Heterogeneous Overreaction in Expectation Formation: Evidence and Theory[*]

HENG CHEN
University of Hong Kong
XU LI
University of Hong Kong
GUANGYU PEI
Chinese University of Hong Kong
QIAN XIN
Harbin Institute of Technology

December 15, 2023

Abstract. Using firm-level earnings forecasts and managerial guidance data, we construct guidance surprises for analysts, i.e., differences between managerial guidance and analysts' initial forecasts. We document new evidence on expectation formation: (i) analysts overreact to managerial guidance and the overreaction is state-dependent, i.e., it is stronger for negative guidance surprises but weaker for surprises that are larger in size; and (ii) forecast revisions are neither symmetric in guidance surprises nor monotonic. We organize these facts with a model where analysts are uncertain about the quality of managerial guidance. We show that a reasonable degree of ambiguity aversion is necessary to account for the documented heterogeneous overreaction pattern.

Keywords. overreaction, expectation formation, managerial guidance, forecast revision, asymmetry, nonmonotone, ambiguity aversion

JEL Classification. C53, D83, D84

---

[*]Heng Chen: The Faculty of Business and Economics, University of Hong Kong; Xu Li: The Faculty of Business and Economics, University of Hong Kong; Guangyu Pei: Department of Economics, Chinese University of Hong Kong; and Qian Xin: School of Economics and Management, Harbin Institute of Technology, Shenzhen. This research project is partly funded by the Seed Fund for Basic Research of Hong Kong University (Project No. 202011159014).

1. Introduction

The mechanisms underlying expectation formation are crucial for understanding economic decisions. While it is documented that individuals in general overreact to information (Bordalo, Gennaioli, Ma, and Shleifer 2020), there has been growing interest in the circumstances under which the overreaction is stronger or weaker. In this paper, we provide new evidence that the degree of overreaction can be heterogeneous across individual forecasters, even when they receive the same information. To organize the facts, we propose a forecasting model where agents make forecasts based on noisy information and are uncertain about information quality.

To test how agents form expectations in general and how they react to new information in particular, it would be ideal to have a testing ground in which (i) the new information acquired by agents is observable and measurable, and (ii) agents' forecasts before and after receiving the new information are available. We consider an environment that is fairly close to this: financial analysts forecast the earnings of firms, firms release managerial guidance for earnings, and then analysts update their earnings forecasts. Forecast revisions are then defined to be the differences between analysts' updated forecasts after receiving managerial guidance and their initial forecasts before receiving it. That is, forecast revisions are constructed to reflect the impact of the guidance on earnings.

Using earnings forecasts data (individual analysts' EPS forecasts from the I/B/E/S Estimates) and managerial guidance data (the I/B/E/S Guidance data) from 1994 to 2017, we provide a number of findings. First, analysts' forecasts overreact to information that arrives during the time window that is constructed to encompass managerial guidance. We show that forecast revisions are negatively correlated with forecast errors, which are defined to be the differences between realized earnings and analysts' updated forecasts. This suggests that upward (downward) revisions can predict negative (positive) forecast errors, i.e., there is too much revision relative to the rational benchmark. This result is consistent with the existing findings of Bordalo, Gennaioli, Ma, and Shleifer (2020) using macroeconomic survey data.

Second, our new finding in this paper is that the overreaction is heterogeneous across analysts. We define *guidance surprises* to be the differences between the managerial guidance and analysts' initial forecasts. We construct surprises at the firm-quarter-analyst level, rank those surprises from the most negative to the most positive and then group them into deciles. Estimating the degree of overreaction in each decile subsample, we find that overreaction is stronger when surprises are negative; overreaction tends to be weaker when surprises are larger in size.

Third, we further directly explore how forecast revisions respond to guidance surprises with nonparametric estimations. We find that forecast revisions are asymmetric in surprises: forecast revisions are stronger when the surprises are negative than those when the surprises are of the same magnitude but positive. Furthermore, forecast revisions are not monotonically increasing in surprises either: when the surprises are large enough, forecast revisions decrease in surprises. Thus, the estimated relationship between forecast revisions and surprises displays a pattern of asymmetry and non-monotonicity. It is worth pointing out that the two new facts corroborate with each other.[1]

The new evidence on the documented heterogeneous overreaction pattern calls for a new theory, in which optimal response to new information have to be *state-dependent*. We consider a forecasting model where analysts would receive managerial guidance for earnings from the firm and update their forecasts in response. The key departures from standard forecasting models are (a) that analysts are ambiguous about the quality of the managerial guidance and (b) that they are ambiguity averse and the degree of ambiguity aversion is finite. The former requires that analysts should update their beliefs about the quality of guidance based on the guidance itself and then update their beliefs about earnings for any possible quality. The latter implies that analysts wish to act in a robust fashion.

In this model, the extent to which analysts overreact (or even underreact) to information while revising their forecasts depends critically on how analysts perceive the quality of managerial guidance. Specifically, analysts behave as if, in their posterior beliefs, they optimally overweigh the states of the world where their expected utility is low. When surprises are negative, analysts would subjectively "overcount" the quality of guidance, which leads to a more pronounced overreaction. In addition, when surprises are sufficiently large in size, analysts would infer that the quality of guidance is less likely to be high (the standard Bayesian mechanism), which leads to a more moderated overreaction (or potentially an under-reaction). Both model mechanisms are consistent with the pattern of heterogeneous overreaction found in the data.

In Section 4, we demonstrate that it is crucial to allow agents to possess a *finite* degree of ambiguity aversion to simultaneously capture both nonmonotonicity and asymmetry in the relationship between forecast revisions and surprises. Without ambiguity aversion, analysts' forecast revisions are symmetric, despite the sign of surprises. With extreme ambiguity aversion (i.e., the Wald (1950) maxmin criterion), analysts' forecast revisions are monotonic in surprises, despite the uncertainty in information quality. Furthermore, we characterize an explicit model counterpart for the

[1] If forecast revisions are linear in surprises, then the extent of overreaction to new information cannot be heterogeneous; and if overreaction is heterogeneous in size and direction of surprises, then forecast revisions cannot be linear in surprises. This connection will be characterized in Section 4.3.

coefficients that quantify the extent of heterogenous overreaction documented in the data.

In Section 5.1, we present a quantitative rendition of our model, demonstrating that our estimated model can produce a cross-sectional overreaction pattern consistent with the data. In Section 5.2, we examine two auxiliary predictions of our model using the data, which help corroborate our model mechanisms. While our study is the first to discover and rationalize this set of facts, there might be other mechanisms contributing to the documented patterns. To underscore our theoretical contributions to the literature, we compare our model with several existing theories in Section 5.3, including diagnostic beliefs, overconfidence, loss aversion, and agency theory.

Both the facts documented and the mechanisms characterized in this paper are relevant for the expectation formation literature in general and studies concerning overreactions to information in particular. The empirical part of this paper builds on a new literature that empirically explores information frictions and expectation formation (Coibion and Gorodnichenko 2015). Using macroeconomic survey data, Bordalo, Gennaioli, Ma, and Shleifer (2020) and Broer and Kohlhas (2022) find that forecasters overreact to information in general.[2] In an experimental setting, Afrouzi, Kwon, Landier, Ma, and Thesmar (2022) establish that the overreaction is stronger for a less persistent data generation process and stronger for longer forecast horizons.

In contrast, we document the heterogeneous overreaction among analysts, taking a step beyond the existing literature. Additionally, we develop a complementary empirical approach that directly investigates the relationship between forecast revisions and observable new information, which can prove to be a valuable tool for the literature. It's worth highlighting that we establish a novel empirical setting for studying expectation formation, which holds significance for other related research in this field.[3]

Our new theory adds to the literature of expectation formation by explicitly scrutinizing how forecasters react to noisy data of uncertain quality. Both Epstein and Schneider (2008) and Baqaee (2020) characterize the process of expectation formation when agents have an extreme ambiguity-averse preference (i.e., multiple priors) and show that belief updating is asymmetric in the contexts of asset pricing and business

[2]Other recent studies also provide evidence on the forecasts of financial market participants, such as Bordalo, Gennaioli, Porta, and Shleifer (2019), Bouchaud, Krueger, Landier, and Thesmar (2019), Amromin and Sharpe (2014), Barrero (2022), Ma, Ropele, Sraer, and Thesmar (2020) and Greenwood and Shleifer (2014). Farmer, Nakamura, and Steinsson (2021) study a dynamic environment in which slow learning over the unit root long-run trend can rationalize a set of forecasting anomalies at the consensus level. Binder, Kuang, and Tang (2023), and Kuang, Mitra, Tang, and Xie (2023) use survey experiments to study the effects of economic policies on the forecasts of financial variables.

[3]We use managerial guidance to facilitate the exploration because this is among the very few kinds of information that are observable, measurable and systematically accessible to econometricians. Management earnings guidance is one of the most significant events that releases new information to the market during a quarter.

cycles, respectively. In contrast, our work allows for a finite degree of aversion in the smooth model of ambiguity following Klibanoff, Marinacci, and Mukerji (2005) and Cerreia-Vioglio, Maccheroni, and Marinacci (2022). Focusing on ambiguity about the second moments of the data generating process, our model offers theoretical predictions that are qualitatively different from the aforementioned works and that are also empirically relevant.

In general, there is a growing interest in understanding how agents' use of information deviates from the rational expectation benchmark. Prominent examples include diagnostic expectations (Bordalo, Gennaioli, and Shleifer 2018, Bianchi, Ilut, and Saijo 2022), overconfidence (Broer and Kohlhas 2022), cognitive discounting (Gabaix 2020), level-K thinking (García-Schmidt and Woodford 2019, Farhi and Werning 2019), narrow thinking (Lian 2020), adaptive learning (Adam, Kuang, and Marcet 2012, Kuang and Mitra 2016), autocorrelation averaging (Wang 2020) and loss aversion (Elliott, Komunjer, and Timmermann 2008, Capistrán and Timmermann 2009). A common feature of those models in a Gaussian environment is that forecast revisions are increasing in surprises and the direction of surprises does not matter. Our model differs in both aspects.[4]

2. Evidence

2.1. Data, Sample and Timing

In this section, we explore how analysts revise their earnings forecasts upon newly received information. Our goal is to construct a scenario where the information flow is observable, measurable and accessible to the econometrician.

Toward this end, we focus on managerial guidance, which is among the very few information sources that satisfy such criteria. In financial markets, the management teams of publicly listed firms issue guidance for the earnings of the current quarter between the last quarter's and current quarter's earnings announcements. That is a crucial opportunity for firms to provide information about earnings to market participants, such as financial analysts. Because of its importance, earnings guidance often triggers analysts' forecast updates: analysts likely revise their forecasts a few days after receiving earnings guidance, i.e., on average 4 days in our sample (constructed in this section).[5] Furthermore, it is common that firms continue to provide earnings

[4] State-dependent forecasting behavior can be a consequence of strategic information provision. Nimark and Pitschner (2019) demonstrate that asymmetry in forecasting may arise when an information provider slants negative news, making it more salient when reported. In contrast, in our model, asymmetry arises due to ambiguity in information quality and ambiguity aversion, without strategic behavior from information providers.

[5] On average, analysts publish their initial forecasts 43 days before earnings guidance becomes available to the market.

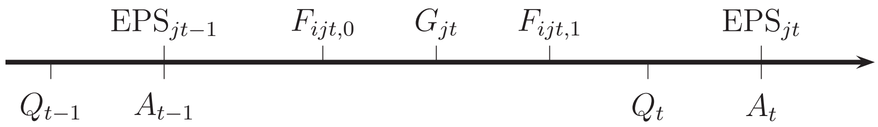

Figure 1. *Timeline. We consider managerial guidance G_t issued between A_{t-1} and A_t. If the guidance for EPS in quarter t is released on the date of A_{t-1} or within two days after A_{t-1}, then it is bundled. If the guidance is released between Q_t and A_t, it is a preannouncement. If more than one guidance is released between A_{t-1} and A_t, we choose the latest one.*

guidance for an extensive period of time, and the discontinuation in earnings guidance is typically perceived unfavorably by the market (Chen, Matsumoto, and Rajgopal 2011). Earnings guidance includes various forms, such as point estimates and range estimates.

The Thomson Reuters I/B/E/S Guidance data provides quantitative managerial expectations, such as earnings per share, from press releases and transcripts of corporate events. The data cover managerial guidance from more than 6,000 companies in North America that can date back to as early as 1994. Furthermore, the I/B/E/S Guidance data are available on the same accounting basis as the I/B/E/S Estimates that provide individual analysts' forecast data. This makes it feasible to rigorously identify the timing of events and to compare managerial guidance and analysts' forecasts for the same firm in a certain period. Our sample construction based on the I/B/E/S Guidance and Estimates data is elaborated and relegated to Appendix A.1.

We stress that we intentionally construct a time window where analysts' initial and updated forecasts encompass the earnings guidance of the current quarter. This construction allows us to analyze how forecasts are updated in response to information observable to the econometrician. The construction procedure can be better apprehended with the aid of Figure 1, which delineates the sequence of major events. Analyst i learns firm j's EPS for quarter $t-1$ at the date of A_{t-1}, which is $\mathrm{EPS}_{j,t-1}$. Then he or she issues a forecast $F_{ijt,0}$ for firm j's EPS in quarter t. Firm j offers guidance G_{jt} for firm j's earnings in quarter t. Then, analyst i updates his or her forecast for firm j's EPS in quarter t (i.e., $F_{ijt,1}$). Quarter t ends at the date of Q_t, and firm j announces its EPS for quarter t at the date of A_t. In sum, in this setting, both initial and updated forecasts are made within the same period, after A_{t-1} and before A_t.

Our full sample consists of 110,895 pairs of individual analysts' forecasts (initial and updated forecasts) issued by 6,987 different analysts for 3,226 district firms over the period from 1994 to 2017. A summary of statistics is reported in Appendix A.2.

Table 1. Forecast Error on Forecast Revision

	Outcome Variable: Forecast Error FE_i					
	Winsorization at the 1% and 99%			Winsorization at the 2.5% and 97.5%		
	Baseline	Control	Unscaled	Baseline	Control	Unscaled
	(1)	(2)	(3)	(4)	(5)	(6)
FR_i	-0.0961***	-0.0963***	-0.0924***	-0.0914***	-0.0914***	-0.0736***
	(0.0142)	(0.0143)	(0.0124)	(0.0117)	(0.0117)	(0.0101)
Earnings of the Last Quarter		0.0036			0.0010	
		(0.0070)			(0.0050)	
Firm FEs	YES	YES	YES	YES	YES	YES
Obs.	110,895	110,895	110,895	110,895	110,895	110,895
Adj. R-sq	0.2125	0.2125	0.1796	0.1947	0.1947	0.1812

The standard errors are clustered on firm and calendar year-quarter following Petersen (2009).*** p<0.01, ** p<0.05, * p<0.1

2.2. Overreaction

Our investigation of how analysts revise their forecasts starts by following the approach proposed by Bordalo, Gennaioli, Ma, and Shleifer (2020), in which they examine professional analysts' forecasts of macro variables. That is, we regress ex post analyst forecast errors on ex ante analyst forecast revisions at the individual level. To this end, we construct both forecast error FE_{ijt} and forecast revision FR_{ijt}. The former is the difference between the realized earnings per share for firm j in quarter t and the revised EPS forecast by individual analyst i for firm j in quarter t. The latter is the difference between the revised forecast after guidance and the initial forecast before guidance issued by the same analyst i for firm j in quarter t. To avoid the heterogeneity embedded in EPS across firms, both FE_{ijt} and FR_{ijt} are scaled by the stock price at the beginning of quarter t. To mitigate the impact of potential outliers, both of them are winsorized at the 1% and 99% level of their respective distributions. We estimate the following equation:

$$FE_{ijt} = b_0 + b_1 FR_{ijt} + \delta_j + \omega_{ijt}, \tag{1}$$

where we control for firm fixed effect (δ_j) to absorb time-invariant firm characteristics. Following Petersen (2009), the standard errors are clustered at the firm and calendar year-quarter to adjust for both inter-temporal and cross-sectional correlations.

The results from estimating Equation (1) are presented in column (1) of Table 1. We find that forecast errors are negatively correlated with forecast revisions at the individual analyst level and statistically significant at less than the 1% level. The negative coefficient indicates that analysts overreact to new information over the period that the managerial guidance is received by analysts. Despite the settings being entirely different, this result is consistent with those found in Bordalo, Gennaioli, Ma, and Shleifer (2020) and Broer and Kohlhas (2022).

We add the earnings in the last quarter ($t-1$) of firm j to the right-hand side of Equation (1) and report the results in column (2) of Table 1. The change in the estimated coefficient on forecast revision is negligible, and the coefficient on the earnings in the last quarter is close to zero and not significant. This suggests that the information about earnings in past quarters is fully utilized by analysts to form either initial or updated forecasts. That is the key difference from studies using SPF data, where initial and updated forecasts are made in two separate periods.

To ensure that our results are robust to data construction, we present results by not scaling earnings and forecasts by stock prices. The estimate for forecast revisions is robust, which is reported in column (3). To test whether our results are driven by outliers, we winsorize FE_{ijt} and FR_{ijt} at the 2.5% and 97.5% levels of their respective distributions and re-do the aforementioned exercises. Those results are reported in columns (4)-(6) of Table 1, which demonstrate the robustness of our findings.[6]

Two comments on the specification of Equation (1) are in order. First, incorporating the firm fixed effect into this regression is crucial for identifying the average overreaction. If there is a systematic bias in EPS forecasts that varies across firms, the coefficient (b_1) would be biased. This is because Equation (1) is estimated by pooling the firms. Second, to address concerns about the potential Nickel bias arising from the inclusion of the firm fixed effect, we conduct separate estimations of Equation (1) for each firm, following the approach proposed by Bordalo, Gennaioli, Ma, and Shleifer (2020). The median estimation of b_1 is then reported in Table A4 of Appendix B. The median coefficient is similar to that reported in Table 1.

2.3. Heterogeneous Overreaction

One unique feature of our setting is that the guidance is common for all analysts, but the surprises contained in the guidance are not common across analysts due to their heterogeneous initial forecasts. Analysts can be surprised to different extents and even in different directions. One natural question arises: Do analysts overreact differently to the same information? In this section we explore such heterogeneity of overreaction across analysts.

First, we construct a variable *guidance surprise* (i.e., Surprise_{ijt}) to capture the *observable* surprise in managerial guidance for individual analysts. It is defined and measured by the difference between the value of guidance (i.e., G_{jt}) issued by firm

[6] Appendix B.1 reports additional robustness tests demonstrating that our results remain robust with different sample selection and trimming the outliers instead of winsorizing. The estimated coefficients in the aforementioned exercises are qualitatively unchanged and only different in magnitude.

j in quarter t and analyst i's corresponding initial forecast (i.e., F_{0ijt}) for firm j in quarter t before guidance. That is, $\text{Surprise}_{ijt} \equiv G_{jt} - F_{0ijt}$. For each individual analyst, the managerial guidance can be *unfavorable* or *favorable* if it falls below or exceeds the analyst's initial forecast before guidance, and the managerial guidance can be *large* or *small* if it is far from or close to the analyst's initial forecast before guidance.

Second, we remove outliers by trimming forecast errors, forecast revisions and surprises at the 2.5% and 97.5% levels of their respective distributions (to be consistent with the nonparametric estimations in the next section). We then rank surprises from the most negative to the most positive, sort them into deciles and label them from 1 to 10 according to the decile rank. To enlarge the subsample size and smooth estimates, we define a running decile window j such that (1) window j covers decile $j-1$, j, and $j+1$ if $j \neq 1$ or $j \neq 10$; (2) running decile window 1 covers deciles 1 and 2; and (3) running decile window 10 covers deciles 9 and 10.

Third, for each subsample of a running decile window, we re-estimate Equation (1) (i.e., regressing forecast errors on forecast revisions). We plot the estimated coefficients and confidence intervals in Figure 2 against their window ranks. We find that analysts overreact to information in each subsample, i.e., the estimated coefficient b_1 is negative and significant. However, the degree of overreaction is not constant and is U-shaped in surprises and skewed to the left. This implies that the overreaction is stronger when the surprises are negative and the overreaction is weaker when the surprises are larger in size.[7]

In summary, on the one hand, we confirm that analysts overreact to information in this particular setting. Given that the forecast revisions are constructed around managerial guidance, analysts are likely to overreact to guidance surprises. On the other hand, we discover that the way that analysts react to information depends on the characteristics of the surprises that they receive, such as the size and direction of the surprises.

2.4. Forecast Revisions and Surprises: Mechanisms

In this section, we set out to uncover the mechanisms that underlie the heterogeneous overreaction pattern. To this end, we directly investigate the relationship between forecast revisions and surprises. Note that if forecast revisions are linear in surprises, then the degree of overreaction to new information cannot be heterogeneous (charac-

[7]To examine whether our results are robust, we rerun the exercises with a sample where forecast errors, forecast revisions and surprises are trimmed at the 1% and 99% levels of their respective distributions. We also re-estimate Equation (1) for each decile of surprises without using running windows. The patterns found are rather similar. We relegate them to Appendix B.2 (see Figures A1 and A2, respectively).

