\rho}.$$

Conjecture a linear equilibrium

$$k_i = \alpha_i s_i + \beta_i s, \qquad K = \alpha\theta + \beta s,$$

where $\alpha = \int_i \alpha_i di$ and $\beta = \int_i \beta_i di$.

Since $\mathbb{E}_i[\theta \mid s_i, s] = A_i s_i + B_i s$, the conjectured aggregate action $K = \alpha\theta + \beta s$ implies

$$\mathbb{E}_i[K \mid s_i, s] = \alpha\, \mathbb{E}_i[\theta \mid s_i, s] + \beta s = \alpha(A_i s_i + B_i s) + \beta s.$$

Hence

$$\begin{aligned} k_i &= (1-r)(A_i s_i + B_i s) + r\big[\alpha(A_i s_i + B_i s) + \beta s\big] \\ &= (1 - r + r\alpha)A_i\, s_i + \big[(1 - r + r\alpha)B_i + r\beta\big] s. \end{aligned}$$

Matching coefficients:

$$\alpha_i = (1 - r + r\alpha)A_i, \qquad \beta_i = (1 - r + r\alpha)B_i + r\beta.$$

Integrating the first equation across i:

$$\alpha = (1 - r + r\alpha)A \implies \alpha = \frac{A(1-r)}{1 - rA}.$$

Therefore

$$1 - r + r\alpha = \frac{1-r}{1-rA},$$

so

$$\alpha_i = \frac{1-r}{1-rA} A_i.$$

Integrating the second equation:

$$\beta = (1 - r + r\alpha)B + r\beta \implies \beta = \frac{B}{1 - rA}.$$

Substituting this back:

$$\beta_i = \frac{(1-r)B_i + rB}{1 - rA}.$$

Thus a linear equilibrium exists with exactly the coefficients in the lemma.

Uniqueness follows because (α, β) solve a linear system with determinant proportional to $1 - rA > 0$ (again $A \in (0,1)$ and $r < 1$), and then (α_i, β_i) are pinned down uniquely by the two coefficient-matching equations above. □

A.4 Proof of Proposition 2

Player i chooses $(\delta_{e,i}, \delta_{\eta,i})$ to maximize (18). Taking others' beliefs as given (so B is parametric for player i),

$$TU_i(\delta_{e,i}, \delta_{\eta,i}) \propto -\frac{1}{D_i} - \frac{\rho B^2}{\delta_{\eta,i}\tau_\eta} - \psi \frac{N_i}{D_i^2},$$

where

$$D_i := \tau + \delta_{e,i}\tau_e + \delta_{\eta,i}\tau_\eta, \qquad N_i := \tau + \delta_{e,i}^2\tau_e + \delta_{\eta,i}^2\tau_\eta.$$

For an interior optimum, first-order conditions are

$$\frac{\partial TU_i}{\partial \delta_{e,i}} \propto \frac{\tau_e}{D_i^2} - 2\psi \frac{\tau_e(\delta_{e,i} D_i - N_i)}{D_i^3} = 0,$$

$$\frac{\partial TU_i}{\partial \delta_{\eta,i}} \propto \frac{\tau_\eta}{D_i^2} + \frac{\rho B^2}{\delta_{\eta,i}^2 \tau_\eta} - 2\psi \frac{\tau_\eta(\delta_{\eta,i} D_i - N_i)}{D_i^3} = 0.$$

In a symmetric equilibrium, $\delta_{e,i} = \delta_e$ and $\delta_{\eta,i} = \delta_\eta$ for all i, so

$$B = \frac{\delta_\eta \tau_\eta}{\tau + \delta_e \tau_e + \delta_\eta \tau_\eta} = \frac{\delta_\eta \tau_\eta}{D}, \qquad D := \tau + \delta_e \tau_e + \delta_\eta \tau_\eta.$$

Hence

$$\frac{\rho B^2}{\delta_\eta^2 \tau_\eta} = \frac{\rho \tau_\eta}{D^2},$$

and the two FOCs become

$$D - 2\psi(\delta_e D - N) = 0,$$

$$(1+\rho)D - 2\psi(\delta_\eta D - N) = 0,$$

where $N := \tau + \delta_e^2 \tau_e + \delta_\eta^2 \tau_\eta$. Subtracting the first equation from the second gives

$$\rho D - 2\psi(\delta_\eta - \delta_e)D = 0 \implies \delta_\eta = \delta_e + \frac{\rho}{2\psi}.$$