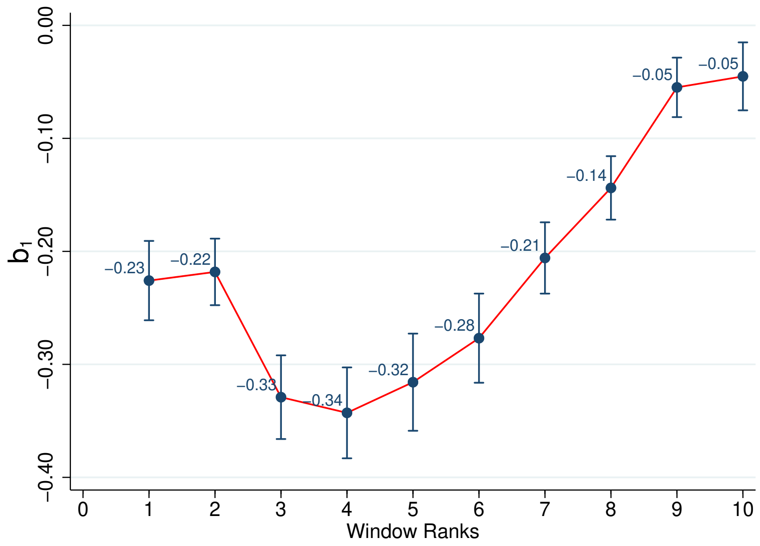

Figure 2. *Heterogeneous Overreaction. The estimated coefficients of the FE-on-FR regressions b_1 and the 95% confidence interval for each running decile window are plotted against the window rank. Running decile window j covers decile $j-1$, j, and $j+1$ if $j \neq 1$ or $j \neq 9$; running decile window 1 covers deciles 1 and 2, and running decile window 10 covers deciles 9 and 10.*

terized in Section 4.3); and if overreaction is heterogeneous in the size and direction of surprises, then forecast revisions cannot be linear in surprises.

To estimate the relationship in a more reliable fashion, we resort to the nonparametric estimation approach. Using the standard tool of local polynomial regression, we estimate the relationship between forecast revisions and surprises by using the Epanechnikov kernel and the third degree of the smoothing polynomial.

Because we are interested in "large" surprises and because we estimate the relationship with local polynomials, the results can be affected and biased by winsorization of the data. To alleviate this concern, we instead trim both forecast revisions and surprises at the 2.5% and 97.5% levels of their respective distributions and residualize them by controlling for time, analyst, and firm fixed effects. We estimate their relationship using the local polynomial specification, and the results are presented in Figure 3(a). Forecast revisions are decreasing, increasing and decreasing in surprises and are asymmetric around the origin. Figure 3(b) illustrates its derivatives with respect to surprises. The derivatives are negative when the surprises are large enough and positive when they are small. Forecast revisions respond more strongly to negative surprises than to positive surprises of the same magnitude. In Appendix B.3, we present a range of robustness checks, and the empirical findings are robust.

To quantify the degree of asymmetry in the estimated relationship, we compute the percentage deviations of forecast revisions to negative surprises (i.e., $\text{Surp}_{ijt} < 0$)

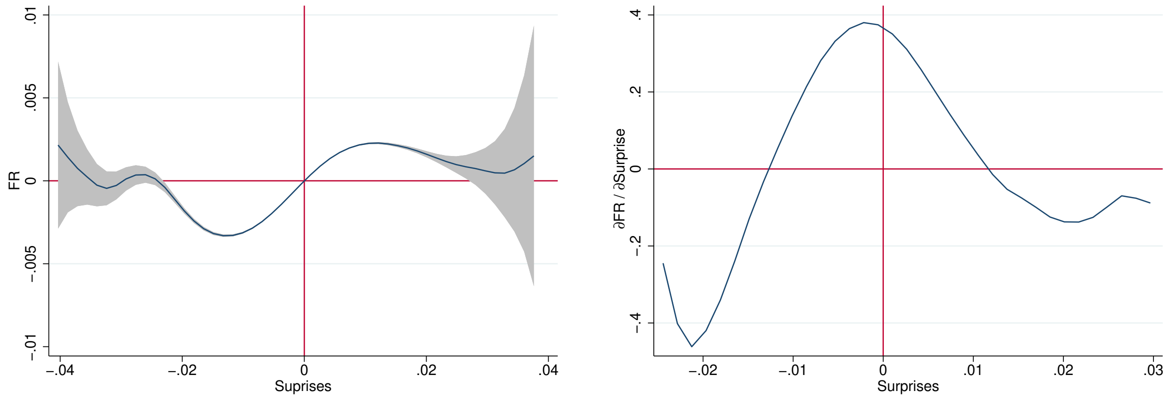

(a) Trimming, Nonparametric estimation

(b) Trimming, Derivative: marginal effect

Figure 3. *Nonparametric estimation, 5% trimming (2.5%, 97.5%). Panel (a) illustrates the relationship between forecast revisions and surprises in managerial guidances (both trimmed at 5%) that is nonparametrically estimated using the Epanechnikov kernel and the third degree of the smoothing polynomial. It is decreasing, increasing and decreasing and asymmetric around the origin. The shaded areas represent the 95% confidence intervals for the respective estimations. Panel (b) illustrates its derivatives with respect to surprises for the range where the non-parametric estimation and the numerical derivative are relatively precise, i.e., when surprises are between* $[-0.025, 0.030]$. *The derivatives are negative when the surprises are large enough and positive when they are small. Forecast revisions respond more strongly to negative surprises than to positive surprises of the same magnitude.*

from forecast revisions to positive ones of the same magnitude (i.e., $-\text{Surp}_{ijt} < 0$) and construct an average conditional on surprises are negative. That is,

$$\Xi \equiv \int_{-\infty}^{0} \frac{\left|\text{FR}\left(\text{Surp}_{ijt}\right)\right| - \left|\text{FR}\left(-\text{Surp}_{ijt}\right)\right|}{\left|\text{FR}\left(-\text{Surp}_{ijt}\right)\right|} \text{dP}\left(\text{Surp}_{ijt} | \text{Surp}_{ijt} < 0\right), \tag{2}$$

where $\text{P}\left(\text{Surp}_{ijt} | \text{Surp}_{ijt} < 0\right)$ is the conditional distribution of surprise that can be directly inferred from the data. The asymmetry measure Ξ is positive if, on average, negative surprises (i.e., $\text{Surp}_{ijt} < 0$) result in larger forecast revisions compared to positive ones. If Ξ is zero, the response to forecast revisions to surprises are symmetric.

In our baseline sample, we find that $\Xi = 0.18$, indicating that, on average, negative surprises result in revisions that are approximately 18% stronger than revisions triggered by positive surprises of similar magnitude.

The facts documented in sections 2.3 and 2.4 would be puzzling if one assumed that analysts know the quality of managerial guidance with certainty. In such a case, forecast revisions would be linear in surprises within the Gaussian environment, and the degree of overreaction would also be constant. Once we relax this assumption and accommodate the conjecture that the quality of information can be uncertain to analysts, those documented facts can be reasonable and consistent with each other. To account for those facts in a unifying framework, we propose a model where analysts are uncertain about the quality of information that they receive.

3. The Model

3.1. Setup

Consider a one-period static model where there exists a continuum of analysts, indexed by $i \in [0,1]$, and a firm. The firm's earnings θ are stochastic. Analyst i makes a forecast F_{0i} about the earnings at the beginning of the period and makes an updated forecast F_i at the end of the period.

Utility function. In the context of forecasting problems, we impose one restriction that analysts' optimal forecast is precisely $F^* = \theta$, conditional on analysts' information being complete (i.e., the earnings θ are known to the analysts). Any utility functions that satisfy this restriction can be approximated by a utility function $U(\cdot,\cdot)$ that is quadratic in both forecasts and earnings. In the main text, we consider one particular case among this class of quadratic utility functions, which is given by:

$$U(F,\theta) = -(F-\theta)^2 + \beta\theta, \tag{3}$$

where β is a constant. To interpret parameter β, consider the scenario where analysts have complete information. They can minimize the forecasting errors to zero, but the realized earnings may still matter for analysts in our model. The parameter $\beta > 0$ ($\beta < 0$) implies that analysts would be better (worse) off if the realized earnings θ were higher. The parameter β will be estimated and interpreted in Appendix C.[8]

This utility function is used for ease of exposition and highlighting our new mechanisms. In Technical Appendix, we present a full characterization of the model with the most general quadratic utility function of this class. We show that it is qualitatively similar and provide evidence that the additional parameters in the general case are empirically irrelevant in this setting.

Information structure. We assume that the earnings follow a normal distribution with mean 0 and variance σ_θ^2, i.e., $\theta \sim N\left(0,\sigma_\theta^2\right)$; let $\tau_\theta = 1/\sigma_\theta^2$. The distribution of earnings is known to all analysts. To have a direct mapping with the data, we allow each analyst i to be endowed with private information about the earnings before making the initial forecasts, as follows:

$$z_{0i} = \theta + \iota_i,$$

[8] Appendix C provides discussions on empirical evidence that analysts' utility can be dependent on earnings. In this section, we provide a characterization in which β can take any value.

where ι_i is normally distributed with mean 0 and variance σ_z^2, i.e., $\iota_i \sim N\left(0, \sigma_z^2\right)$; let $\tau_z = 1/\sigma_z^2$. Analyst i makes forecast F_{0i} with heterogeneous information z_{0i}.

Analysts then receive managerial guidance released by the firm, which is a noisy signal about earnings:

$$y = \theta + \eta.$$

where η is normally distributed with mean 0 and variance σ_Y^2, i.e., $\eta \sim N\left(0, \sigma_Y^2\right)$; let $\tau_Y = 1/\sigma_Y^2$. After analysts have made their updated forecasts, the earnings announcement is made, and the payoffs to analysts are realized.

The information structure in this model warrants discussion. First, in this paper, we focus on a static model without modeling the dynamics of earnings across periods. As discussed in Section 2.1, analysts have perfect information about earnings in the last quarter. Both the initial and updated forecasts in the data are made after the earnings in the last quarter are known to analysts. In this case, forecasts of the last period's earnings are not relevant in this period, conditional on the last quarter's earnings themselves.[9] Note that the updated earnings forecasts of the last period are *not* the initial forecasts for earnings in this period. Second, for simplicity, we assume that unobservable private information (such as new information from analysts' research or acquired from other sources) is absent between the two rounds of forecasts. In Technical Appendix, we fully characterize a generalized model by allowing the presence of private information and show that all the qualitative properties remain.

Ambiguity-averse preferences. The key departure of this model from the existing forecasting literature is that we assume that analysts are uncertain or ambiguous about the quality of the managerial guidance or their objective precision (i.e., τ_Y). Therefore, they have to form their own subjective belief about its precision (i.e., τ_y). Such an assumption is reasonable. Analysts may not know the quality of the guidance with complete certainty because management has incentives not to release the best possible information at hand and because even the best possible estimates from the management can be plagued with noise but analysts are not certain about its structure.

Specifically, we let Γ_y be the range of support for the possible precision τ_y of managerial guidance. Analysts believe that $\tau_y \in \Gamma_y$ and possess some prior belief over Γ_y, whose density distribution is given by $p\left(\tau_y\right)$. We say that one particular τ_y represents a *model* that generates the managerial guidance y.

Furthermore, we assume that analysts dislike uncertainty in the quality of the managerial guidance or are ambiguity averse. In this model, we capture such a preference of analysts by using the *smooth model of ambiguity* as proposed in Klibanoff, Marinacci,

[9]In fact, we show in Section 2.2 that earnings in the last quarter cannot predict forecast errors in the current quarter conditional on forecast revisions and are orthogonal to forecast revisions in the data.

and Mukerji (2005). That is, analyst i maximizes the objective function:

$$\int_{\Gamma_y} \phi\left(\mathbb{E}^{\tau_y}\left[U\left(F_i, \theta\right) | z_{0i}, y\right]\right) p\left(\tau_y | z_{0i}, y\right) \mathrm{d}\tau_y, \tag{4}$$

where $\phi(\cdot)$ is some increasing, concave and twice continuously differentiable function. In addition, $\mathbb{E}^{\tau_y}\left[U\left(F_i, \theta\right) | z_{0i}, y\right]$ denotes the mathematical expectation conditional on analyst i's information set (z_{0i}, y) for a particular model τ_y (or a certain precision of managerial guidance). In what follows, we use $\mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]$ to denote the expected utility of analyst i, unless it causes confusion. The density of the posterior belief over possible models is assumed to be Bayesian and denoted by $p\left(\tau_y | z_{0i}, y\right)$.

The curvature of function $\phi(\cdot)$ captures an aversion to mean-preserving spreads in $\mathbb{E}_i^{\tau_y}$ induced by ambiguity in τ_y.[10] The more concave the function $\phi(\cdot)$ is, the stronger the ambiguity aversion. In other words, it characterizes analysts' taste for ambiguity. In this paper, we consider a function $\phi(\cdot)$ that features constant absolute ambiguity aversion (CAAA) following Cerreia-Vioglio, Maccheroni, and Marinacci (2022) throughout:

$$\phi(t) = -\frac{1}{\lambda} e^{-\lambda t}, \tag{5}$$

where $\lambda \geq 0$ measures the degree of ambiguity aversion. Two special cases are nested. When $\lambda = 0$ and $\phi(\cdot)$ is linear, this corresponds to the case where analysts are ambiguity neutral or fully Bayesian. When $\lambda \to +\infty$, this corresponds to the case where analysts' aversion to ambiguity is infinite, which is the classic Wald (1950) maxmin criterion.[11]

3.2. Noisy Information Expectations: RE Benchmark

Our framework is a generalized version of the standard forecasting problem in which analysts possess noisy information and minimize the mean-squared error of their forecasts of the random variable. In other words, the noisy information benchmark is a special case of our model when agents are ambiguity neutral (i.e., $\lambda = 0$) and there exists no uncertainty in information quality (i.e., Γ_y is singleton).[12] In this section, we characterize such a special case and illustrate why it fails to account for the empirical patterns documented in Section 2.3 and 2.4 and why deviations from this benchmark are necessary.

[10] Ambiguity aversion differs from risk aversion, which is implicitly captured by $U\left(F_i, \theta\right)$. In this model, it is the aversion to ambiguity rather than the aversion to risk that drives our results.

[11] The model with extreme ambiguity aversion is a special case of the multiple priors preference proposed by Gilboa and Schmeidler (1989), where the priori set of priors include all Dirac measures of each model.

[12] In the noisy information benchmark, the parameter β in Equation (A12) plays no role at all. However, it is important for the optimal forecasts when agents have ambiguity averse preferences.

With noisy information expectations, the optimal initial and updated forecasts are such that

$$\mathrm{F}_{0i}^{\mathrm{NI}} = \mathbb{E}\left[\theta|z_{0i}\right]; \qquad \mathrm{F}_{i}^{\mathrm{NI}} = \mathbb{E}\left[\theta|z_{0i}, y\right],$$

where $\mathbb{E}\left[\theta|\mathcal{I}_i\right]$ denotes the conditional expectations (i.e., Bayesian posterior). The relationship between $\mathrm{F}_{0i}^{\mathrm{NI}}$ and $\mathrm{F}_{i}^{\mathrm{NI}}$ is therefore given by:

$$\mathrm{F}_{i}^{\mathrm{NI}} = \left(1 - \kappa_y\right) \mathrm{F}_{0i}^{\mathrm{NI}} + \kappa^{\mathrm{RE}} y,$$

where κ^{RE} is the relevant weight assigned to the public information:

$$\kappa^{\mathrm{RE}} \equiv \frac{\tau_Y}{\tau_\theta + \tau_z + \tau_Y} > 0. \tag{6}$$

Therefore, the relevant forecast revision is given by

$$\mathrm{FR}_{i}^{\mathrm{NI}} \equiv \mathrm{F}_{i}^{\mathrm{NI}} - \mathrm{F}_{0i}^{\mathrm{NI}} = \kappa^{\mathrm{RE}} \left(y_i - \mathrm{F}_{0i}^{\mathrm{NI}}\right), \tag{7}$$

and forecast error is given by

$$\mathrm{FE}_{i}^{\mathrm{NI}} \equiv \theta - \mathrm{F}_{i}^{\mathrm{NI}} = \kappa_\theta \theta - \kappa_z \iota_i - \kappa^{\mathrm{RE}} \eta, \tag{8}$$

where $\kappa_\theta \equiv \frac{\tau_\theta}{\tau_\theta + \tau_z + \tau_Y} > 0$ and $\kappa_z \equiv \frac{\tau_z}{\tau_\theta + \tau_z + \tau_Y} > 0$.

Lemma 1 (FR-on-Surprise and FE-on-FR)**.** *In the noisy expectation benchmark, forecast revisions are linear in guidance surprises and uncorrelated with forecast errors,*

$$Cov\left(\mathrm{FE}_{i}^{\mathrm{NI}}, \mathrm{FR}_{i}^{\mathrm{NI}}\right) = 0.$$

Observe that the term $(y - \mathrm{F}_{0i}^{\mathrm{NI}})$ in Equation (7) is the theory counterpart of managerial guidance surprises in our empirical exercise. Equation (7) predicts that forecast revisions should be linear in guidance surprises. However, this prediction contradicts the non-monotone and asymmetric relationship documented in Section 2.4.

Further, using Equations (7) and (8), it is evident that forecast revisions and forecast errors are uncorrelated. It then predicts that the estimated coefficient in the FE-on-FR regression should be 0, i.e., no over-reaction at the individual level. This prediction contradicts evidence that analysts overreact to new information (documented in Section 2.2) and that such overreaction varies in a non-monotonic and asymmetric fashion (documented in Section 2.3).

The key to the failure that the noisy information benchmark cannot capture the empirical patterns, is that the optimal forecasting rule is state-independent and determined by constant signal-to-noise ratios. That is, the weight κ^{RE} assigned to the public signal (i.e., managerial guidance in this context) is constant and independent of the realization of the public signal. However, evidence suggests that the weight should vary depending on the realization of public signal in a particular way: the weight should be larger when the surprise is negative than when it is positive but of the same magnitude; and the weight should be negative (instead of positive) when surprises are large enough. In the following section, we demonstrate that our framework, featuring the ambiguous information quality and ambiguity aversion towards uncertainty, can generate a state-dependent forecasting rule that is consistent with data.

3.3. Equilibrium Characterization

In this section, we turn to the characterization of analysts' optimal forecasts. The initial forecast of each analyst F_{0i}^* is derived by Bayes' rule:

$$F_{0i}^* = \frac{\tau_z}{\tau_z + \tau_\theta} z_{0i}. \tag{9}$$

To choose the optimal updated forecast F_i^* after obtaining a new set of information, analysts maximize the objective in Equation (4). That is, the optimal forecast F_i^* is such that the first-order condition holds:

$$F_i = \int_{\Gamma_y} \left(\frac{\tau_z z_{0i} + \tau_y y}{\tau_\theta + \tau_z + \tau_y} \right) \tilde{p}\left(\tau_y | z_{0i}, y; F_i\right) \mathrm{d}\tau_y, \tag{10}$$

where the distorted posterior belief $\tilde{p}$ is such that

$$\tilde{p}\left(\tau_y | z_{0i}, y; F_i\right) \propto \underbrace{\phi'\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]\right)}_{\text{Pessimistic Distortion}} \underbrace{p\left(z_{0i}, y | \tau_y\right) p\left(\tau_y\right)}_{\text{Bayesian Kernel}}. \tag{11}$$

The term with the combined fraction in Equation (10) captures the posterior mean of the random variable θ for a particular model τ_y, where the weights assigned to observations (z_{0i}, y) are dictated by Bayes' rule.

The distribution of τ_y is updated by following equation (11). When analysts are ambiguity neutral (i.e., $\lambda = 0$), $\phi'(\cdot)$ is constant and the posterior distribution of τ_y simply follows Bayes' rule. When analysts are ambiguity averse (i.e., $\lambda > 0$), the posterior distribution of τ_y is distorted by their pessimistic attitude: its density is reweighted by the term $\phi'\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]\right)$.

To understand such pessimism, consider analyst i who obtains observations (z_{0i}, y) and contemplates releasing a forecast F_i. She views model τ_y as the more likely model if she is worse off under such a model. That is, a model with τ_y that generates a lower expected utility for analyst i is given a higher weight in her distorted posterior belief. Recall that $\phi'(\cdot) > 0$ and $\phi''(\cdot) < 0$. Consequently, the posterior belief $\tilde{p}\left(\tau_y|z_{0i}, y; F_i\right)$ depends on her forecast F_i. Such a dependence is the key difference from the standard forecasting problems.

To facilitate the subsequent analysis and characterize the pessimism, define the *surprise* of managerial guidance y for analyst i by $s_i \equiv y - F_{0i}^*$, i.e., the difference between the guidance y and the analyst's initial forecast F_{0i}^*. The optimality condition of Equation (10) is represented by:

$$F_i = F_{0i}^* + \kappa\left(F_{0i}^*, s_i, F_i\right) \cdot s_i, \tag{12}$$

where

$$\kappa\left(F_{0i}^*, s_i, F_i\right) \equiv \left[\int_{\Gamma_y} \left(\frac{\tau_y}{\tau_\theta + \tau_z + \tau_y}\right) \tilde{p}\left(\tau_y|F_{0i}^*, s_i; F_i\right) \mathrm{d}\tau_y\right], \tag{13}$$

and the distorted posterior belief is such that

$$\tilde{p}\left(\tau_y|F_{0i}^*, s_i; F_i\right) \equiv \tilde{p}\left(\tau_y|z_{0i}, F_{0i}^*, s_i + F_{0i}^*; F_i\right). \tag{14}$$

For any particular model τ_y, the optimal response to the surprise s_i is $\frac{\tau_y}{\tau_\theta+\tau_z+\tau_y}$, which is dictated by Bayes' rule and increasing in τ_y (the quality of managerial guidance). The response to the surprise (represented by κ) is a weighted average over the model space by using the distorted distribution $\tilde{p}\left(\tau_y|F_{0i}^*, s_i; F_i\right)$, and therefore it is bounded between 0 and 1. In this representation, the pessimistic preference of analysts is specifically captured by the following lemma.

Lemma 2 (Pessimism)**.** *Consider any $F_i' > F_i$ and the likelihood ratio*

$$L\left(\tau_y\right) \equiv \frac{\tilde{p}\left(\tau_y|F_{0i}^*, s_i; F_i'\right)}{\tilde{p}\left(\tau_y|F_{0i}^*, s_i; F_i\right)}.$$

If the surprise s_i is positive, $L\left(\tau_y\right)$ decreases in τ_y; if it is negative, $L\left(\tau_y\right)$ increases in τ_y.

All proofs are collected in Appendix E. Suppose that the surprise s_i is positive. An analyst i who contemplates a higher forecast F_i' would consider the positive surprise to be less likely to be informative and assign a lower probability density for models with a high τ_y in her distorted belief $\tilde{p}$. Therefore, κ is decreasing in F_i. In contrast, suppose that the surprise s_i is negative. An analyst i who contemplates a higher forecast would consider the negative surprise to be more likely to be informative and therefore assign a higher probability density to models with high τ_y in her distorted belief. Therefore,

κ is increasing in F_i.

As implied by Lemma 2, the right-hand side of Equation (12) always decreases in F_i. The optimal forecast F_i^* is the fixed point of Equation (12). The following proposition summarizes the equilibrium existence and uniqueness of the forecasting problem.

Proposition 1 (Existence and Uniqueness)**.** *If analysts are ambiguity averse ($\lambda > 0$), there always exists a unique optimal forecast $F_i^*\left(F_{0i}^*, s_i\right)$ that satisfies (12) and a unique optimal response $\kappa^*\left(s_i\right) \equiv \kappa\left(F_{0i}^*, s_i, F_i^*\right)$ associated with it.*

An interesting special case is nested in this framework: if analysts are ambiguity neutral, there is no dependence of analyst i's posterior belief $\tilde{p}$ on F_i. Bayes' rule dictates that the posterior distribution of τ_y only depends on the magnitude of the surprise, but not its sign. Therefore, the response to surprises in managerial guidance should always be symmetric.

4. Equilibrium Analysis

This section presents a set of equilibrium analyses corresponding to the empirical facts documented in Section 2. We demonstrate that the two basic model mechanisms (uncertainty in quality and aversion to uncertainty) and their interaction can help account for those empirical patterns.

4.1. Asymmetry

We first characterize the impacts of ambiguity aversion on analysts' asymmetric responses to negative and positive surprises in managerial guidance. To state this formally, let a pair of surprises be $\left(s_i^-, s_i^+\right)$, such that $s_i^- < 0 < s_i^+$ and $s_i^- + s_i^+ = 0$.

Proposition 2. *If analysts are ambiguity averse, forecast revisions in response to surprises are asymmetric. Specifically,*

$$\left(\kappa^*\left(s_i^-\right) - \kappa^*\left(s_i^+\right)\right)\beta \geq 0,$$

where the equality holds if and only if $\beta = 0$.

To illustrate this, consider the case where analysts are better off when the earnings realization is high (i.e., $\beta > 0$). That is, analysts consider the news that suggests higher realizations of earnings to be favorable.

Proposition 2 states that analysts would always be less responsive to positive surprises (i.e., s_i^+, favorable news) than to negative surprises (i.e., s_i^-, unfavorable news). The mechanism is as follows. In this model, analysts are uncertain about the quality of the information source and, therefore, need to assess its quality based on the news itself. Given that favorable news improves analyst i's expected utility, she would behave with more caution (due to her ambiguity-averse preferences) and "discount" the quality of favorable news. Conversely, given that negative surprises or unfavorable news reduce her expected utility, she would "over-count" its quality, i.e., assign a high probability density to models with high quality τ_y. Therefore, analyst i responds to negative surprises to a larger extent than to positive surprises of the same magnitude, that is, $\kappa^*\left(s_i^-\right) > \kappa^*\left(s_i^+\right)$.

4.2. Nonmonotonicity

Next, we show that the model also features a nonmonotonic relationship between forecast revisions and surprises. Two key take-away messages are as follows. First, the nonmonotonicity does not rely on ambiguity aversion but instead on ambiguity (uncertainty) in quality. Second, in fact, the nonmonotonicity disappears when the degree of ambiguity aversion becomes extreme. Proposition 3 formalizes the former, and proposition 4 characterizes the latter. To simultaneously capture both nonmonotonicity and asymmetry, neither ambiguity-neutral preferences nor extreme ambiguity aversion is feasible.

Proposition 3. *If analysts are ambiguity neutral ($\lambda = 0$), the optimal forecast revision $F_i^* - F_{0i}^*$ increases in s_i conditional on surprise s_i being small in magnitude and decreases in s_i conditional on surprise s_i being sufficiently large in magnitude. The forecast revision at the individual level $F_i^* - F_{0i}^*$ is always symmetric around the origin.*

Given that the quality of guidance is uncertain, analyst i updates her belief through two mechanisms. First, for any given quality τ_y, analyst i updates her belief about the earnings upon receiving the guidance. This mechanism dictates that positive (negative) surprises raise (suppress) forecasts. Second, she also updates her belief about the distribution of quality. When the surprise is large, Bayesian analysts will assign a higher probability density to low qualities. That is, they tend to believe that large surprises are of low quality. Crudely, this is because low-quality information sources would have fatter tails and be more likely to generate large surprises. In other words, the posterior distribution of information quality given a small surprise first-order stochastically dominates the posterior distribution given a large surprise. Therefore, this mechanism implies that forecast revisions can be less responsive to surprises when they are larger.

For small enough surprises, the second mechanism (i.e., updating the distribution of quality) is less consequential, and therefore forecast revisions increase in surprises. For large enough surprises, the second mechanism dominates the first, and, as a result,

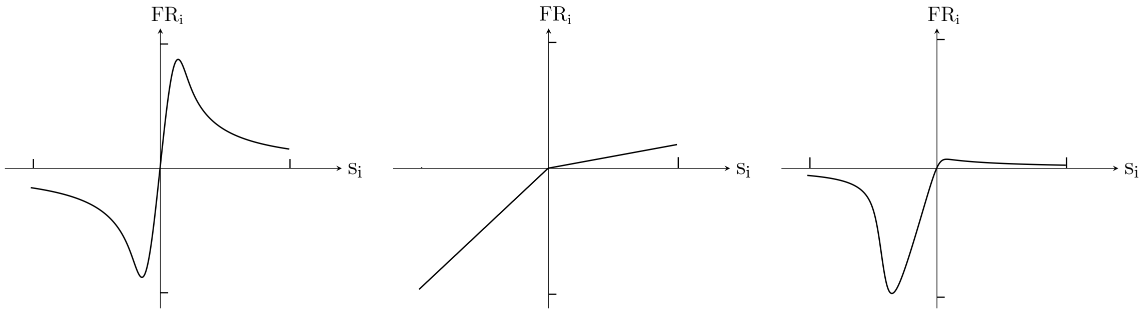

(a) Ambiguity Neutral $\lambda = 0$ (b) Max-min $\lambda \to +\infty$ (c) Ambiguity Aversion $\lambda > 0$

Figure 4. *Monotonicity and the degree of ambiguity aversion. Panel (a) illustrates the case where analysts are ambiguity neutral. Forecast revisions are decreasing, increasing and decreasing in surprises. Panel (b) illustrates the case where analysts have extreme degree of ambiguity aversion ($\lambda \to +\infty$). Note that $\beta > 0$. Forecast revisions are increasing in surprises and asymmetric. Panel (c) illustrates the case where analysts' ambiguity aversion is moderate. Both asymmetry and nonmonotonicity are present.*

forecast revisions decrease in surprises. Figure 4(a) illustrates this pattern that forecast revisions decrease and increase and then decrease in surprises. The symmetry is trivial given that analysts are Bayesian.

Now, we turn to the other polar cases: extreme ambiguity aversion ($\lambda \to +\infty$) or the classic max-min criterion.

Proposition 4. *If analysts have extreme degree of ambiguity aversion ($\lambda \to +\infty$), the optimal forecast revision $F_i^* - F_{0i}^*$ is increasing in surprise s_i.*

When surprises are relatively small in magnitude, the Bayesian mechanism dictates that forecast revisions increase in surprises (Proposition 3). Furthermore, the ambiguity aversion mechanism also dictates an increasing relationship. Analyst i tends to believe that negative surprises are of higher (lower) quality than positive surprises of the same magnitude if $\beta > 0$ ($\beta < 0$). Given that the ambiguity aversion is extreme, analysts believe that the quality of negative news is of the highest possible value and that of positive news is of the lowest possible value if $\beta > 0$ and vice versa. Figure 4(b) illustrates the case where $\beta > 0$ and $\lambda \to +\infty$. In this case, analyst i believes that negative surprises are of the highest quality and positive surprises are of the lowest quality. Therefore, forecast revisions increase in surprises with a flatter slope when surprises are positive and with a steeper slope when surprises are negative.

When surprises are very large in magnitude, the Bayesian mechanism dictates that forecast revisions decrease in surprises (Proposition 3). However, this is dominated by the impact of extreme ambiguity aversion. Therefore, forecast revisions always increase in surprises, despite the sign of β.

In summary, the contrast of the two polar cases reveals (i) that ambiguity in guidance quality gives rise to non-monotonicity in surprises and (ii) that aversion to such ambiguity leads to asymmetric responses to negative and positive surprises. Our model of finite ambiguity aversion lies in between. Figure 4(c) illustrates the relationship between forecast revisions and surprises when the degree of ambiguity aversion is moderate. The optimal forecast revision is not monotonically increasing, which resembles the case of ambiguity neutrality. Nevertheless, it is also asymmetric, which resembles the case of extreme ambiguity aversion.

4.3. Heterogeneous Overreaction: Theoretical Counterpart

The preceding two sub-sections characterize how forecast revisions respond to surprises in our model and demonstrate its consistency with the data. In this section, we offer a direct theoretical counterpart for the cross-sectional heterogeneous overreaction pattern documented in Section 2.3.

We begin our investigation by constructing the FE-on-FR coefficients in Equation (1) in the neighborhood of a particular surprise level. This construction is the theoretical counterpart of empirical coefficients in each running decile window shown in Figure (2). It allows us to study how overreaction varies over surprise in the model and directly map those model predictions to the data. Specifically, let counterpart of the concerning coefficient be:

$$\hat{b}_1\left(s_m, \epsilon\right) \equiv \frac{\operatorname{Cov}\left(\mathrm{FE}_i, \mathrm{FR}_i \mid s_i \in \mathbb{I}\left(s_m, \epsilon\right)\right)}{\operatorname{Var}\left(\mathrm{FR}_i \mid s_i \in \mathbb{I}\left(s_m, \epsilon\right)\right)}.$$

The term $\hat{b}_1\left(s_m, \epsilon\right)$ captures the FE-on-FR coefficient on an open interval $\mathbb{I}\left(s_m, \epsilon\right)$, where s_i is its middle point s_m and the width is ϵ, that is, $\mathbb{I}\left(s_m, \epsilon\right) = \left(s_m - \epsilon, s_m + \epsilon\right)$. Observe that when $s_m = 0$ and ϵ goes to ∞, $\hat{b}_1\left(s_m, \epsilon\right)$ converges to the estimated coefficient of the canonical FE-on-FR regression that characterizes average degree of overreaction.

We can further show that for sufficiently small ϵ,

$$\hat{b}_1\left(s_m\right) \equiv \lim_{\epsilon \to 0} \hat{b}_1\left(s_m, \epsilon\right) \approx -1 + \frac{\kappa^{\mathrm{RE}}}{\underbrace{\kappa\left(s_m\right) + \kappa'\left(s_m\right) s_m}_{\text{1st-order approx. of } \partial \mathrm{FR}_i / \partial s_i \text{ at } s_i = s_m}}. \tag{15}$$

where κ^{RE} denotes the responsiveness to guidance surprise in the benchmark model with rational expectation, which is characterized by Bayes' rule (see Equation (6)). The derivation is relegated to Appendix E.

Note that the denominator on the right-hand side of Equation (15) represents the first-order approximation of the marginal effect of guidance surprise s_i on forecast revisions FR_i, evaluated at the midpoint of the interval $\mathbb{I}(s_m, \epsilon)$, specifically $s_i = s_m$, when ϵ is sufficiently small. It is important to highlight that it captures the relation between forecast revisions and surprises around the point $s_i = s_m$. In other words, Equation (15) provides a theoretical mapping between the cross-sectional distribution of FE-on-FR coefficients and the FR-on-Surprise relation.

To illustrate, consider a special case, in which forecast revisions are linear in surprises: $\mathrm{FR}_i = \kappa s_i$ with κ representing the responsiveness to guidance surprise. It is worth noting that a wide range of expectation formation theories exhibit this feature, including noisy information model, diagnostic expectation, overconfidence, and parsimonious forms of loss aversion. Linearity in expectation formation implies that

$$\hat{b}_1(s_m) = -1 + \kappa^{\mathrm{RE}}/\kappa,$$

with equality holds exactly. Analysts would overreact (or underreact) to guidance surprises if and only if the responsiveness, represented by κ, is larger (or smaller) than κ^{RE}. In other words, when forecast revisions to surprises are state-independent, the degree of overreaction (or underreaction) is shown to be homogeneous.

In our model, the optimal response $\kappa(s_i)$ is state dependent. The cross-sectional pattern of heterogeneous overreaction in fact inherits the properties of non-monotonicity and asymmetry displayed in the FR-on-Surprise relation. The following proposition summarizes the results.

Proposition 5. *Suppose analysts are ambiguity averse (i.e., $\lambda > 0$) and prefer better earnings outcomes (i.e., $\beta > 0$).*

i. Analysts' overreaction (or underreaction) to guidance surprise is asymmetric with stronger overreaction (or weaker underreaction) for negative surprises, when surprises are sufficiently small in size:

$$\lim_{s_i \to 0} \frac{\mathrm{d}\hat{b}_1(s_i)}{\mathrm{d}s_i} > 0,$$

where $\hat{b}_1(s_i)$ refers to the FE-on-FR coefficient around the neighborhood of s_i.

ii. Analysts' overreaction (or underreaction) is weaker (or stronger) when guidance surprise is extreme than when it is moderate:

$$\hat{b}_1(0) < \lim_{|s_i| \to \infty} \hat{b}_1(s_i).$$

Note that the term $\lim_{s_i \to 0} \mathrm{d}\hat{b}_1(s_i)/\mathrm{d}s_i$ reduces to zero if analysts' overreaction (or underreaction) to guidance surprise is symmetric. Furthermore, item (ii) in this proposition is a sufficient condition for the non-monotonicity pattern.

5. Discussions

In this section, we address a range of relevant issues concerning our theory. In Section 5.1, we address the issue of whether our mechanism is quantitatively relevant by structurally estimating this model and comparing it with the estimated pattern found in the data. In Section 5.2, we test auxiliary predictions derived from our model, which further corroborates the mechanism proposed. In Section 5.3, we also consider various alternative hypotheses and show that the new empirical patterns documented in this paper cannot be accounted for by existing theories.

5.1. Quantitative Analysis

While we have demonstrated that the qualitative patterns of asymmetry and non-monotonicity in our model align with those observed in the data, a question arises about the model's quantitative informativeness regarding the empirical findings. Additionally, a set of parameters characterizing the utility function and ambiguity aversion play a crucial role in determining the model's qualitative properties. However, these parameters remain unobservable.

To address these two issues, we proceed to structurally estimate this model, using the simulated method of moments to match the relationship between forecast revisions and surprises that is empirically estimated in Section 2.4.[13] Then the estimated model is interpreted and used to revisit the pattern of heterogeneous overreaction (documented in sections 2.2 and 2.3) and inform the key parameters. While Appendix C provides details of the estimation, this section summarizes the key findings.

Unobservable parameters The degree of ambiguity aversion is the key to our model, and its value is estimated to be $\lambda = 449.9$. On the one hand, it is consistent with our model prediction that neither extreme ambiguity aversion ($\lambda \to +\infty$) nor ambiguity-neutral preferences ($\lambda = 0$) would be realistic for analysts in this setting. It is an important finding that justifies the use of smooth model of ambiguity aversion. On the other hand, it is worth noting that, with the reduced form utility specification, the estimated λ is only in proportion to the the actual degree of ambiguity aversion and it can be inflated by a constant shifter in utility function.[14]

[13] In the benchmark model, for simplicity, we do not allow analysts to acquire private information after they release their initial forecasts. Technical Appendix provides full characterization of the model by allowing for private information. In this structurally estimated model, we also allow for private information.

[14] To be specific, consider a model in which the degree of ambiguity aversion is $\hat{\lambda}$ and utility function is $U(F_i, \theta) = -\chi (F_i - \theta)^2 - \chi\beta\theta$, for some positive shifter $\chi > 0$. It can be shown that our model

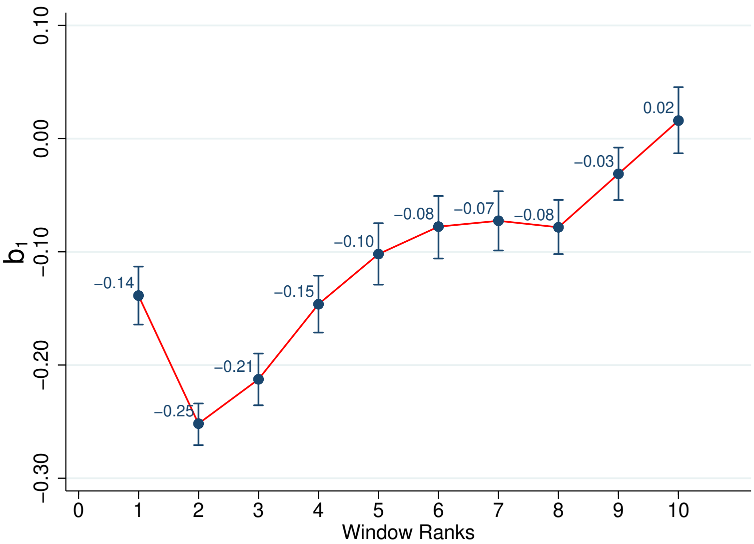

Figure 5. *Overreaction by surprise deciles with simulated data. Using simulated data, we report the estimated coefficients of the FE-on-FR regressions b_1 for each running decile window, and we plot them against the window rank. Running decile window j covers decile $j-1$, j, and $j+1$ if $j \neq 1$ or $j \neq 10$; running decile window* 1 *covers deciles* 1 *and* 2*, and running decile window* 10 *covers deciles* 9 *and* 10*.*

Furthermore, the parameter β that characterizes the utility function is estimated to be positive, i.e., $\beta = 1.37$, indicating that analysts are likely to care about the earnings performance of firms that they cover. Prior empirical studies suggest that it is plausible that β is positive. There are multiple channels through which financial analysts would benefit from better earnings performance of the firms that they cover and therefore view positive surprises in managerial guidance as favorable. For example, stronger earnings performance can be rewarding to financial analysts who make earnings forecasts and recommendations for the underlying stocks through the trading commissions channel.[15]

Quantitative relevancy To examine whether our estimated model can produce the

is isomorphic to it on condition that $\lambda = \chi\hat{\lambda}$. That is, a large χ would inflate the estimated λ in our model. The range for the estimated degree of ambiguity aversion is quite large in the literature and is sensitive to the model setup and estimation method. For example, based on asset pricing evidence, Gallant, Jahan-Parvar, and Liu (2015) estimated a degree of relative ambiguity aversion at about 66, while Collard, Mukerji, Sheppard, and Tallon (2018) calibrated the degree of ambiguity aversion to be around 12.