Substituting into the first FOC and solving for δ_e yields

$$\delta_e = \frac{2\tau\psi(2\psi+1) + (1+\rho)\rho\tau_\eta}{2\psi\left(2\tau\psi - (1+\rho)\tau_\eta - \tau_e\right)}, \qquad \delta_\eta = \delta_e + \frac{\rho}{2\psi},$$

which is exactly (20). Since $\rho > 0$ and $\psi > 0$, $\delta_\eta > \delta_e$.

Define

$$\psi^\star := \frac{\tau_e + (1+\rho)\tau_\eta}{2\tau}.$$

If $\psi > \psi^\star$, then $2\tau\psi - (1+\rho)\tau_\eta - \tau_e > 0$, so the interior solution above is well-defined and unique. Moreover, with this denominator positive,

$$\delta_e - 1 = \frac{\tau + \tau_e + (1+\rho)\tau_\eta + \frac{(1+\rho)\rho\tau_\eta}{2\psi}}{2\tau\psi - (1+\rho)\tau_\eta - \tau_e} > 0,$$

hence $\delta_e > 1$ and therefore $\delta_\eta > 1$.

If $\psi \le \psi^\star$, no finite interior maximizer satisfies both FOCs. For the constrained problem with upper bound M on belief distortions, best responses are at the upper boundary, and as $M \to \infty$ the posterior ignores the prior:

$$A_i \to \frac{\tau_e}{\tau_e + \tau_\eta}, \qquad B_i \to \frac{\tau_\eta}{\tau_e + \tau_\eta}.$$

This gives part (i) of the proposition.

Finally, comparative statics for $\psi > \psi^\star$: direct differentiation of the closed form for δ_e gives $\delta_e'(\psi) < 0$, and since $\delta_\eta = \delta_e + \rho/(2\psi)$,

$$\delta_\eta'(\psi) = \delta_e'(\psi) - \frac{\rho}{2\psi^2} < 0.$$

Thus both δ_e and δ_η are weakly decreasing in ψ (strictly decreasing on the interior region). □

A.5 Proof of Proposition 3

From Lemma 2, in a symmetric equilibrium with (δ_e, δ_η),

$$k_i = \alpha s_i + \beta s, \qquad K = \alpha\theta + \beta s,$$

where

$$\alpha = \frac{1-r}{1-rA}A, \qquad \beta = \frac{B}{1-rA}, \qquad r = \frac{\rho}{1+\rho} \in (0, 1/2),$$

$$A = \frac{\delta_e \tau_e}{D}, \qquad B = \frac{\delta_\eta \tau_\eta}{D}, \qquad D := \tau + \delta_e\tau_e + \delta_\eta\tau_\eta.$$

Using the same covariance identities as in Proposition 1,

$$\mathrm{cov}(FE_i, FR_i) = \frac{(1-\alpha-\beta)(\alpha+\beta)}{\tau} - \frac{\alpha^2}{\tau_e} - \frac{\beta^2}{\tau_\eta},$$

$$\mathrm{cov}(\bar{FE}, \bar{FR}) = \frac{(1-\alpha-\beta)(\alpha+\beta)}{\tau} - \frac{\beta^2}{\tau_\eta},$$

and both variance denominators in CG_i and CG_a are strictly positive. Substituting (A, B, α, β) gives

$$\mathrm{cov}(FE_i, FR_i) = \frac{\tau_e[(1-r)\delta_e - (1-r)^2\delta_e^2] + \tau_\eta\left[\delta_\eta - \delta_\eta^2\right]}{(1-rA)^2 D^2},$$

$$\mathrm{cov}(\bar{FE}, \bar{FR}) = \frac{\tau_e(1-r)\delta_e + \tau_\eta(\delta_\eta - \delta_\eta^2)}{(1-rA)^2 D^2}.$$

Now use Proposition 2. For low ψ, agents ignore priors, so $A + B = 1$ with $A = \frac{\tau_e}{\tau_e+\tau_\eta}$ and $B = \frac{\tau_\eta}{\tau_e+\tau_\eta}$. Hence $1 - \alpha - \beta = 0$, implying

$$\mathrm{cov}(FE_i, FR_i) = -\frac{\alpha^2}{\tau_e} - \frac{\beta^2}{\tau_\eta} < 0, \qquad \mathrm{cov}(\bar{FE}, \bar{FR}) = -\frac{\beta^2}{\tau_\eta} < 0.$$

So for sufficiently small ψ, both individual and consensus forecasts overreact.