[15] Financial analysts aim to boost stock trading and generate trading commissions for their brokerage houses. Positive recommendations based on earnings expectations tend to increase trading volume, benefiting the analysts. Studies like Barber and Odean (2008) show that investors are inclined to follow these positive recommendations, leading to higher trading activity. Additionally, research by Groysberg, Healy, and Maber (2011) and Brown, Call, Clement, and Sharp (2015) reveals that sell-side analysts' compensation is linked to underwriting business and trading commissions of the stocks they cover directly. These analysts often focus on firms with promising earnings prospects (McNichols and O'Brien 1997; Das, Guo, and Zhang 2006), which in turn generate underwriting business and trading commissions for their brokerage houses (Alford and Berger 1999; Niehaus and Zhang 2010).

pattern of heterogeneous overreaction found in the data (in Section 2.3), we utilize the simulated data and construct the surprises observable to the econometrician in the same way as we do with the empirical data. We rank surprises from the most negative to the most positive and sort them into deciles, labelling them from 1 to 10 according to the decile rank. We further define a running decile window j, such that (1) the window j covers decile $j-1$, j, and $j+1$ if $j \neq 1$ or $j \neq 10$; (2) running decile window 1 covers deciles 1 and 2; and (3) running decile window 10 covers deciles 9 and 10. For each subsample, we re-estimate Equation (1). We plot the estimated coefficients and confidence intervals in Figure 5 against their window ranks. In the simulated data, we find that the pattern of heterogeneous overreaction is U-shaped and skewed to the left, which is consistent with our model predictions in Section 4.3 and also close to the pattern in the empirical data (Figure 2).

5.2. Auxiliary Predictions

In this paper, we provide a theory about how the expectation is formed when forecasters are not certain of the quality of the information that they receive. Our theory organizes a number of facts that we document with the earnings forecast data. In this section, we further show that our theory provides two auxiliary predictions that are consistent with the earnings forecast data.

Pessimistic bias. If the analysts in our sample are indeed ambiguity averse, then there should be a pessimistic bias in their beliefs. That is because ambiguity averse analysts react to negative guidance surprises more strongly than positive ones, since they are ambiguous about the precision of manager guidance. The revised forecasts, on average, over-represent negative guidance surprises, leading to a systematic pessimistic bias. The following proposition summarizes the result:

Proposition 6. *In this model, the optimal initial forecasts F_{i0}^* are unbiased, but the revised forecasts F_i^* are pessimistically biased, which leads to systematically positive forecast errors:*

$$\mathbb{E}\left[FE_{0i}\right] \equiv \mathbb{E}\left[\theta - F_{0i}^*\right] = 0; \qquad \mathbb{E}\left[FE_i\right] \equiv \mathbb{E}\left[\theta - F_i^*\right] > 0.$$

where $\mathbb{E}\left[\cdot\right]$ refers to the unconditional expectation with respect to the objective data generating process.

Bias in forecast errors can be obtained by regressing forecast errors on a constant and examining the estimated coefficient. To address the heterogeneity in the data generating process across firms, we run the aforementioned regression on a firm-by-firm basis and report the distribution of estimated coefficients. To ensure an adequate number of observations for each firm, we focus on a subset of firms that have provided

Table 2. Forecast Error on Constant: Median Coefficients ($\times 100$ percent)

	Firm-by-Firm Forecast Error on Constant Regressions	
	Initial Forecasts	Revised Forecasts
Median	-0.040	0.040
(p 2.5, p 97.5)	(-0.085 0.011)	(0.019 0.057)
(p 5.0, p 95.0)	(-0.067 0.06)	(0.022 0.056)
No. of firms.	786	786

We report the 5% (row 2) and 10% (row 3) bootstrapped confidence intervals, the boundaries of which are the 2.5, 5.0, 95.0, and 97.5 percentiles of the estimated median coefficients out of the 500 bootstrap samples. Following Bordalo, Gennaioli, Ma, and Shleifer (2020), these samples are obtained from block bootstrap the panel using blocks of 30 quarters.

earnings guidance for at least 12 consecutive quarters during our sample period. Table 2 presents the median estimates of forecast errors obtained from these firm-specific regressions, along with confidence intervals generated through block bootstrapping the panel data.

Interestingly, in column (1), the median estimate for initial forecasts is negative but insignificant, as indicated by the bootstrapped confidence interval. Column (2) presents the median estimate for revised forecasts, which is positive and significant. These results suggest that initial forecasts are unbiased, while revised forecasts exhibit a systematic pessimistic bias. This pattern aligns with the predictions of our theory.

Heterogeneity in quality. In this paper, our underlying assumption is that the quality of firms' earnings guidance is uncertain. However, this uncertainty likely varies across firms for analysts. Established firms with good reputations may offer high-quality managerial guidance, leading analysts to have minimal doubts about its quality. For these firms with low or no uncertainty in earnings guidance quality, our theory suggests that analysts' forecast revisions should exhibit a close-to-linear relationship with guidance surprises. In other words, the connection is expected to be both monotonic and symmetric. This is because, once the uncertainty in quality is eliminated, analysts update their beliefs solely based on the guidance and do not need to reassess the quality.

To test this prediction using our data, a conceptual challenge arises: the perceived uncertainty in guidance quality is not observable and, therefore, not measurable. To overcome this challenge, we proxy for it using the observed average quality in the data, specifically, the ex post variance of the differences between guidance and actual earnings. Our assumption is that the perceived uncertainty in quality is low if the observed average quality is high.

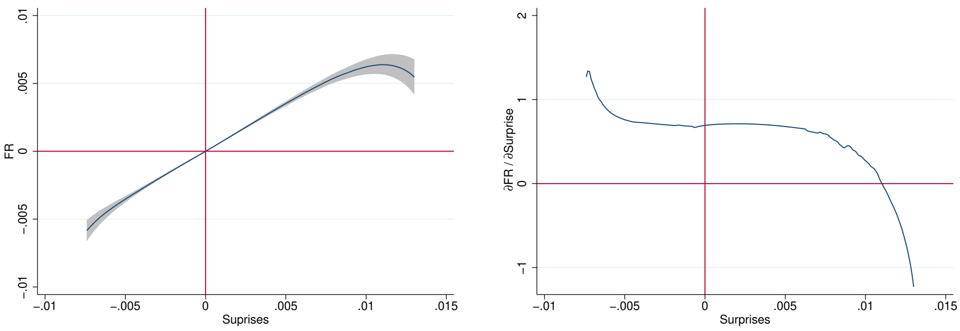

(a) Nonparametric estimation   (b) Derivative: marginal effect

Figure 6. *Nonparametric estimation using a subsample with the top 5% of firms in terms of guidance precision. Panel (a) illustrates the relationship between forecast revisions and surprises in managerial guidance that is nonparametrically estimated using the Epanechnikov kernel and the third degree of the smoothing polynomial. The shaded areas represent the 95% confidence intervals for the respective estimations. Panel (b) illustrates the derivative of forecast revisions with respect to surprises.*

We construct a subset comprising firms providing highly precise earnings guidance, indicating low uncertainty about their quality. Initially, we rank firms based on their average guidance quality within our full sample of 110,895 individual analyst forecasts, encompassing 16,241 firm-quarter observations. To be consistent with previous empirical exercises, we then trim realized earnings and management guidance at the 2.5% and 97.5% percentiles, yielding 15,427 firm-quarter observations. By regressing management guidance on the same-quarter realized earnings, controlling for year-quarter fixed effects, we obtain residuals. Firms present for fewer than 5 quarters are excluded, resulting in a reduced sample of 1,035 firms. Next, we compute standard deviations of the residuals for each firm and sort them accordingly. We concentrate on the top 5% of firms exhibiting the highest average guidance quality, forming a subsample of 2,521 individual analyst forecast revisions and guidance surprises.

Using this subsample, we re-estimate the relationship between forecast revisions and guidance surprises by following the same procedure as detailed in section 2.4. The results are shown in Figure 6(a). The relationship between forecast revisions and surprises is almost linear, unless the surprises are relatively very large and positive. The derivative estimated and shown in Figure 6(b) is close to a constant when the surprises are not too large, thus contrasting with the derivative estimated using the full sample (shown in Figure 3(b)). Additionally, the asymmetry measure Ξ, computed using Equation (2), is found to be -0.01. This value, close to zero, indicates a nearly symmetrical pattern in expectation formation when guidance quality has low or no uncertainty.

5.3. Alternative Hypotheses

In this paper, we provide a simple unified framework to account for new facts regarding how analysts update their forecasts or form expectations. It is important that our estimated model can generate the skewed U-shaped pattern of overreaction that is consistent with the data. This paper is the first that discovers and rationalizes this set of facts in the literature of expectation formation. Nevertheless, we acknowledge that there could be other mechanisms that simultaneously contribute to the observed patterns. We examine several likely candidates in sequence, which helps differentiate our theory from those in the existing literature. This section provides a summary of our investigations and the details are relegated to Appendix D.

Diagnostic expectations and over-confidence. Two related theories are commonly utilized to explain the observed overreaction patterns present in SPF data. Bordalo, Gennaioli, and Shleifer (2018) presents the theory of diagnostic expectations, a non-Bayesian model of belief formation that formalizes the concept of the representativeness heuristic (Kahneman and Tversky 1972): forecasters overweight states that are more likely in light of the arrival of new signals and consequently overreact to new information when forming their expectations. Broer and Kohlhas (2022) show that overconfidence can provide a rationalization for overreaction, i.e., forecasters subjectively believe new signals to be more accurate than they actually are. Once we allow for this set of behavioral features in the noisy information benchmark specified in Section 3.2, over-reaction to new information emerges. However, forecast revisions are still linear in surprises and the degree of over-reaction is constant and does not depend on realizations of surprises. This set of model predictions are inconsistent with the data.

Loss Aversion. Another plausible conjecture is that analysts exhibit loss aversion, instead of ambiguity aversion. To explore this possibility, we consider two widely used variants of loss aversion in the literature. Capistrán and Timmermann (2009) propose a parsimonious setup with analytical solutions, while Elliott, Komunjer, and Timmermann (2008) and Elliott and Timmermann (2008) construct a flexible setup with greater quantitative potential. In Appendix D.2, we demonstrate that regardless of the specifications for loss aversion, the FR-on-Surprise relation remains globally monotonic, whether it is linear or nonlinear. This is inconsistent with the observed non-monotonic FR-on-Surprise relation detailed in Section 2.4.

Dynamic Models. Using the Survey of Professional Forecasters (SPF), Kohlhas and Walther (2021) show that forecasters' expectations overreact to recent realizations of the output growth and therefore display a pattern of extrapolation. To explain this, they propose a model of "asymmetric attention", where Bayesian agents pay more attention to the procyclical component and less attention to the countercyclical component. Afrouzi, Kwon, Landier, Ma, and Thesmar (2022) design an experiment where participants who observe a large number of past realizations of a given AR(1) process make forecasts about future realizations. They show a pattern of overreaction, i.e., the perceived persistence of the AR(1) process is larger than the actual persistence. They propose a "top-of-mind"model, where agents rely excessively on or overreact to the recent realizations, relative to the rational benchmark.

In our empirical setting, both initial and updated forecasts are made within the same period, which encompass the earnings guidance for the current period. We use the variations of surprises contained in the earnings guidance across analysts to explore impacts of surprises' characteristics on forecast revisions. Therefore, dynamic models are not informative about the cross-sectional heterogeneity of overreaction. Appendix D.3 provides evidence for illustrating this particular finding.

Agency issues. This empirical setting is new to the literature and informative about expectation formation. However, one may worry about the role of agency issues between analysts and the managerial teams who might have incentives to misrepresent information. In the literature, it is often shown that managers spin information in self-serving ways to cater to investors and analysts (e.g. Solomon 2012). Given managerial guidance is an important information protocol provided by managers, it is reasonable to conjecture that managers have an incentive to bias their guidance positively, which makes positive managerial guidance less reliable than negative managerial guidance. This skewed information reliability, if exists, may lead to the asymmetry we documented. This conjecture is empirically testable. In Appendix D.4, we present evidence that is against the conjecture that positive guidance is less reliable. The managerial motives can be complex, often unobservable and unpredictable, which constitutes a source for guidance quality to be unreliable. In fact, that is the key motivation for our assumption that guidance quality is uncertain.

The literature also documents that managers could have incentives to manage earnings expectations downwards before the earnings release to make it beatable (e.g., Matsumoto 2002; Cotter, Tuna, and Wysocki 2006; Johnson, Kim, and So 2019). Given that one may imagine that more negatively surprised analysts could adjust their forecasts by more to ensure that the firms beat their earnings forecasts. To investigate this possibility, we rely on Johnson, Kim, and So (2019) who constructed the expectation management index (EMI) that captures the extent to which firms manage investors' earnings expectations. We add EMI as an additional control in our main regressions (reported by Table 1) and in specifications presented by Figure 2. If such a walk-down-to-beatable mechanism is crucial for our investigations, the estimated coefficients from our regressions should be greatly affected in terms of magnitude and significance. However, we find that all our estimations only change marginally at the best (available upon request), which suggests that our findings are unlikely driven only by managerial strategic guidance.

6. Conclusion

This paper documents a set of cross-sectional facts concerning expectation formation using firm-level earnings forecast and managerial guidance data: the overreaction to information is stronger for unfavorable surprises and weaker for larger surprises, and forecast revisions are asymmetric in surprises and nonmonotonic. We present a model of information uncertainty and smoothed ambiguity aversion to account for these facts. This model qualitatively differs from models with extreme ambiguity aversion or those with ambiguity-neutral agents. Our work adds to the literature that studies expectation formation by documenting new facts and providing new theory.

The empirical setting has unique advantages and will be useful for exploiting other aspects of expectation formation. First, analysts have dispersed information before receiving the guidance, summarized by their initial forecasts. The two features combined imply that the same managerial guidance delivers different surprises to analysts with different initial forecasts. The variations in surprises at the analyst level enable us to explore the cross-sectional features of overreaction. Second, in contrast to studies using the Survey of Professional Forecasters, this setting is static: we utilize within-quarter variations in surprises among analysts to uncover how analysts update their forecasts. Therefore, it is cleaner for exploring cross-sectional variations in expectation formation.

References

Adam, K., P. Kuang, and A. Marcet (2012). House price booms and the current account. *NBER Macroeconomics Annual 26*, 77–122.

Afrouzi, H., S. Kwon, A. Landier, Y. Ma, and D. Thesmar (2022). Overreaction in expectations: Evidence and theory. *Working Paper*.

Alford, A. W. and P. G. Berger (1999). A simultaneous equations analysis of forecast accuracy, analyst following, and trading volume. *Journal of Accounting, Auditing & Finance 14*(3), 219–240.

Amromin, G. and S. A. Sharpe (2014). From the horse's mouth: Economic conditions and investor expectations of risk and return. *Management Science 60*(4), 845–866.

Baqaee, D. R. (2020). Asymmetric inflation expectations, downward rigidity of wages, and asymmetric business cycles. *Journal of Monetary Economics 114*, 174–193.

Barber, B. M. and T. Odean (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *The review of financial studies 21*(2), 785–818.

Barrero, J. M. (2022). The micro and macro of managerial beliefs. *Journal of Financial Economics 143*(2), 640–667.

Bhattacharya, S. and P. Pfleiderer (1985). Delegated portfolio management. *Journal of Economic Theory 36*(1), 1–25.

Bianchi, F., C. Ilut, and H. Saijo (2022). Diagnostic business cycles. *The Review of Economic Studies (forthcoming)*.

Binder, C., P. Kuang, and L. Tang (2023). Central bank communication and house price expectations. *Working Paper*.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). Overreaction in macroeconomic expectations. *American Economic Review 110*(9), 2748–82.

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

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

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

Broer, T. and A. N. Kohlhas (2022, 06). Forecaster (Mis-)Behavior. *The Review of Economics and Statistics*, 1–45.

Brown, L. D., A. C. Call, M. B. Clement, and N. Y. Sharp (2015). Inside the "black box" of sell-side financial analysts. *Journal of Accounting Research 53*(1), 1–47.

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

Cerreia-Vioglio, S., F. Maccheroni, and M. Marinacci (2022). Ambiguity aversion and wealth effects. *Journal of Economic Theory 199*, 104898. Symposium Issue on Ambiguity, Robustness, and Model Uncertainty.

Chen, S., D. Matsumoto, and S. Rajgopal (2011). Is silence golden? an empirical analysis of firms that stop giving quarterly earnings guidance. *Journal of Accounting and Economics 51*(1-2), 134–150.

Chernozhukov, V. and H. Hong (2003). An mcmc approach to classical estimation. *Journal of Econometrics 115*(2), 293–346.

Chuk, E., D. Matsumoto, and G. S. Miller (2013). Assessing methods of identifying management forecasts: Cig vs. researcher collected. *Journal of Accounting and Economics 55*(1), 23–42.

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

Collard, F., S. Mukerji, K. Sheppard, and J.-M. Tallon (2018). Ambiguity and the historical equity premium. *Quantitative Economics 9*(2), 945–993.

Cotter, J., I. Tuna, and P. D. Wysocki (2006). Expectations management and beatable targets: How do analysts react to explicit earnings guidance?*. *Contemporary Accounting Research 23*(3), 593–624.

Das, S., R.-J. Guo, and H. Zhang (2006). Analysts' selective coverage and subsequent performance of newly public firms. *The Journal of Finance 61*(3), 1159–1185.

Elliott, G., I. Komunjer, and A. Timmermann (2008). Biases in macroeconomic forecasts: Irrationality or asymmetric loss? *Journal of the European Economic Association 6*(1), 122–157.

Elliott, G. and A. Timmermann (2008, March). Economic forecasting. *Journal of Economic Literature 46*(1), 3–56.

Epstein, L. G. and M. Schneider (2008). Ambiguity, information quality, and asset pricing. *The Journal of Finance 63*(1), 197–228.

Farhi, E. and I. Werning (2019). Monetary policy, bounded rationality, and incomplete markets. *American Economic Review 109*(11), 3887–3928.

Farmer, L., E. Nakamura, and J. Steinsson (2021). Learning about the long run. *Working Paper*.

Feng, M. and S. McVay (2010). Analysts' incentives to overweight management guidance when revising their short-term earnings forecasts. *The Accounting Review 85*(5), 1617–1646.

Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review 110*(8), 2271–2327.

Gallant, A. R., M. Jahan-Parvar, and H. Liu (2015). Measuring ambiguity aversion. *Working Paper*.

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

Gentzkow, M. and J. M. Shapiro (2006). Media bias and reputation. *Journal of political Economy 114*(2), 280–316.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of mathematical economics 18*(2), 141–153.

Greenwood, R. and A. Shleifer (2014). Expectations of returns and expected returns. *Review of Financial Studies 27*(3), 714–746.

Groysberg, B., P. M. Healy, and D. A. Maber (2011). What drives sell-side analyst compensation at high-status investment banks? *Journal of Accounting Research 49*(4), 969–1000.

Johnson, T. L., J. Kim, and E. C. So (2019, 11). Expectations Management and Stock Returns. *The Review of Financial Studies 33*(10), 4580–4626.

Kahneman, D. and A. Tversky (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology 3*(3), 430–454.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica 73*(6), 1849–1892.

Kohlhas, A. and A. Walther (2021). Asymmetric attention. *American Economic Review (Forthcoming)*.

Kuang, P. and K. Mitra (2016). Long-run growth uncertainty. *Journal of Monetary Economics 79*, 67–80.

Kuang, P., K. Mitra, L. Tang, and S. Xie (2023). Macroprudential policy and housing market expectations. *Working Paper*.

Lian, C. (2020, 12). A Theory of Narrow Thinking. *The Review of Economic Studies*.

Ma, Y., T. Ropele, D. Sraer, and D. Thesmar (2020). A quantitative analysis of distortions in managerial forecasts. *Working Paper*.

Maslar, D. A., M. Serfling, and S. Shaikh (2021). Economic downturns and the informativeness of management earnings forecasts. *Journal of Accounting Research 59*(4), 1481–1520.

Matsumoto, D. A. (2002). Management's incentives to avoid negative earnings surprises. *The Accounting Review 77*(3), 483–514.

McNichols, M. and P. C. O'Brien (1997). Self-selection and analyst coverage. *Journal of Accounting Research 35*, 167–199.

Niehaus, G. and D. Zhang (2010). The impact of sell-side analyst research coverage on an affiliated broker's market share of trading volume. *Journal of Banking and Finance 34*(4), 776–787.

Nimark, K. P. and S. Pitschner (2019). News media and delegated information choice. *Journal of Economic Theory 181*, 160–196.

Petersen, M. A. (2009). Estimating standard errors in finance panel data sets: Comparing approaches. *The Review of financial studies 22*(1), 435–480.

Rogers, J. L. and A. Van Buskirk (2013). Bundled forecasts in empirical accounting research. *Journal of Accounting and Economics 55*(1), 43–65.

Solomon, D. H. (2012). Selective publicity and stock prices. *The Journal of Finance 67*(2), 599–638.

Wald, A. (1950). Statistical decision functions. *Wiley*.

Wang, C. (2020). Under- and over-reaction in yield curve expectations. *Working paper*.

Appendix

A. Data

A.1. Sample Construction

First, we retrieve all quarterly earnings guidance from the I/B/E/S Guidance Detail file issued for the current quarter by firm management from 1994 to 2017. The sample starts in 1994 as this is the first year when the I/B/E/S systematically collected information on managerial guidance.[16] We only keep closed-ended managerial guidance, including point and range forecasts, to quantify and compare them with analysts' forecasts. Consistent with the literature, the value of the guidance is set to equal the midpoint if it is a range forecast.

Second, given that our focus is on analysts' belief-updating process upon receiving new information from firm management, we exclude all managerial guidance bundled with earnings announcements.[17] We only consider unbundled guidance, partly because it is nearly impossible to distinguish whether a forecast revision reflects information gained from forward-looking managerial guidance or from the realized prior earnings when both of them occur simultaneously.

Third, for firm-quarters in which managers provide multiple rounds of earnings guidance (at different dates during the period from two days after the prior quarter earnings announcement date and the current quarter earnings announcement date), we only retain the latest guidance before the current quarter earnings announcement.[18]

Fourth, we then obtain individual analysts' EPS forecasts for a firm-quarter from the I/B/E/S Estimates (the Unadjusted Detail History file) and match them with the I/B/E/S Guidance data using the same firm identifier (I/B/E/S ticker). Because earnings projections in the I/B/E/S Guidance Detail file are provided on a split-adjusted basis, we manually split-adjust both individual analysts' forecasts and managerial projections so that they are comparable with the ultimate realized earnings announced for

[16] The coverage bias in the management forecast data documented by Chuk, Matsumoto, and Miller (2013) is less of a concern in this particular setting. First, we obtain management forecast data from the I/B/E/S Guidance Detail file rather than the problematic First Call CIG database. Second, the focus of this paper is to understand how analysts update their beliefs given new information, i.e., management guidance in our setting. While the decision on the issuance of management guidance itself is also an important research question, it is not the focus of this paper. Third, the fact that we require at least one analyst issuing forecasts for a firm alleviates the concern that guidance data are more likely to be collected for firms with analyst coverage. Fourth, our results are robust to starting the sample period in 1998, after which the coverage bias has been shown to be relatively small.

[17] Bundled guidance is defined as the managerial forecasts issued within 2 days around the actual earnings announcement date (Rogers and Van Buskirk 2013).

[18] However, our results are not sensitive to this specific choice and are qualitatively unchanged if we either keep the earliest guidance issued during a quarter or discard all firm-quarters with multiple guidance.

the forecasted quarter. The realized earnings data are also obtained from the I/B/E/S Estimates. Following a standard practice in the literature, we deflate the EPS estimates by the stock price at the beginning of the quarter using data retrieved from the CRSP.[19] To avoid the small price deflator problem that may distort the distribution, we exclude observations with a stock price of less than one dollar.

Finally, in these data, the initial analyst forecasts are defined and constructed by individual analyst forecasts that are issued after the prior quarter earnings announcement date and are the most updated forecasts before the earnings guidance. The revised analyst forecasts are defined as those issued by the same set of analysts on or immediately after the earnings guidance date. For analysts who initially offer forecasts but provide no forecast revisions until the earnings announcement, we assume that their revised forecasts remain the same as their initial forecasts, a practice consistent with prior literature (Feng and McVay 2010; Maslar, Serfling, and Shaikh 2021).

Suppose that a typical fiscal quarter ends at Q_t, and its realized earnings are usually announced at A_t after the end of the quarter Q_t (The Securities and Exchange Commission requires public firms to file their financial statements within 45 days after the end of the fiscal quarter). Similarly, the earnings announcement date A_{t-1} for quarter $t-1$ would also happen after Q_{t-1}. In this paper, we retrieve earnings guidance that is issued by firm management on any date between A_{t-1} and A_t. Because an increasing number of firms bundle their earnings projections for quarter t with the announcement of the realized earnings for quarter $t-1$, we further require the guidance to be unbundled (as justified earlier). That is, we only consider guidances issued between two dates, i.e., A_{t-1} and A_t. Given earnings guidance G_t, we can accordingly identify the sequence of analysts' earnings forecasts for the same quarter. We define analysts' forecasts that are issued after A_{t-1} but at the latest before G_t as their initial forecast and the forecast that is issued on or after G_t but before A_t as their revised forecast. As noted above, for analysts who provide an initial forecast but do not revise, we assume that the revised forecast remains the same as the initial forecast. There are two exceptions to this general timing. First, it might be the case that G_t lies between Q_t and A_t, in which case we term the guidance a preannouncement following the convention in the literature. Second, firm management can offer more than one earnings guidance, and therefore, G_t may appear multiple times during the period. In this case, we only retain the latest guidance before A_t.

[19]We provide a robustness check for our empirical results without deflating the EPS estimates with stock prices and show that this practice does not affect our findings.

A.2. Summary of Statistics

Table A1. *Summary of Statistics*

	(1) N	(2) mean	(3) sd	(4) p25	(5) p50	(6) p75
Initial forecasts	110,895	0.0120	0.0129	0.0070	0.0123	0.0180
Revised forecasts	110,895	0.0104	0.0149	0.0057	0.0113	0.0173
Forecast revision	110,895	-0.0016	0.0055	-0.0017	0.0000	0.0000
Forecast errors	110,895	-0.0000	0.0047	0.0000	0.0003	0.0011
Surprise	110,895	-0.0040	0.0171	-0.0062	-0.0012	0.0003
Managerial guidance	16,241	0.0067	0.0293	0.0027	0.0089	0.0160
Earnings	16,241	0.0089	0.0197	0.0044	0.0112	0.0177

B. Robustness

B.1. Robustness: Overreaction

Table A2. *Forecast Error on Forecast Revision: Samples*

	Outcome Variable: Forecast Error FE_i					
	Winsorization at the 1% and 99%			Winsorization at the 2.5% and 97.5%		
	Excl Pre-anc	Excl Multiple	Excl Both	Excl Pre-anc	Excl Multiple	Excl Both
	(1)	(2)	(3)	(4)	(5)	(6)
FR_i	-0.0731***	-0.1614***	-0.1579***	-0.0684***	-0.1584***	-0.1590***
	(0.0272)	(0.0210)	(0.0491)	(0.0224)	(0.0165)	(0.0377)
Firm FEs	YES	YES	YES	YES	YES	YES
Obs.	51,692	47,311	19,162	51,692	47,311	19,162
Adj R-sq.	0.2380	0.2856	0.3348	0.2344	0.2636	0.3179

The standard errors are clustered on firm and calendar year-quarter following Petersen (2009).*** p<0.01, ** p<0.05, * p<0.1

Subsamples. The result in column (1) of Table A2 is based on a sample excluding all firm-quarters with pre-announcement guidance, which is defined as the guidance issued between firms' fiscal quarter-end and the earnings announcement date for the quarter. The result in column (2) of Table A2 is based on a sample excluding all firm-quarters with multiple guidances. The result in column (3) of Table A2 is based on a sample excluding all firm-quarters with either pre-announcement guidance or multiple guidances. To ensure that our results are not driven by outliers, we winsorize FE_{ijt} and FR_{ijt} at the 2.5% and 97.5% of their respective distributions and re-estimate Equa-

Table A3. Forecast Error on Forecast Revision: Trimming Outliers

	Outcome Variable: Forecast Error FE_i							
	Trimmed at 1% and 99%				Trimmed at 2.5% and 97.5%			
	Full	Excl Pre-anc	Excl Multiple	Excl Both	Full	Excl Pre-anc	Excl Multiple	Excl Both
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
FR_i	-0.1023***	-0.0903***	-0.1694***	-0.1890***	-0.0836***	-0.0709***	-0.1537***	-0.1641***
	(0.0107)	(0.0217)	(0.0133)	(0.0299)	(0.0083)	(0.0136)	(0.0102)	(0.0186)
Firm FEs	YES	YES	YES	YES	YES	YES	YES	YES
Obs.	106,714	50,099	44,667	18,303	100,577	47,559	41,149	17,020
Adj R-sq.	0.1880	0.2340	0.2577	0.3239	0.1722	0.2241	0.2372	0.2918

The standard errors are clustered on firm and calendar year-quarter following Petersen (2009).*** p<0.01, ** p<0.05, * p<0.1

Table A4. Forecast Error on Forecast Revision: Median Coefficients

	Firm-by-Firm Forecast Error on Forecast Revision Regressions			
	Winsorization at the 1% and 99%		Winsorization at the 2.5% and 97.5%	
	Baseline	Unscaled	Baseline	Unscaled
	(1)	(2)	(3)	(4)
Median	-0.1616	-0.1747	-0.1647	-0.1612
(p 2.5, p 97.5)	(-0.2256, -0.1456)	(-0.2396, -0.1617)	(-0.2281, -0.1429)	(-0.2244, -0.1519)
(p 5.0, p 95.0)	(-0.2173, -0.1483)	(-0.2317, -0.1633)	(-0.2107, -0.1477)	(-0.2172, -0.1537)
No. of firms.	2849	2849	2849	2849

We report the 5% (row 2) and 10% (row 3) bootstrapped confidence intervals, the boundaries of which are the 2.5, 5.0, 95.0, and 97.5 percentiles of the estimated median coefficients out of the 500 bootstrap samples. Following Bordalo, Gennaioli, Ma, and Shleifer (2020), these samples are obtained from block bootstrap the panel using blocks of 30 quarters.

tion (1). The results for the corresponding subsamples are reported in column (4), (5) and (6).

Trimming Outliers. In the main text, we estimate Equation (1) with winsorized data to mitigate the influence of outlier observations. In this Appendix, we re-estimate Equation (1) with trimmed data and examine the robustness of our results reported in the main text. The corresponding results are summarized in Table A3. All results are robust, thus suggesting that our results are not sensitive to the way in which we handle outliers.

Nickel Bias. In a dynamic panel, the presence of firm fixed effects may raise concerns about "Nickel bias." To address this issue, we follow Bordalo, Gennaioli, Ma, and Shleifer (2020) and conduct the FE-on-FR regression on a firm-by-firm basis. We report the median coefficient from the distribution of b_1 estimates. The corresponding results are summarized in Table A4. The median coefficient from the firm-by-firm regressions aligns with that of the pooled regression with firm fixed effects, mitigating

Table A5. Forecast Error on Forecast Revision: Analyst Fixed Effects

	Outcome Variable: Forecast Error FE_i					
	Winsorization at the 1% and 99%			Winsorization at the 2.5% and 97.5%		
	Baseline	Control	Unscaled	Baseline	Control	Unscaled
	(1)	(2)	(3)	(4)	(5)	(6)
FR_i	-0.0876***	-0.0877***	-0.0875***	-0.0836***	-0.0835***	-0.0696***
	(0.0141)	(0.0141)	(0.0120)	(0.0115)	(0.0115)	(0.0097)
Earnings of the Last Quarter		0.0003			-0.0020	
		(0.0071)			(0.0050)	
Analyst FEs	YES	YES	YES	YES	YES	YES
Firm FEs	YES	YES	YES	YES	YES	YES
Obs.	110,895	110,895	110,895	110,895	110,895	110,895
Adj. R-sq	0.2355	0.2354	0.2075	0.2202	0.2203	0.2125

The standard errors are clustered on firm and calendar year-quarter following Petersen (2009).*** p<0.01, ** p<0.05, * p<0.1

concerns about Nickel bias in identifying the average overreaction.

Analyst Fixed Effects. To mitigate concerns regarding potential impacts of time-invariant analyst characteristics, we incorporate additional controls for analyst fixed effects in the estimation of Equation (1). The summarized results in Table A5 remain robust, indicating that our findings are not sensitive to analyst fixed effects.

B.2. Robustness: Heterogeneous Overreaction

Heterogeneous Overreaction with Trimming 2% Outlier observations.

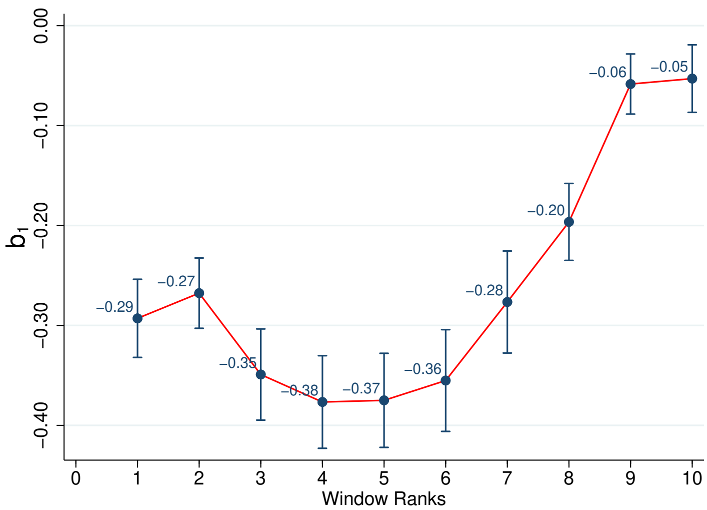

Figure A1. *Heterogeneous Overreaction, Trimming 2% Outliers. The estimated coefficients of the FE-on-FR regressions b_1 and 95% confidence interval for each running decile window is plotted against the window rank. A running decile window j covers decile $j-1$, j, and $j+1$ if $j \neq 1$ or $j \neq 9$; the running decile window 1 covers deciles 1 and 2 and the running decile window 10 covers deciles 9 and 10.*

Heterogeneous Overreaction for Each Decile of Surprises.

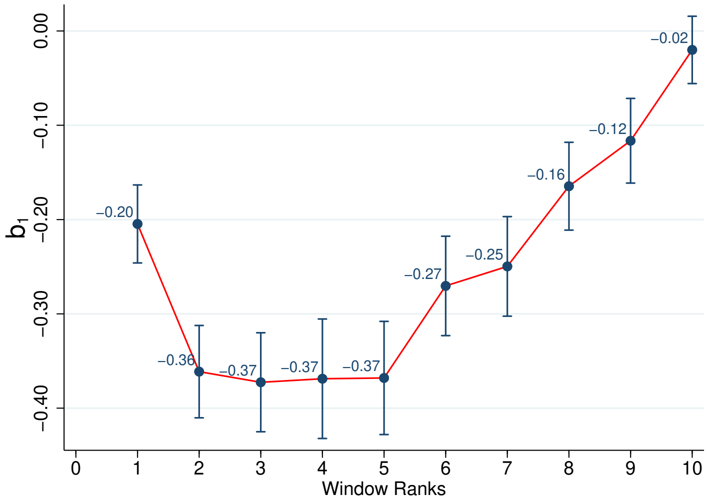

Figure A2. *Heterogeneous Overreaction. The estimated coefficients of the FE-on-FR regressions b_1 and 95% confidence interval for each decile of surprises, without using running windows.*

B.3. Robustness: Forecast Revisions and Surprises

One valid concern is that the decreasing arms of the estimated relationship might be driven by a small number of observations in the tails. Ultimately, the confidence intervals become very wide when the surprises are relatively large in magnitude. However, we find that this is not the case. In our baseline setup, there are 3,932 observations to the left of the trough and 2,218 observations to the right of the peak, which account for close to 6% of the total observations used in this estimation. Given the number of observations utilized, this concern is alleviated.

Another potential issue is that whether to offer earnings guidance could be strategically chosen by firms, which could affect our estimations. First, this is unlikely because firms do not make decisions about whether they disclose the earnings guidance on a quarterly basis and typically continue to provide earnings guidance for an extended period of time (Chen, Matsumoto, and Rajgopal 2011). Second, we construct a subsample in which we only include earnings forecasts conditional on firms (whose earnings are being forecasted) having to release earnings guidance for at least 12 consecutive quarters during our sample period.[20] We nonparametrically re-estimate the rela-

[20]Based on the initial full sample of management guidance, we select a quarterly management guidance for our subsample if it lies in any series of at least 12 consecutive quarters where managers provide earnings forecasts in each quarter. For example, the guidance issued in 2012Q4 is selected if it is in a series of 12 consecutive quarters from 2011Q1 to 2013Q4 with management guidance. The subsample consists of 49,116 observations with 5,601 firm-quarters. We also vary the threshold for the number of consecutive quarters, such as 8 and 16. The results are rather similar.

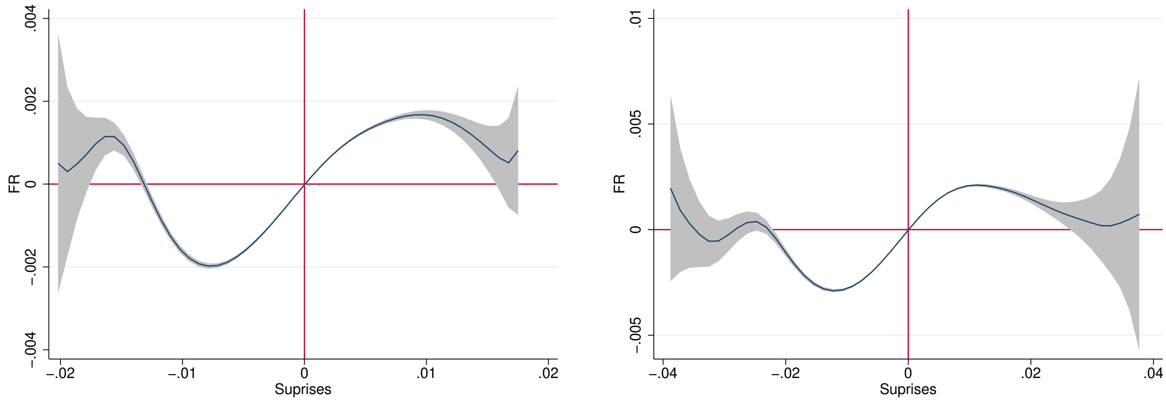

(a) Firms with at least 12 consecutive guidances

(b) Exclusion of financial crisis

Figure A3. *The figure illustrates the relationship between forecast revisions and surprises in managerial guidances, both of which are trimmed at 2.5% and 97.5%. The relation is non-parametrically estimated using the Epanechnikov kernel and the third degree of the polynomial smooth. Panel (a) reports the estimation for the subsample that only includes earnings forecasts on condition that firms release earnings guidances more than 12 consecutive quarters during our sample period. Panel (b) reports the estimation for the subsample that excludes observations during the financial crisis. In both cases, the forecast revision is decreasing, increasing and decreasing and asymmetric around the origin.*

tionship between forecast revisions and surprises following the procedure described above. The results are presented in Panel (a) of Figure A3. They are rather similar to those obtained using the full sample, and the two key characteristics are even more pronounced. The calculated asymmetry measure Ξ is 0.23 (see Equation (2)). This implies that, on average, unfavorable surprises result in revisions that are 23% higher than revisions for favorable surprises. Therefore, the concern of strategic disclosure is inconsequential for our findings.

Our data cover the period of the 2007-2009 financial crisis, and it is not unlikely that financial market participants behaved abnormally during that period, which could affect the relationship in which we are interested. To investigate this possibility, we remove the data from 2007 to 2009, i.e., the financial crisis period, and re-estimate the relationship. The results are presented in Panel (b) of Figure A3, which are very similar to those obtained by using the full sample.

C. Quantitative Analysis

We provide the details of our quantitative analysis in this section. In the benchmark model presented in the main text, we do not consider any private information available to analysts, since it would not affect our analytical results qualitatively. Technical Appendix provides the characterization of the model with such unobservable private information. In our quantitative analysis, we intend to directly relate the relationship characterized in our model (Section 3) to that in the data (Section 2.3 and 2.4). Therefore, we explicitly model the unobservable private information. Specifically, we assume analyst i is endowed with private information x_i after releasing her initial forecast:

$$x_i = \theta + \varepsilon_i; \qquad \varepsilon_i \sim N\left(0, 1/\tau_x\right). \tag{A1}$$

where τ_x refers to the precision of private information x_i. Note that in our quantitative analysis, we construct and work with surprises observable to the econometrician from our simulated data and estimate the relationship between forecast revisions and surprises in the same way as we do with the data.

Our model is fully specified by two sets of parameters and one distribution. First, two parameters characterize the preferences of analysts, i.e., ambiguity aversion λ, and analysts' attitude toward earnings β. Second, there is a set of volatilities, i.e., the objective volatility of earnings σ_θ, the objective volatility of managerial guidance σ_Y, the volatility of initial endowed information about earnings before the initial forecast σ_z, and the volatility of private information σ_x. Third, the analysts' prior belief about guidance quality $p\left(\tau_y\right)$ defined in Section 3.1 also needs to be specified. We assume that the ratio $\delta \equiv \tau_y/(\tau_\theta + \tau_z + \tau_x + \tau_y)$ is a uniform distribution over $[L, U]$, where $0 \leq L < U \leq 1$. The advantage of this transformation is that we can entertain the possibility that τ_y is very large, without dealing with a very wide support for τ_y, which economizes our computation. The upper bound U (lower bound L) regulates the perceived largest (smallest) possible precision for managerial guidance.