For high ψ, Proposition 2 implies $(\delta_e, \delta_\eta) \to (1, 1)$, i.e., the rational-expectations benchmark. By Proposition 1 (with $\rho > 0$), in that limit

$$\mathrm{cov}(FE_i, FR_i) > 0, \qquad \mathrm{cov}(\bar{FE}, \bar{FR}) > 0,$$

hence both CG_i and CG_a are positive for sufficiently high ψ.

By continuity in (δ_e, δ_η) and in ψ, sign changes occur at threshold values of ψ. Also,

$$\mathrm{cov}(\bar{FE}, \bar{FR}) - \mathrm{cov}(FE_i, FR_i) = \frac{\alpha^2}{\tau_e} > 0.$$

Therefore, whenever $CG_i > 0$ we must also have $CG_a > 0$; equivalently, the consensus covariance crosses zero weakly before the individual one. Hence only three sign regions are possible, ordered as ψ increases:

1. low ψ: $CG_i < 0$ and $CG_a < 0$;
2. intermediate ψ: $CG_i < 0 < CG_a$;
3. high ψ: $CG_i > 0$ and $CG_a > 0$.

Finally, because Proposition 2 gives that both distortions decrease with ψ, and the two covariance numerators above move continuously from negative (low ψ) to positive (high ψ), both CG_i and CG_a are weakly increasing in ψ. This proves Proposition 3. □

A.6 Proof of Proposition 4

In a symmetric equilibrium with experienced-utility penalty, both distortions are chosen endogenously. From Proposition 2,

$$\delta_e(\psi) = \frac{2\tau\psi(2\psi+1) + (1+\rho)\rho\tau_\eta}{2\psi\left(2\tau\psi - (1+\rho)\tau_\eta - \tau_e\right)}, \tag{22}$$

$$\delta_\eta(\psi) = \delta_e(\psi) + \frac{\rho}{2\psi}, \tag{23}$$

with admissible region $\psi > \psi^* := \frac{(1+\rho)\tau_\eta+\tau_e}{2\tau}$, so

$$z := 2\tau\psi - (1+\rho)\tau_\eta - \tau_e > 0.$$

By Lemma 2,

$$k_i = \alpha s_i + \beta s, \qquad K = \alpha\theta + \beta s,$$

and, under symmetry,

$$\alpha = \frac{A}{1+\rho-\rho A}, \qquad \beta = \frac{(1+\rho)B}{1+\rho-\rho A},$$

with $A = \frac{\delta_e\tau_e}{\tau+\delta_e\tau_e+\delta_\eta\tau_\eta}$ and $B = \frac{\delta_\eta\tau_\eta}{\tau+\delta_e\tau_e+\delta_\eta\tau_\eta}$. Equivalently,

$$\alpha = \frac{\delta_e\tau_e}{(1+\rho)\tau + (1+\rho)\delta_\eta\tau_\eta + \delta_e\tau_e}, \qquad \beta = \frac{(1+\rho)\delta_\eta\tau_\eta}{(1+\rho)\tau + (1+\rho)\delta_\eta\tau_\eta + \delta_e\tau_e}.$$

1) Dispersion in actions decreases with ψ. Cross-sectional dispersion is

$$\sigma_k^2 := \int (k_i - K)^2 di = \int (\alpha\varepsilon_i)^2 di = \frac{\alpha^2}{\tau_e}.$$