To estimate the set of parameters $\Theta = \{\lambda, \beta, L, U, \sigma_\theta, \sigma_x, \sigma_Y, \sigma_z\}$, we follow Chernozhukov and Hong (2003) in computing Laplace type estimators (LTE) with an MCMC approach, and the "distance" between the empirical and simulated revision-surprise relationships is constructed in the fashion of the method of simulated moments.

To define the distance, we first choose $N = 50$ equally spaced points for surprises between $[-0.025, 0.030]$, within which the empirical relationship (nonparametrically estimated in section 2.4) decrease, increase and then decrease. Then, we derive the corresponding values of forecast revisions from the estimated revision-surprise rela-

Table A6. *Targeted Moments of FR-on-Surprise Relation*

Surp	FR	Surp	FR	Surp	FR	Surp	FR	Surp	FR
-0.02500	0.00004	-0.01378	-0.00313	-0.00255	-0.00098	0.00867	0.00207	0.01990	0.00155
-0.02388	-0.00000	-0.01265	-0.00325	-0.00143	-0.00054	0.00980	0.00214	0.02102	0.00140
-0.02276	-0.00015	-0.01153	-0.00330	-0.00031	-0.00011	0.01092	0.00219	0.02214	0.00128
-0.02163	-0.00044	-0.01041	-0.00323	0.00082	0.00031	0.01204	0.00220	0.02327	0.00118
-0.02051	-0.00082	-0.00929	-0.00307	0.00194	0.00068	0.01316	0.00218	0.02439	0.00110
-0.01939	-0.00126	-0.00816	-0.00286	0.00306	0.00102	0.01429	0.00214	0.02551	0.00099
-0.01827	-0.00175	-0.00704	-0.00258	0.00418	0.00132	0.01541	0.00207	0.02663	0.00085
-0.01714	-0.00220	-0.00592	-0.00222	0.00531	0.00157	0.01653	0.00196	0.02776	0.00068
-0.01602	-0.00259	-0.00480	-0.00184	0.00643	0.00178	0.01765	0.00184	0.02888	0.00050
-0.01490	-0.00293	-0.00367	-0.00142	0.00755	0.00195	0.01878	0.00170	0.03000	0.00036

Table A7. *Estimated Model Parameters*

	Mean	90% HPDI	95% HPDI
λ	449.9	$(411.9, 504.0)$	$(379.5, 504.2)$
β	1.379	$(0.773, 1.971)$	$(0.694, 2.092)$
U	0.772	$(0.676, 0.855)$	$(0.674, 0.875)$
L	0.082	$(0.036, 0.119)$	$(0.030, 0.121)$
$100\sigma_x$	0.472	$(0.332, 0.593)$	$(0.305, 0.625)$
$100\sigma_z$	0.186	$(0.140, 0.234)$	$(0.137, 0.240)$
$100\sigma_Y$	0.435	$(0.416, 0.453)$	$(0.411, 0.453)$

tionship and denote them with the vector $\hat{m}$. Table A6 presents these targeted moments.

We further construct the vector m, i.e., the model counterpart of $\hat{m}$, which is estimated by using our simulated dataset. Specifically, for each set of model parameters, we simulate our model and estimate the revision-surprise relationship with the same nonparametric regression as for the empirical data (see Section 2.4). We then obtain the vector m from the estimated relationship between forecast revisions and surprises observable to the econometrician. The distance that we construct is:

$$\Lambda(\Theta) = \frac{1}{N} \left[m(\Theta) - \hat{m} \right]' \hat{W} \left[m(\Theta) - \hat{m} \right].$$

where $N = 50$ is the length of the vector of targeted moments $\hat{m}$ and $\hat{W}$ is the weighting matrix with diagonal elements being the precision of moments $\hat{m}$. Our goal is to choose model parameters to "minimize" the distance $\Lambda(\Theta)$ in a pseudo Bayesian manner by using MCMC with the Metropolis-Hastings algorithm.

A few remarks regarding the simulation procedure are in order. First, we choose

σ_θ, i.e., the standard deviation of θ, to exactly match the empirical counterpart of an unconditional standard deviation of realized earnings (after removing the firm and time fixed effects). As a result, the calibrated value of $100\sigma_\theta$ is 0.985. Second, when we simulate the model, we feed surprises (to the econometrician) uncovered from the empirical data into our simulation. We recover the corresponding surprises to the analysts and then obtain updated forecasts by using decision rules in our model. Third, in this model, the unconditional volatility of surprises to the econometrician is determined by both σ_Y and $1/\sigma_\theta^2 + 1/\sigma_z^2$. We directly estimate $1/\sigma_\theta^2 + 1/\sigma_z^2$ in the estimation and back out σ_Y by requiring that the unconditional volatility of surprises matches its empirical counterpart, an internal consistency condition for our estimation strategy.

The estimated parameters are reported in Table A7 together with the 90% and 95% high posterior density interval (HPDI). The relative magnitude of the estimated volatilities appears to be reasonable. The volatility of private information is larger than that of earnings. The managerial guidance is much more precise than the private information. This is likely because there may not be much private information that arrives during the time window that we construct (i.e., between the two forecasts around the date of managerial guidance release). Based on this set of parameters, the response of forecast revisions to surprises under noisy rational expectation (i.e., κ^{RE}) is 0.132. The upper bound for the subjective belief on managerial guidance precision is 0.772, and the lower bound is 0.082. The support is large enough to allow sufficient ambiguity and encompass κ^{RE}. The degree of ambiguity aversion is the key, and its value is estimated to be $\lambda = 449.9$. Finally, the parameter β that characterizes the utility function is estimated to be positive (i.e., slightly larger than 1), indicating that analysts are likely to care about the earnings performance of firms that they cover.

Using this set of estimated parameters, we simulate the model and nonparametrically estimate the revision-surprise relationship with the simulation data. In Figure 4(a), we display the relationship, together with its empirical counterparts (previously shown in Figure 3(a)). In Figure 4(b), we illustrate its implied derivative with respect to surprises. Our model can successfully capture both features of nonmonotonicity and asymmetry.

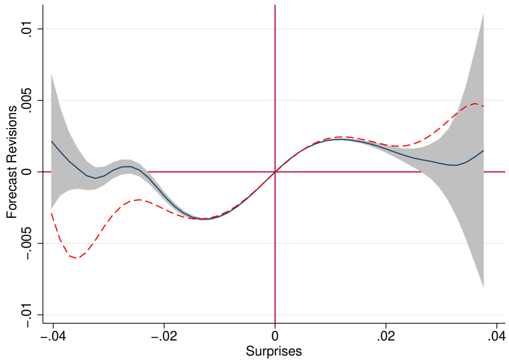

(a) $\lambda = 449.9$, Nonparametric estimation

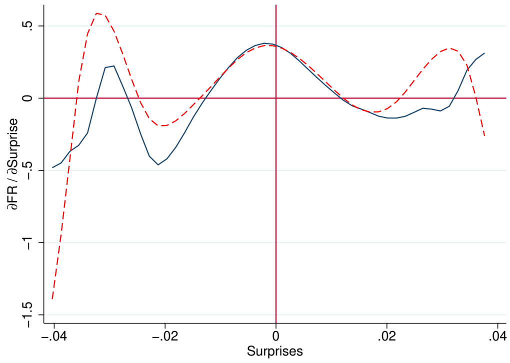

(b) $\lambda = 449.9$, Derivative

Figure A4. *The revision-surprise relationship nonparametrically estimated with simulation data. We simulate the model with the set of parameters reported in Table A7 and, in particular, $\lambda = 449.9$. The dashed line in panel (a) illustrates the revision-surprise relationship estimated with simulation data. The empirical counterpart (i.e., the solid line) and its confidence interval (i.e., the shaded area) are also plotted for comparison. Panel (b) illustrates the derivative of the revision-surprise relationship with the dashed line. Its empirical counterpart is illustrated with the solid line.*

D. Discussions: Alternative Hypotheses

In this section, we provide a detailed discussion of several likely candidates to account for the empirical patterns documented in our paper: diagnostic expectation, over-confidence, loss aversion, dynamic models, and agency issues.

D.1. Diagnostic Expectations and Over-confidence

Bordalo, Gennaioli, Ma, and Shleifer (2020) show that forecasters with diagnostic expectations over-react to new information at the individual level. Diagnostic expectations proposed by Bordalo, Gennaioli, and Shleifer (2018) is a non-Bayesian model of belief formation that formalizes representativeness heuristic (Kahneman and Tversky 1972): agents overweight states that are more likely in light of the arrival of new signals. As a consequence, agents over-react to new information when forming expectations. Specifically, agents update their beliefs using the following distorted posterior density $f^{\psi}\left(\theta|\mathcal{I}_{it}\right)$:

$$f^{\psi}\left(\theta|\mathcal{I}_{it}\right) \propto f\left(\theta|\mathcal{I}_{it}\right)\left(\frac{f\left(\theta|\mathcal{I}_{it}\right)}{f\left(\theta|\mathcal{I}_{it-1}\right)}\right)^{\psi},$$

where $f\left(\theta|\mathcal{I}_{it}\right)$ denotes the Bayesian posterior density and the constant $\psi \geq 0$ measures the extent to which the posterior of agents with diagnostic expectations are distorted away from Bayesian benchmark. When $\psi = 0$, the model with diagnostic expectations reduces to the noisy information benchmark. Observe that $R_t\left(\theta\right) \equiv \frac{f(\theta|\mathcal{I}_{it})}{f(\theta|\mathcal{I}_{it-1})}$ measures the representativeness of the new information defined as the gap between

$\mathcal{I}_{it}$ and $\mathcal{I}_{it-1}$. When $\psi > 0$, the distorted posterior belief overweights states θ featuring $R_t(\theta) \geq 1$, which leads to overreaction to the arrival of new information (Bordalo, Gennaioli, Ma, and Shleifer 2020).

In this case, the optimal initial and updated forecasts are given by:

$$\mathrm{F}_{0i}^{\mathrm{DE}} = \mathbb{E}[\theta|z_{0i}] + \psi\left(\mathbb{E}[\theta|z_{0i}] - \mathbb{E}[\theta]\right),$$

and

$$\mathrm{F}_{i}^{\mathrm{DE}} = \mathbb{E}[\theta|z_{0i}, x_i, y] + \psi\left(\mathbb{E}[\theta|z_{0i}, x_i, y] - \mathbb{E}[\theta|z_{0i}]\right),$$

where $\mathbb{E}[\theta|\mathcal{I}_i]$ denotes the conditional expectations (i.e., Bayesian posterior). As a consequence, forecast revisions under diagnostic expectations are given by:

$$\mathrm{FR}_{i}^{\mathrm{DE}} = (1+\psi)\,\kappa^{\mathrm{RE}}\left(y - \mathrm{F}_{0i}^{\mathrm{DE}}\right) + (1+\psi)\,\kappa_x\left(x - \mathrm{F}_{0i}^{\mathrm{DE}}\right) + \psi\left[\left(\kappa_x + \kappa^{\mathrm{RE}}\right) - \frac{1}{1+\psi}\right]\mathrm{F}_{0i}^{\mathrm{DE}}, \tag{A2}$$

where $\kappa_x \equiv \frac{\tau_x}{\tau_\theta+\tau_z+\tau_x+\tau_Y} > 0$ and $\kappa^{\mathrm{RE}} \equiv \frac{\tau_Y}{\tau_\theta+\tau_z+\tau_x+\tau_Y} > 0$ are the respective weights for the private and public information in the noisy information benchmark. Similar to the noisy information benchmark, the term $y - \mathrm{F}_{0i}^{\mathrm{DE}}$ is the theoretical counterpart to the manager guidance surprises in our empirical exercise. Observe that the relationship between forecast revisions and surprises in Equation (A2) remains linear and state-independent.

Broer and Kohlhas (2022) show that over-confidence can help rationalize over-reaction documented with SPF data, in which they assume that forecasters subjectively believe that new signals are more precise than they actually are. Interestingly, once we allow for such behavioral feature in the noisy information benchmark specified in Section 3.2, the FR-on-Surprise relation is still linear, while overreaction to new information appears. To see this, specifically we assume that analysts subjective believe that $\eta \sim \mathbb{N}\left(0, 1/\bar{\tau}_y\right)$ such that $\bar{\tau}_y > \tau_Y$ and derive the relationship between forecast revisions and surprises as follows:

$$\mathrm{FR}_{i}^{\mathrm{OC}} \equiv \mathrm{F}_{i}^{\mathrm{OC}} - \mathrm{F}_{0i} = \bar{\kappa}_y\left(y_i - \mathrm{F}_{0i}\right) + \bar{\kappa}_x\left(x_i - \mathrm{F}_{0i}\right),$$

where $\bar{\kappa}_x \equiv \frac{\tau_x}{\tau_\theta+\tau_z+\tau_x+\bar{\tau}_y} > 0$ and $\bar{\kappa}_y \equiv \frac{\bar{\tau}_y}{\tau_\theta+\tau_z+\tau_x+\bar{\tau}_y} > 0$.

D.2. Loss Aversion

Another plausible conjecture is that analysts are loss-averse, which may also likely generate the empirical pattern documented in Section 2.4, given that they behave in a pessimistic way. To investigate this possibility, we consider two commonly used specifications of loss aversion: one parsimonious setup with analytical solutions (Capistrán and Timmermann 2009) and another flexible setup with more quantitative potentials (Elliott, Komunjer, and Timmermann 2008, Elliott and Timmermann 2008). In this section, we show that (1) the parsimonious setup predicts a linearly increasing relation between forecast revisions and surprises and that (2) the flexible setup predicts a monotone increasing relation between forecast revisions and surprises.

The Parsimonious Setup. We follow Capistrán and Timmermann (2009) and specify the loss function of analysts to be:

$$L\left(\mathrm{F}_i, \theta; \phi\right) = \frac{1}{\phi^2}\left[\exp\left(\phi\left(\theta - \mathrm{F}_i\right)\right) - \phi\left(\theta - \mathrm{F}_i\right) - 1\right],$$

where F_i stands for the forecast of analyst i and the parameter ϕ is a constant that captures asymmetries in the loss function. If $\phi > 0$, analysts dislike negative forecast error $\theta - \mathrm{F}_i < 0$ more than positive forecast error $\theta - \mathrm{F}_i > 0$. If ϕ goes to zero, the loss function is reduced to the standard MSE function. Information structure is the same as that of the noisy information benchmark.

Analyst i chooses the optimal forecasts F_i^L to minimax the loss function conditional on her information set, which leads to her decision rule:

$$\mathrm{F}_i^L = \mathbb{E}_i\left[\theta\right] - \frac{1}{2}\phi \mathrm{Var}_i\left[\theta\right].$$

Relative to the noisy information expectations benchmark ($\phi = 0$), the loss-averse analyst i ($\phi > 0$) would like to inflate the forecast errors and bias her forecast downward by $\frac{1}{2}\phi \mathrm{Var}_i\left[\theta\right]$.

Accordingly, the initial optimal forecast is given by:

$$\mathrm{F}_{0i}^L = \mathbb{E}\left[\theta | z_{0i}\right] - \frac{1}{2}\phi \mathrm{Var}\left[\theta | z_{0i}\right] = \frac{\tau_z}{\tau_\theta + \tau_z} z_{0i} - \frac{1}{2}\phi \frac{1}{\tau_\theta + \tau_z}.$$

and the updated optimal forecast is given by:

$$\mathrm{F}_i^L = \mathbb{E}\left[\theta | z_{0i}, x_i, y\right] - \frac{1}{2}\phi \mathrm{Var}\left[\theta | z_{0i}, x_i, y\right] = \frac{\tau_z z_i + \tau_x x_i + \tau_Y y}{\tau_\theta + \tau_z + \tau_x + \tau_Y} - \frac{1}{2}\phi \frac{1}{\tau_\theta + \tau_z + \tau_x + \tau_Y}.$$

Therefore, the forecast revision is such that

$$\mathrm{FR}_i^{\mathrm{L}} \equiv \mathrm{F}_i^L - \mathrm{F}_{0i}^L = \kappa_y\left(y - \mathrm{F}_{i0}^{\mathrm{L}}\right) + \kappa_x\left(x_i - \mathrm{F}_{i0}^{\mathrm{L}}\right). \tag{A3}$$

where $\kappa_x \equiv \frac{\tau_x}{\tau_\theta + \tau_z + \tau_x + \tau_y} > 0$ and $\kappa_y \equiv \frac{\tau_Y}{\tau_\theta + \tau_z + \tau_x + \tau_y} > 0$ are the relevant weights for the private and public information under noisy information benchmark. Observe that the relation between forecast revisions and guidance surprises are still linear.

The Flexible Setup. Following Elliott, Komunjer, and Timmermann (2008) and Elliott and Timmermann (2008), we specify the loss function of analysts to be:

$$L_p\left(\mathrm{F}_i, \theta; \alpha\right) = \left[\alpha + (1-2\alpha)\,\mathbb{1}\left\{\theta - \mathrm{F}_i < 0\right\}\right]\left|\theta - \mathrm{F}_i\right|^p, \tag{A4}$$

where $\mathbb{1}\{\cdot\}$ denotes an indicator function, the parameter $\alpha \in (0,1)$ is a constant captures asymmetries in the loss function, and the parameter $p \geq 1$ is another constant that controls the curvature of the loss functions.

This specification of loss function is flexible and can be reduced to many commonly used loss functions in the literature (as shown in Elliott, Komunjer, and Timmermann 2008 and Elliott and Timmermann 2008). For example, if $\alpha = \frac{1}{2}$ the loss function is symmetric. It can be further reduced to the standard MSE loss function if $p = 2$ or mean absolute error function if $p = 1$. In particular, when $\alpha < \frac{1}{2}$, negative forecast error $(\theta - \mathrm{F}_i < 0)$ disproportionately induces more loss than positive forecast error $(\theta - \mathrm{F}_i > 0)$, indicating that analysts are loss averse.

For the ease of exposition, we focus our analysis on the parameter space that $p = 2$, that is, the loss function is of a generalized MSE form. However, it is noted that all results presented below can be extended to the general case that $p \geq 1$.

Analyst i chooses the optimal forecasts F_i^L to minimax the loss function conditional on her information set. Implicitly, her optimal decision rule is characterized by:

$$\int_{-\infty}^{+\infty}\left(\theta - \mathrm{F}_i^L\right) f\left(\theta|\mathcal{I}_i\right)\mathrm{d}\theta + \frac{1-2\alpha}{\alpha}\int_{-\infty}^{\mathrm{F}_i^L}\left(\theta - \mathrm{F}_i^L\right) f\left(\theta|\mathcal{I}_i\right)\mathrm{d}\theta = 0, \tag{A5}$$

where $f\left(\theta|\mathcal{I}_i\right)$ denotes the posterior density of the fundamental θ with respect to the information set $\mathcal{I}_i$.

It is worth noting that when $\alpha = \frac{1}{2}$ (i.e., the loss function is symmetric), only the first term of the LHS of (A5) is relevant. Therefore, the optimal forecast is just the conditional expectation as in the noisy information setup:

$$\int_{-\infty}^{+\infty}\left(\theta - \mathrm{F}_i^L\right) f\left(\theta|\mathcal{I}_i\right)\mathrm{d}\theta = 0 \Rightarrow \mathrm{F}_i^L = \mathbb{E}\left[\theta|\mathcal{I}_i\right].$$

where $\mathbb{E}\left[\theta|\mathcal{I}_i\right]$ denotes the conditional mean under Bayesian posterior.[21]

[21]For the general case $p \geq 1$, the same result holds. Intuitively, as long as the information structure is symmetric, any symmetric loss function implies that optimal forecasts are the conditional expectations (Bhattacharya and Pfleiderer 1985).

Table A8. Forecast Errors, Forecast Revisions and Earnings in the Last Quarter

	Outcome Variable: in Quarter t for Firm j, analyst i's			
	Forecast Errors		Forecast Revisions	
	1% and 99%	2.5% and 97.5%	1% and 99%	2.5% and 97.5%
	(1)	(2)	(3)	(4)
Earnings in the Last Quarter ($t-1$)	0.0024	0.0006	-0.0048	-0.0058
	(0.0070)	(0.0048)	(0.0066)	(0.0053)
Surprise$_i$			0.1468***	0.2445***
			(0.0125)	(0.0128)
Constant	-0.0000	0.0001**	-0.0009***	-0.0004***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Time FEs	NO	NO	YES	YES
Analyst FEs	NO	NO	YES	YES
Firm FEs	YES	YES	YES	YES
Obs.	110,895	110,895	110,895	110,895
Adj. R-sq	0.2341	0.1850	0.3943	0.4588

The standard errors are clustered on firm and calendar year-quarter following (Petersen 2009).
*** p<0.01, ** p<0.05, * p<0.1

Furthermore, observe that the second term of the LHS of (A5) is negative if and only if analysts are loss averse ($\alpha < \frac{1}{2}$). Therefore, as in the parsimonious setup, loss averse analysts would like to inflate the forecast errors by biasing their forecasts downwards

$$\mathrm{F}_i^L < \mathbb{E}\left[\theta|\mathcal{I}_i\right].$$

Proposition 7. *With the flexible specification of the loss function* (A4)*, forecast revisions are globally monotone in surprises.*

To understand this lemma, we note that a sufficient condition for the global monotonicity is that the optimal forecast F_i^L is globally monotone in signals. According to (A5), the optimal forecasts can be written as the summation of the Bayesian posterior mean and a bias:

$$\mathrm{F}_i^L = \underbrace{\mathbb{E}\left[\theta|\mathcal{I}_i\right]}_{\text{Bayesian}} + \underbrace{\frac{1-2\alpha}{\alpha}\int_{-\infty}^{\mathrm{F}_i^L}\left(\theta - \mathrm{F}_i^L\right) f\left(\theta|\mathcal{I}_i\right)\mathrm{d}\theta}_{\text{Bias}}. \tag{A6}$$

In the proof of Lemma 7, we demonstrate that both the Bayesian posterior mean and the bias are increasing in signals of the fundamental θ, which is the same as the parsimonious setup. Though the FR-on-Surprise relation can be non-linear in this case, it is still globally monotone, which is still inconsistent with the documented non-monotone FR-on-Surprise relation in Section 2.4.

D.3. Dynamic models.

To illustrate, recall that we show in Section 2.2 that earnings in the last quarter cannot predict forecast errors conditional on forecast revisions (see Table 1). How do forecast errors react to earnings in the last quarter without controlling forecast revisions? We run a regression of forecast errors (i.e., analyst i's forecast error in quarter t for firm j) on earnings in the last quarter (quarter $t-1$ and firm j) with a full set of fixed effects as in Equation (1). We report the estimation results in Table A8. Column (1) shows that the estimated coefficient is very small and insignificant, suggesting that earnings in the last quarter cannot predict analysts' forecast errors. To ensure robustness, we winsorize the FE_{ijt} and the last quarter earnings at 2.5% and 97.5% of their respective distributions and re-estimate Equation (1). We report the results in column (2), which are consistent with those in column (1). This result is in contrast with both Kohlhas and Walther (2021) and Afrouzi, Kwon, Landier, Ma, and Thesmar (2022).

Furthermore, in this setting, we predict that forecast revisions would not be affected by earnings in the last quarter. To confirm this, we run forecast revisions on surprises as well as earnings in the last quarter. The results are reported in columns (3) and (4) of Table A8 for different levels of winsorization. The estimated coefficient on earnings in the last quarter is very small and insignificant, suggesting that they do not affect analysts' forecast revisions in the current period either.

This set of results is intuitive: the initial forecast in this setting absorbed information contained in earnings in the last quarter, which do not impact forecast revisions that take place in the current quarter. Therefore, forecast revisions reflect the impact of earnings guidance, instead of the impact of earnings in pervious quarters. By contrast, in studies using SPF data, "initial forecasts" for a random variable x_{t+k} in period $t+k$ are made in period $t-1$ and "updated forecasts" are made in period t after observing the current realization of the variable x_t.

D.4. Agency issues

Skewed information reliability. It is empirically testable whether positive guidance is of lower quality on average. We regress a measure of guidance quality on guidance negativity and report the results in Table A9. Specifically, the dependent variable is the absolute difference between managerial guidance and actual realized earnings per share for firm i in quarter t, scaled by the stock price at the beginning of quarter t. The independent variable of our interest, Negative Guidance, is an indicator, which is equal to 1 if the managerial guidance is smaller than the median of individual analysts's initial forecasts before guidance, and 0 otherwise. We control for firm and calendar year-quarter fixed effects so that the results cannot be explained by any

Table A9. *Guidance Quality and Negativity*

	Outcome Variable: Absolute Difference between Guidance and Earnings			
	Sample: Full		Exclude Conforming	
	1% and 99%	2.5% and 97.5%	1% and 99%	2.5% and 97.5%
	(1)	(2)	(3)	(4)
Negative Guidance	0.0012***	0.0008***	0.0010***	0.0003
	(5.1339)	(3.7038)	(3.1044)	(1.2632)
Constant	0.0050***	0.0048***	0.0057***	0.0056***
	(35.1519)	(38.5143)	(26.4874)	(30.2028)
Time FEs	YES	YES	YES	YES
Firm FEs	YES	YES	YES	YES
Observations	15,528	15,528	13,476	13,500
Adjusted R-squared	0.6105	0.5395	0.6151	0.5428

Notes: The observation numbers in columns (3) and (4) vary because the numbers of conforming cases vary due to Winsorization. The standard errors are clustered on firm and year-quarter. *** p<0.01, ** p<0.05, * p<0.1.

time-invariant firm characteristics and across-quarter differences.

Column (1) reports the regression results based on the full sample of 15,528 firm-quarter observations, while column (2) presents results of the same specification except that we winsorize the managerial guidance and the difference between guidance and earnings at the 2.5% and 97.5 levels to mitigate potential bias driven by extreme observations. Furthermore, we repeat the same set of exercises (reported in columns (1) to (2)), by excluding all cases where the managerial guidance coincides with the prevailing median analysts' forecast (i.e., conforming cases), and show the respective results in columns (3) to (4).

The coefficient on Negative Guidance is positive and significant in columns (1), implying that the quality of guidance, which is inversely related to the magnitude of differences between guidance and realized earnings, is on average slightly lower on condition that the managerial guidance is negative. We worry that that is driven by outliers, but results in column (2) suggest that it is unlikely. The result, reported in column (3), remains the same, once we exclude conforming cases. The effect becomes insignificant, reported in column (4), if we exclude conforming cases and winsorize at 5%. In any case, the evidence does not favor the conjecture that positive guidance is less reliable. The managerial motives can be complex, often unobservable and unpredictable, which constitutes a source for guidance quality to be unreliable. That is the key motivation for our assumption that guidance quality is uncertain.

E. Proofs and Derivations

Proof of Lemma 1. The fact that forecast revisions are linear in guidance surprises directly follows from Equation (7). Further, forecast errors and forecast revisions are not correlated, because of rationality in noisy information expectation, that is, forecast errors are uncorrelated with any observables in the information set including forecast revisions. To demonstrate it mathematically, notice that

$$\begin{aligned}
\mathrm{FE}_i^{\mathrm{NI}} =& \frac{\tau_\theta}{\tau_\theta+\tau_z+\tau_y}\theta - \frac{\tau_z}{\tau_\theta+\tau_z+\tau_y}\iota_i - \frac{\tau_Y}{\tau_\theta+\tau_z+\tau_y}\eta; \\
\mathrm{FR}_i^{\mathrm{NI}} =& \frac{\tau_Y}{\tau_\theta+\tau_z+\tau_Y}\left(\frac{\tau_\theta}{\tau_\theta+\tau_z}\theta + \eta - \frac{\tau_z}{\tau_\theta+\tau_z}\iota_i\right).
\end{aligned}$$

The covariance between FE and FR is then given by

$$\mathrm{Cov}\left(\mathrm{FE}_i^{\mathrm{NI}}, \mathrm{FR}_i^{\mathrm{NI}}\right) \propto \frac{\tau_\theta}{\tau_\theta+\tau_z+\tau_y}\frac{\tau_\theta}{\tau_\theta+\tau_z}\frac{1}{\tau_\theta} - \frac{\tau_Y}{\tau_\theta+\tau_z+\tau_Y}\frac{1}{\tau_Y} + \frac{\tau_z}{\tau_\theta+\tau_z+\tau_y}\frac{\tau_z}{\tau_\theta+\tau_z}\frac{1}{\tau_z} = 0,$$

which completes the proof.

■

Derivation of Equation (12)-(14)**.** Denote $\delta \equiv \frac{\tau_y}{\tau_\theta+\tau_z+\tau_y}$. Then, it can be shown that

$$\begin{aligned}
\tilde{p}\left(\tau_y|F_{0i}^*, s_i; F_i\right) \equiv& \tilde{p}\left(\tau_y|z_{0i}, F_{0i}^*, s_i + F_{0i}^*; F_i\right), \\
=& \tilde{p}\left(\tau_y|z_{0i}, x_i, y\right), \\
\propto& \underbrace{\exp\left(-\lambda\left\{-F_i^2 + (2F_i+\beta)\left(F_{0i}^*+\delta s_i\right) - \left[\left(F_{0i}^*+\delta s_i\right)^2 + \frac{1-\delta}{\tau_\theta+\tau_z}\right]\right\}\right)}_{\phi'\left(\mathbb{E}_i^{\tau_y}[U(F_i,\theta)]\right)} \\
&\times \underbrace{p\left(F_{0i}^*\right)p\left(s_i|\tau_y\right)}_{=p\left(z_{0i},y|\tau_y\right)} p\left(\tau_y\right), \\
\propto& \exp\left(-\lambda\left[(2F_i+\beta)\,\delta s_i - \left(2F_{0i}^*\delta s_i + \delta^2 s_i^2 - \frac{\delta}{\tau_\theta+\tau_z}\right)\right]\right) p\left(s_i|\tau_y\right) p\left(\tau_y\right),
\end{aligned}$$

where the third line uses the fact that F_{0i}^* and s_i are independent with only the distribution of s_i affected by τ_y; and the last line drops all terms that are not a function of τ_y. Then, the optimality condition (10) can be compactly written as

$$F_i = F_{0i}^* + \kappa\left(F_{0i}^*, s_i, F_i\right)\cdot s_i,$$

where

$$\kappa\left(F_{0i}^{*}, s_i, F_i\right) = \int_{\Gamma_y} \left(\frac{\tau_y}{\tau_\theta + \tau_z + \tau_y}\right) \tilde{p}\left(\tau_y | F_{0i}^{*}, s_i; F_i\right) \mathrm{d}\tau_y.$$

■

Proof of Lemma 2. The log-likelihood ratio can be specifically written by:

$$\log\left(L\left(\tau_y\right)\right) = -\lambda s_i \left[2\left(F_i' - F_i\right)\left(\frac{\tau_y}{\tau_\theta + \tau_z + \tau_y}\right)\right] + \text{constant}.$$

Given the fact that $\tau_y / \left(\tau_\theta + \tau_z + \tau_y\right)$ increases in τ_y and that $F_i' - F_i > 0$, $L(\tau_y)$ decreases in τ_y, if and only if $s_i > 0$; and $L(\tau_y)$ increases in τ_y, if and only if $s_i < 0$. The lemma is shown. ■

Proof of Proposition 1. The optimality condition (10) is equivalent to (12):

$$F_i = F_{0i}^{*} + \underbrace{\left[\int_{\Gamma_y} \left(\frac{\tau_y}{\tau_\theta + \tau_z + \tau_y}\right) \tilde{p}\left(\tau_y | z_{0i}, y; F_i\right) d\tau_y\right]}_{\kappa} \cdot s_i. \tag{A7}$$

To obtain the second equality, we use the definition of F_{0i}^{*} and s_i and the definition of $\tilde{p}\left(\tau_y | z_{0i}, y; F_i\right)$ specified in the main text.

We first demonstrate that the right-hand side of (A7) is decreasing in F_i. Towards this end, we show

$$\begin{aligned}
\frac{1}{2}\frac{\partial\kappa}{\partial F_i} s_i = & \left\{ \int_{\Gamma_y} \left(\frac{\tau_y}{\tau_\theta + \tau_z + \tau_y}\right) \frac{\phi''\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]\right)}{\phi'\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]\right)} \frac{\partial \mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]}{\partial F_i} \tilde{p}\left(\tau_y | z_{0i}, y; F_i\right) \mathrm{d}\tau_y \right. \\
& \left. - \kappa \left[\int_{\Gamma_y} \frac{\phi''\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]\right)}{\phi'\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]\right)} \frac{\partial \mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]}{\partial F_i} \tilde{p}\left(\tau_y | z_{0i}, y; F_i\right) \mathrm{d}\tau_y\right] \right\} s_i, \\
= & \int_{\Gamma_y} \frac{\phi''\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]\right)}{\phi'\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]\right)} \left(\frac{\partial \mathbb{E}_i^{\tau_y}\left[U\left(F_i, \theta\right)\right]}{\partial F_i}\right)^2 \tilde{p}\left(\tau_y | z_{0i}, y; F_i\right) \mathrm{d}\tau_y < 0.
\end{aligned}$$

The first equality is obtained by using the definition of κ and the expression of $\partial\tilde{p}/\partial F_i$.

That is,

$$\begin{aligned}\frac{\partial \tilde{p}\left(\tau_y|z_{0i},y;F_i\right)}{\partial F_i} =& \frac{\phi''\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i,\theta\right)\right]\right)}{\phi'\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i,\theta\right)\right]\right)}\frac{\partial \mathbb{E}_i^{\tau_y}\left[U\left(F_i,\theta\right)\right]}{\partial F_i}\tilde{p}\left(\tau_y|z_{0i},y;F_i\right) \\ &- \tilde{p}\left(\tau_y|z_{0i},y;F_i\right)\left[\int_{\Gamma_y}\frac{\phi''\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i,\theta\right)\right]\right)}{\phi'\left(\mathbb{E}_i^{\tau_y}\left[U\left(F_i,\theta\right)\right]\right)}\frac{\partial \mathbb{E}_i^{\tau_y}\left[U\left(F_i,\theta\right)\right]}{\partial F_i}\tilde{p}\left(\tau_y|z_{0i},y;F_i\right)\mathrm{d}\tau_y\right].\end{aligned}$$

To get the second equality, we use the expression of $\partial \mathbb{E}_i^{\tau_y}\left[U\left(F_i,\theta\right)\right]/\partial F_i$. That is,

$$\frac{\partial \mathbb{E}_i^{\tau_y}\left[U\left(F_i,\theta\right)\right]}{\partial F_i} = \left(\frac{\tau_y}{\tau_\theta+\tau_z+\tau_y}-\kappa\right)s_i.$$

The third inequality holds given $\phi'(\cdot)>0$ and $\phi''(\cdot)<0$.

We then notice that κ is bounded between 0 and 1. Therefore, the right-hand side of Equation (12) goes to ∞, when F_i approaches $-\infty$; and it goes to $-\infty$ when F_i approaches ∞. Both existence and uniqueness are implied.

Next we show that the optimal response κ^* only depends on s_i. Observe that

$$\begin{aligned}\tilde{p}\left(\tau_y|F_{0i}^*,s_i;F_i\right) =& \tilde{p}\left(\tau_y|s_i;\kappa\right), \\ &\propto \exp\left(-\lambda\left[\beta\delta s_i+2\kappa\delta s_i^2-\left(\delta^2 s_i^2-\frac{\delta}{\tau_\theta+\tau_z}\right)\right]\right)p\left(s_i|\tau_y\right)p\left(\tau_y\right).\end{aligned}$$

To derive the first equality, we use the Equation (12) to replace F_i, and therefore F_{0i}^* drops out. Therefore, κ^* is the fixed point of the following condition:

$$\kappa^* = \int_{\Gamma_y}\left(\frac{\tau_y}{\tau_\theta+\tau_z+\tau_y}\right)\tilde{p}\left(\tau_y|s_i;\kappa^*\right)\mathrm{d}\tau_y.$$

Therefore, it is the case that κ^* is only a function of s_i. ■

Proof of Proposition 2. By using the definition F_i^*, the difference in the expected utilities is explicitly given by:

$$\begin{aligned}&\mathbb{E}^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^+\right),\theta\right)\right]-\mathbb{E}^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^-\right),\theta\right)\right] \\ =&2\beta\delta s_i^+ + \left[\left(\kappa^*\left(s_i^-\right)-\delta\right)^2-\left(\kappa^*\left(s_i^+\right)-\delta\right)^2\right]\left(s_i^+\right)^2.\end{aligned}$$

where $\delta\equiv\tau_y/\left(\tau_\theta+\tau_z+\tau_y\right)$. Let $T\left(\beta\right)\equiv\kappa^*\left(s_i^-\right)-\kappa^*\left(s_i^+\right)$.

Claim 1: If $\beta=0$, then $T\left(\beta\right)=0$.

We guess and verify that it holds that $\kappa^*\left(s_i^-\right)=\kappa^*\left(s_i^+\right)$. If this is true, we establish that $\mathbb{E}_i^{\tau_y}\left[U\left(F_i,\theta\right)\right]$ is symmetric in s_i: for any τ_y and any pair of $\left(s_i^-,s_i^+\right)$, we have:

$$\mathbb{E}^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^+\right),\theta\right)\right]=\mathbb{E}^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^-\right),\theta\right)\right].$$

In other words, for any τ_y, we have:

$$\phi'\left(\mathbb{E}_i^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^+\right),\theta\right)\right]\right) \quad = \quad \phi'\left(\mathbb{E}_i^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^-\right),\theta\right)\right]\right).$$

By the definition of κ, this implies:

$$\kappa^*\left(s_i^-\right) \quad = \quad \kappa^*\left(s_i^+\right).$$

which implies that $\beta=0$ is a solution to $T\left(\beta\right)=0$. Further, according to Proposition 1, both $\kappa^*\left(s_i^-\right)$ and $\kappa^*\left(s_i^+\right)$ are unique.

Claim 2: If $\beta\neq 0$, $T\left(\beta\right)\neq 0$.

Suppose towards a contradiction that there exists some $\beta'>0$, such that $T\left(\beta'\right)=0$. This implies that $\kappa^*\left(s_i^-\right)=\kappa^*\left(s_i^+\right)=\kappa'$. For any pair of $\left(s_i^-,s_i^+\right)$, we have:

$$\frac{\partial\log\left(\frac{\tilde{p}\left(\tau_y|F_{0i}^*,s_i^-;F_{0i}^*+\kappa's_i^-\right)}{\tilde{p}\left(\tau_y|F_{0i}^*,s_i^+;F_{0i}^*+\kappa's_i^+\right)}\right)}{\partial\tau_y}=\lambda\left(\frac{\partial\mathbb{E}_i^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^+;\kappa'\right),\theta\right)\right]}{\partial\tau_y}-\frac{\partial\mathbb{E}_i^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^-;\kappa'\right),\theta\right)\right]}{\partial\tau_y}\right)>0.$$

The last inequality is obtained by using the fact that

$$\mathbb{E}_i^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^+\right),\theta\right)\right]-\mathbb{E}_i^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^-\right),\theta\right)\right] \quad = \quad 2\beta'\delta\left(s_i^+-s_i^-\right)>0.$$

In other words, $\tilde{p}\left(\tau_y|F_{0i}^*,s_i^-;F_{0i}^*+\kappa's_i^-\right)$ first-order stochastically dominates $\tilde{p}\left(\tau_y|F_{0i}^*,s_i^+;F_{0i}^*+\kappa's_i^+\right)$. By the definition of κ, this implies:

$$\kappa^*\left(s_i^-\right)>\kappa^*\left(s_i^+\right).$$

A contradiction. Similarly, suppose towards a contradiction that there exists some $\beta'<0$ such that $T\left(\beta'\right)=0$. It implies that $\kappa^*\left(s_i^+\right)>\kappa^*\left(s_i^-\right)$. A contradiction. The claim is shown.

Claim 3: If β goes to ∞, $T\left(\beta\right)>0$.

When β goes to $\to\infty$, both $\kappa^*\left(s_i^-\right)$ and $\kappa^*\left(s_i^+\right)$ are bounded. Therefore,

$$\mathbb{E}_i^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^+\right),\theta\right)\right]-\mathbb{E}_i^{\tau_y}\left[U\left(F^*\left(F_{0i}^*,s_i^-\right),\theta\right)\right] \quad \to \quad 2\beta\delta\left(s_i^+-s_i^-\right)>0.$$

$\tilde{p}\left(\tau_y|F_{0i}^*,s_i^-;F_{0i}^*+\kappa's_i^-\right)$ first-order stochastically dominates $\tilde{p}\left(\tau_y|F_{0i}^*,s_i^+;F_{0i}^*+\kappa's_i^+\right)$, given

$\beta \to \infty$. Therefore, by the definition of κ, it implies that

$$\kappa^* \left(s_i^-\right) > \kappa^* \left(s_i^+\right).$$

That is, $T(\beta) > 0$. The claim is shown.