Thus it suffices to sign $d\alpha/d\psi$. Direct differentiation of the closed form above gives

$$\frac{d\alpha}{d\psi} = -\frac{\tau_e(1+\rho)P_\alpha(\psi)}{2\psi^2\tau\left(2\psi\rho\tau + 2\psi\rho\tau_\eta + 2\psi\tau + 2\psi\tau_e + 2\psi\tau_\eta - \rho\tau_e\right)^2},$$

where

$$\begin{aligned} P_\alpha(\psi) = {} & 4\psi^2\tau^2 + 4\psi^2\tau\tau_e + 4\psi^2\tau\tau_\eta + 4\psi\rho^2\tau\tau_\eta + 4\psi\rho^2\tau_\eta^2 \\ & + 4\psi\rho\tau\tau_\eta + 4\psi\rho\tau_e\tau_\eta + 4\psi\rho\tau_\eta^2 - \rho^2\tau_e\tau_\eta. \end{aligned}$$

Set $\psi = \frac{z+(1+\rho)\tau_\eta+\tau_e}{2\tau}$. After substitution, P_α becomes a polynomial in z with strictly positive coefficients; since $z > 0$, $P_\alpha(\psi) > 0$. Hence $\frac{d\alpha}{d\psi} < 0$, and therefore

$$\frac{d\sigma_k^2}{d\psi} = \frac{2\alpha}{\tau_e}\frac{d\alpha}{d\psi} < 0.$$

2) Variance of consensus forecasts decreases with ψ. Define

$$\nu^2 := \mathbb{E}[K^2] = \mathrm{var}(K) = \frac{(\alpha+\beta)^2}{\tau} + \frac{\beta^2}{\tau_\eta},$$

since $K = (\alpha + \beta)\theta + \beta\eta$ and $\theta \perp \eta$. It is enough to show $\frac{d(\alpha+\beta)}{d\psi} < 0$ and $\frac{d\beta}{d\psi} < 0$. From the closed forms, direct differentiation yields

$$\frac{d(\alpha+\beta)}{d\psi} = -\frac{2(1+\rho)C}{\left(2\psi\rho\tau + 2\psi\rho\tau_\eta + 2\psi\tau + 2\psi\tau_e + 2\psi\tau_\eta - \rho\tau_e\right)^2} < 0,$$

where

$$C = \rho^2\tau\tau_\eta + \rho^2\tau_\eta^2 + 2\rho\tau\tau_\eta + 2\rho\tau_e\tau_\eta + 2\rho\tau_\eta^2 + \tau\tau_e + \tau\tau_\eta + \tau_e^2 + 2\tau_e\tau_\eta + \tau_\eta^2 > 0,$$

and

$$\frac{d\beta}{d\psi} = -\frac{\tau_\eta(1+\rho)P_\beta(\psi)}{2\psi^2\tau\left(2\psi\rho\tau + 2\psi\rho\tau_\eta + 2\psi\tau + 2\psi\tau_e + 2\psi\tau_\eta - \rho\tau_e\right)^2},$$

with $P_\beta(\psi)$ transforming (under the same z-substitution) into a polynomial in z with strictly positive coefficients; hence $P_\beta(\psi) > 0$ and $\frac{d\beta}{d\psi} < 0$. Therefore

$$\frac{d\nu^2}{d\psi} = \frac{2(\alpha+\beta)}{\tau}\frac{d(\alpha+\beta)}{d\psi} + \frac{2\beta}{\tau_\eta}\frac{d\beta}{d\psi} < 0,$$

because $\alpha + \beta > 0$ and $\beta > 0$.

Thus, with both distortions endogenous, dispersion in actions and variance of consensus forecasts both decrease with the cost of belief distortion ψ. □

B Additional Tables

Table B.1: Consensus EPS CG Coefficient (1989–2015)

This table compares the consensus-level CG coefficient from Bouchaud et al. (2019) (Table III, Panel A, Column 1) with our estimate using the same sample period (fiscal years 1989–2015). Both columns report pooled OLS regressions of consensus forecast errors on consensus forecast revisions, scaled by lagged stock price, with standard errors clustered by firm and year. Bouchaud et al. (2019) define consensus as the median forecast; we use the mean. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)
Forecast Revision	0.165***	0.166***
	(0.016)	(0.016)
Observations	54,090	38,351
R-squared	0.030	0.042
Cluster	Firm-Year	Firm-Year