Claims 1 and 2 imply that $T(\beta)$ crosses zero once and only at $\beta = 0$. Combined with Claim 3, it further implies that $\beta T(\beta) \geq 0$, where the equality holds only when $\beta = 0$. The proposition is shown. ■

Proof of Proposition 3 . If forecasters are ambiguity neutral, the optimal forecasts are such that

$$F_i^* = F_{0i}^* + \left[\int_{\Gamma_y} \delta p \left(\tau_y | s_i\right) \mathrm{d}s_i\right] s_i,$$

where $\delta \equiv \tau_y / (\tau_\theta + \tau_z + \tau_y)$ and the posterior belief $p\left(\tau_y | s_i\right)$ is given by

$$p\left(\tau_y | s_i\right) \propto \sqrt{\delta} \exp\left(-\frac{1}{2}\left(\tau_\theta + \tau_z\right) s_i^2 \delta\right) p\left(\tau_y\right),$$

Taking the derivative of F_i^* w.r.t s_i leads to

$$\frac{\partial F_i^*}{\partial s_i} = \int_{\Gamma_y} \delta p\left(\tau_y | s_i\right) ds_i - \left(\tau_\theta + \tau_z\right) \mathrm{Var}\left(\delta | s_i\right) s_i^2,$$

where $\mathrm{Var}\left(\delta | s_i\right)$ denotes the conditional volatility of δ under probability density $p\left(\tau_y | s_i\right)$.

It is then straightforward to show that:

$$\lim_{|s_i| \to 0} \frac{\partial F_i^*}{\partial s_i} = \lim_{|s_i| \to 0} \int_{\Gamma_y} \delta p\left(\tau_y | s_i\right) \mathrm{d}s_i > 0.$$

Furthermore, when $|s_i| \to +\infty$, $p\left(\tau_y | s_i\right)$ converges to $p_\infty\left(\tau_y\right)$ and is given by:

$$p_\infty\left(\tau_y\right) \propto \sqrt{\delta} p\left(\tau_y\right).$$

Then it must be the case that $\lim_{|s_i| \to +\infty} \mathrm{Var}\left(\delta | s_i\right) s_i^2 \to +\infty$. Further using the fact that $\int_{\Gamma_y} \delta p\left(\tau_y | s_i\right) \mathrm{d}s_i$ is bounded above by $\delta_{\max}$, it is straightforward to demonstrate that

$$\lim_{|s_i| \to +\infty} \frac{\partial F_i^*}{\partial s_i} \to -\infty.$$

Finally, the symmetry of $F_i^* - F_{i0}^*$ around the origin directly follows from the fact that $\int_{\Gamma_y} \delta p\left(\tau_y | s_i\right) \mathrm{d}s_i$ is symmetric, since $p\left(\tau_y | s_i\right) = p\left(\tau_y | -s_i\right)$ for $\forall s_i \in \mathbb{R}$.

■

Proof of Proposition 4 . The objective function (4) under the maxmin criterion becomes:

$$\max_{F\in\mathbb{R}}\min_{\tau_y\in\Gamma_y}\mathbb{E}\left[-(F-\theta)^2+\beta\theta|z_i,y;\tau_y\right],$$

where Γ_y is the full support for τ_y. Let the upper bound be $\tau_y^{\max}$ and the lower bound be $\tau_y^{\min}$. For ease of notation, denote the subjective relative precision of guidance to be

$$\delta\equiv\frac{\tau_y}{\tau_\theta+\tau_z+\tau_y}.$$

and accordingly, it is bounded by

$$\delta_{\min}\equiv\frac{\tau_y^{\min}}{\tau_\theta+\tau_z+\tau_y^{\min}}\quad\text{and}\quad\delta_{\max}\equiv\frac{\tau_y^{\max}}{\tau_\theta+\tau_z+\tau_y^{\max}}.$$

To prove the proposition, we first characterize the optimal forecasting rule under the maxmin criterion. Then, we proceed to prove that $F_i^*-F_{0i}^*$ is non-decreasing in s_i.

First of all, it can be shown that

$$\bar{\theta}_{\tau_y}=F_{0i}^*+\delta s_i,\qquad\mathbb{E}_i\left[\theta^2|z_i,y;\tau_y\right]=\left(F_{0i}^*+\delta s_i\right)^2+(1-\delta)\left(\frac{1}{\tau_\theta+\tau_z}\right).$$

Then, the problem can be transformed into

$$\max_{\kappa\in\mathbb{R}}\min_{\delta\in\Delta}V(\kappa,\delta),$$

where $\Delta\equiv[\delta_{\min},\delta_{\max}]$ and the value function $V(\kappa,\delta)$ is given by

$$V(\kappa,\delta)\equiv-\left(F_{0i}^*+\kappa s_i\right)^2+\left[2\left(F_{0i}^*+\kappa s_i\right)+\beta\right]\left(F_{0i}^*+\delta s_i\right)-\left[\left(F_{0i}^*+\delta s_i\right)^2+(1-\delta)\frac{1}{\tau_\theta+\tau_z}\right],$$

where we have used the fact that $F=F_{0i}^*+\kappa s_i$. Notice that $V(\kappa,\delta)$ is quadratic in κ and δ. Also note that $V(\kappa,\delta)$ is concave in δ. Therefore, we have that for any $\kappa\in\mathbb{R}$:

$$\operatorname*{argmin}_{\delta\in\Delta}V(\kappa,\delta)\in\{\delta_{\min},\delta_{\max}\}.$$

Notice that

$$
\begin{aligned}
& V(\kappa, \delta_{\max}) - V(\kappa, \delta_{\min}) \\
= & (2\kappa s_i + \beta) s_i (\delta_{\max} - \delta_{\min}) + \frac{1}{\tau_\theta + \tau_z} (\delta_{\max} - \delta_{\min}) - s_i^2 \left(\delta_{\max}^2 - \delta_{\min}^2\right).
\end{aligned}
$$

It can then be shown that

$$
V(\kappa, \delta_{\max}) - V(\kappa, \delta_{\min}) > 0 \Leftrightarrow \kappa > T(s_i) \equiv \frac{(\delta_{\max} + \delta_{\min})}{2} - \frac{\beta s_i + \frac{1}{\tau_\theta + \tau_z}}{2 s_i^2}.
$$

In what follows, we characterize the optimal forecasting rule for three exclusive cases:

- If $\delta_{\min} > T(s_i)$, it can be shown that
  - when $\kappa \in (-\infty, T(s_i)]$, $\min_{\delta \in \Delta} V(\kappa, \delta) = V(\kappa, \delta_{\max})$. Hence, $\min_{\delta \in \Delta} V(\kappa, \delta)$ is increasing in κ.
  - when $\kappa > T(s_i)$, $\min_{\delta \in \Delta} V(\kappa, \delta) = V(\kappa, \delta_{\min})$. Hence, $\min_{\delta \in \Delta} V(\kappa, \delta)$ is first increasing in κ and then decreasing in κ. It achieves its maximum at $\kappa = \delta_{\min}$.

  Figure 5(a) graphically illustrates the value function under the worst case scenario when $\delta_{\max} < T(s_i)$. Therefore, it must be the case that the optimal $\kappa^* = \delta_{\min}$ when $\delta_{\min} > T(s_i)$.

- If $\delta_{\max} < T(s_i)$, it can be shown that
  - when $\kappa \in (-\infty, T(s_i)]$, $\min_{\delta \in \Delta} V(\kappa, \delta) = V(\kappa, \delta_{\max})$. Hence, $\min_{\delta \in \Delta} V(\kappa, \delta)$ is first increasing in κ and then decreasing in κ. It achieves its maximum at $\kappa = \delta_{\max}$.
  - when $\kappa \in [T(s_i), +\infty)$, $\min_{\delta \in \Delta} V(\kappa, \delta) = V(\kappa, \delta_{\min})$. Hence, $\min_{\delta \in \Delta} V(\kappa, \delta)$ is decreasing in κ.

  Figure 5(b) graphically illustrates the value function under the worst case scenario when $\delta_{\max} < T(s_i)$. Therefore, it must be the case that the optimal $\kappa^* = \delta_{\max}$ when $\delta_{\max} < T(s_i)$.

- If $\delta_{\min} < T(s_i) < \delta_{\max}$, it is then straightforward to show the following:
  - when $\kappa \in (-\infty, T(s_i)]$, $\min_{\delta \in \Delta} V(F, \delta) = V(F, \delta_{\max})$. Hence, $\min_{\delta \in \Delta} V(F, \delta)$ is increasing in κ.

– when $\kappa \in [T(s_i), +\infty)$, $\min_{\delta\in\Delta} V(F,\delta) = V(F,\delta_{\min})$. Hence, $\min_{\delta\in\Delta} V(F,\delta)$ is decreasing in κ.

Figure 5(c) graphically illustrates the value function under the worst case scenario when $\delta_{\min} < T(s_i) < \delta_{\max}$. Therefore, it must be the case that the optimal $\kappa^* = T(s_i)$ when $\delta_{\min} < T(s_i) < \delta_{\max}$.

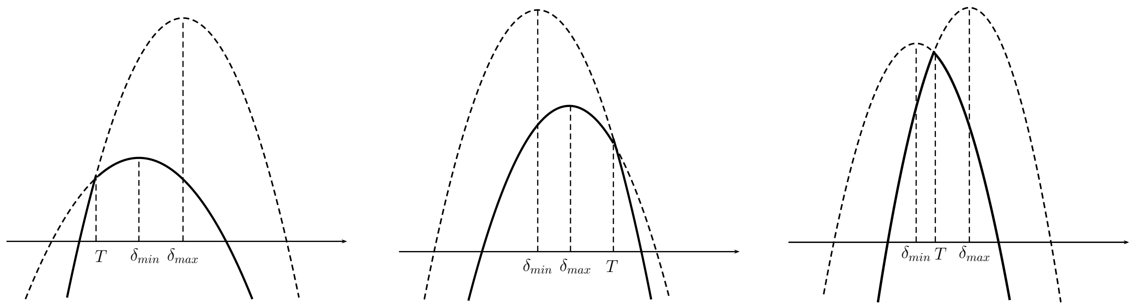

(a) $\delta_{\min} > T(s_i)$ (b) $\delta_{\max} < T(s_i)$ (c) $\delta_{\min} < T(s_i) < \delta_{\max}$

Figure A5. *The value function under the worst case scenario:* $\min_{\tau_y\in\Gamma_y} V(\kappa,\delta)$.

To summarize, we have the following optimal forecasting rule under the maxmin criterion:

$$\kappa^* = \begin{cases} \delta_{\min} & \text{if } \delta_{\min} > T(s_i); \\ \delta_{\max} & \text{if } \delta_{\max} < T(s_i); \\ T(s_i) & \text{otherwise.} \end{cases} \tag{A8}$$

Or equivalently,

$$F^* - X_i = \begin{cases} \delta_{\min} s_i & \text{if } \delta_{\min} > T(s_i); \\ \delta_{\max} s_i & \text{if } \delta_{\max} < T(s_i); \\ T(s_i)\, s_i & \text{otherwise.} \end{cases} \tag{A9}$$

Note that $T(s_i)\, s_i$ is always increasing in s_i. Therefore, given the continuity of $F_i^* - F_{0i}^*$ with respect to s_i, it must be the case that $F_i^* - F_{0i}^*$ is non-decreasing in s_i.

■

Derivation of Equation (15)**.** Following the definition of $\hat{b}_1(s_m,\epsilon)$, we have

$$\begin{aligned}
\hat{b}_1(s_m,\epsilon) \equiv & \frac{\text{Cov}(\text{FE}_i,\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))}{\text{Var}(\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))}, \\
= & \frac{\text{Cov}(\theta-F_{0i}-\text{FR}_i,\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))}{\text{Var}(\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))}, \\
= & -1+\frac{\text{Cov}(\theta,\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))}{\text{Var}(\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))}-\frac{\text{Cov}(F_{0i},\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))}{\text{Var}(\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))}, \\
= & -1+\frac{\text{Cov}(\theta,\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))}{\text{Var}(\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))},
\end{aligned}$$

where in the last equality we use the fact that $\text{Cov}(F_{0i},\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))=0$.

To see why this is the case, notice that the unconditional covariance between initial forecasts F_{0i} and guidance surprise s_i is zero: $\text{Cov}(F_{0i},s_i)=0$. Further using the fact that both F_{0i} and s_i are normally distributed, we know that initial forecasts F_{0i} and guidance surprise s_i are independent. Moreover, since forecast revisions FR_i is a (non-linear) function of guidance surprise s_i only, it is then straightforward to show that $\text{Cov}(F_{0i},\text{FR}_i|s_i\in\mathbb{I}(s_m,\epsilon))=0$.

For any $s_i\in\mathbb{I}(s_m,\epsilon)$, a first-order approximation of FR_i around the $s_i=s_m$ implies

$$\begin{aligned}
\text{FR}_i &\approx \kappa(s_m)s_m+\left[\kappa(s_m)+\kappa'(s_m)s_m\right](s_i-s_m), \\
&= -\kappa'(s_m)s_m^2+\left[\kappa(s_m)+\kappa'(s_m)s_m\right]s_i.
\end{aligned}$$

Substituting it in the expression of $\hat{b}_1(s_m)$, we obtain:

$$\begin{aligned}
\hat{b}_1(s_m) &\equiv \lim_{\epsilon\to 0}\hat{b}_1(s_m,\epsilon), \\
&\approx -1+\lim_{\epsilon\to 0}\left(\frac{\text{Cov}(\theta,s_i|s_i\in\mathbb{I}(s_m,\epsilon))}{\text{Var}(s_i|s_i\in\mathbb{I}(s_m,\epsilon))}/\left[\kappa(s_m)+\kappa'(s_m)s_m\right]\right), \\
&= -1+\frac{\kappa^{\text{RE}}}{\kappa(s_m)+\kappa'(s_m)s_m},
\end{aligned}$$

where we use Equation (A10) to obtain at the last equality.

In the following, we provide an analysis in the special of noisy rational expectations. Under rationality, it can be shown that

$$\text{Cov}\left(\text{FE}_i^{\text{NI}},\text{FR}_i^{\text{NI}}\right)=0.$$

Given FE_i^{NI} and FR_i^{NI} are normally distributed, it can be shown that FE_i^{NI} and FR_i^{NI} are independent, which implies that $\text{Cov}\left(\text{FE}_i^{\text{NI}},\text{FR}_i^{\text{NI}}|s_i\in\mathbb{I}(s_m,\epsilon)\right)=0$. Therefore, we have the following:

$$\hat{b}_1^{\mathrm{RE}}(s_m,\epsilon) = -1 + \frac{1}{\kappa^{\mathrm{RE}}}\frac{\mathrm{Cov}(\theta, s_i | s_i \in \mathbb{I}(s_m,\epsilon))}{\mathrm{Var}(s_i | s_i \in \mathbb{I}(s_m,\epsilon))} = 0, \tag{A10}$$

where we use the fact that $\mathrm{FR}_i^{\mathrm{NI}} = \kappa^{\mathrm{RE}} s_i$ with κ^{RE} given by Equation (6). ■

Proof of Proposition 5. To prove part (i) of the proposition, based on the approximation of Equation (15), it is sufficient to prove that

$$\lim_{s_i\to 0} \frac{\mathrm{d}\kappa(s_i) + \kappa'(s_i)s_i}{\mathrm{d}s_i} = 2\lim_{s_i\to 0}\kappa'(s_i) < 0.$$

Notice that $\kappa(s_i) = \int_{\Gamma_y}\left(\frac{\tau_y}{\tau_\theta+\tau_z+\tau_y}\right)\tilde{p}(\tau_y|s_i;\kappa(s_i))\,\mathrm{d}\tau_y$ where the distorted posterior $\tilde{p}(\tau_y|s_i;\kappa(s_i))$ is such that

$$\tilde{p}(\tau_y|s_i;\kappa) \propto \exp\left(-\lambda\left[\beta\delta s_i + 2\kappa\delta s_i^2 - \left(\delta^2 s_i^2 - \frac{\delta}{\tau_\theta+\tau_z}\right)\right]\right) p(s_i|\tau_y)\,p(\tau_y).$$

Some algebra implies that

$$\begin{aligned}\frac{\mathrm{d}\kappa(s_i)}{\mathrm{d}s_i} &= \int_{\Gamma_y}\delta\frac{\mathrm{d}\tilde{p}(\tau_y|s_i;\kappa)}{\mathrm{d}s_i}\mathrm{d}\tau_y,\\ &= -\lambda(\beta + 4\kappa(s_i)s_i + (\tau_\theta+\tau_z)s_i)\widetilde{\mathrm{Var}}_i(\delta) + 2\lambda s_i\widetilde{\mathrm{Cov}}_i\left(\delta,\delta^2\right) - 2\lambda s_i^2\widetilde{\mathrm{Var}}_i(\delta)\frac{\mathrm{d}\kappa(s_i)}{\mathrm{d}s_i},\end{aligned}$$

where $\widetilde{\mathrm{Var}}_i(\cdot)$ and $\widetilde{\mathrm{Cov}}_i(\cdot)$ denote the variance and covariance under the distorted posterior $\tilde{p}(\tau_y|s_i;\kappa)$. It is then straight-forward to see that

$$\lim_{s\to 0}\frac{\mathrm{d}\kappa(s_i)}{\mathrm{d}s_i} = \lim_{s\to 0} -\lambda\beta\widetilde{\mathrm{Var}}_i(\delta) < 0,$$

which completes the proof of part (i).

To prove part (ii) of the proposition, notice that when $|s_i|$ goes to infinity, the distorted belief $\tilde{p}(\tau_y|s_i;\kappa)$ is degenerate that puts probability 1 on the lowest possible precision for manager guidance:

$$\lim_{|s_i|\to\infty}\kappa(s_i) = \delta_{\min} \equiv \frac{\tau_y^{\min}}{\tau_\theta+\tau_z+\tau_y^{\min}}.$$

Hence we have that $\lim_{|s_i|\to\infty}\kappa'(s_i) = 0$. It is then can be shown that

$$\lim_{|s_i|\to\infty}\hat{b}_1(s_i) = -1 + 1/\delta_{\min}.$$

Further using the fact that $\hat{b}_1(0) = \kappa(0) > \delta_{\min}$, it is straight-forward to prove that

$$\hat{b}_1(0) < \lim_{|s_i| \to \infty} \hat{b}_1(s_i).$$

■

Proof of Proposition 6. It is straight-forward to show that the fundamental θ and the initial forecasts F_{0i}^* are both unconditionally mean zero:

$$\mathbb{E}[\theta] = 0, \qquad \mathbb{E}[F_{0i}^*] = 0.$$

Furthermore, observe that

$$\begin{aligned}
\mathbb{E}[\mathrm{FR}_i] &= \int_{-\infty}^{+\infty} \kappa(s_i)\, s_i p(s_i)\, \mathrm{d}s_i, \\
&= \int_{-\infty}^{0} \kappa(s_i)\, s_i p(s_i)\, \mathrm{d}s_i + \int_{0}^{+\infty} \kappa(s_i)\, s_i p(s_i)\, \mathrm{d}s_i, \\
&= -\int_{0}^{+\infty} \kappa(-s_i)\, s_i p(s_i)\, \mathrm{d}s_i + \int_{0}^{+\infty} \kappa(s_i)\, s_i p(s_i)\, \mathrm{d}s_i, \\
&= \int_{0}^{+\infty} [\kappa(s_i) - \kappa(-s_i)]\, s_i p(s_i)\, \mathrm{d}s_i < 0,
\end{aligned}$$

where $p(s_i)$ denotes the probability density of guidance surprises in the objective environment. To arrive at the third equality, we use the fact that $p(s_i)$ is symmetric and the last inequality follows Proposition 2 such that $\kappa(s_i) < \kappa(-s_i)$.

Finally, using the fact that $\mathrm{FE}_i = \theta - F_{0i}^* - \mathrm{FR}_i$, it is straight-forward to prove that $\mathbb{E}[\mathrm{FE}_i] > 0$. ■

Proof of Proposition 7. The optimal forecast F_i^L is globally monotone in the signal is a sufficient condition for the global monotonicity. According to (A5), the optimal forecasts can be expressed as

$$\mathrm{F}_i^L = \mathbb{E}[\theta|\mathcal{I}_i] + \frac{1-2\alpha}{\alpha} \int_{-\infty}^{\mathrm{F}_i^L} \left(\theta - \mathrm{F}_i^L\right) f(\theta|\mathcal{I}_i)\, \mathrm{d}\theta. \tag{A11}$$

Assume that $\mathcal{I}_i = \{x_i\}$ with $x_i \sim N\left(\theta, \sigma_x^2\right)$, where the fundamental $\theta \sim N\left(0, \sigma_\theta^2\right)$.

In what follows, we prove that $\frac{\mathrm{dF}_i^L}{\mathrm{d}x_i} > 0$. Taking total derivative w.r.t x_i on both sides of (A11) and re-arranging leads to

$$\begin{aligned}\frac{\mathrm{dF}_i^L}{\mathrm{d}x_i} =& \frac{\mathrm{d}\mathbb{E}\left[\theta|x_i\right]}{\mathrm{d}x_i} - \frac{1-2\alpha}{\alpha}\,\mathrm{P}\left(\theta < \mathrm{F}_i^L \middle| x_i\right)\frac{\mathrm{dF}_i^L}{\mathrm{d}x_i} \\ &+ \frac{1-2\alpha}{\alpha}\left(\frac{1}{\sigma_\theta^2}+\frac{1}{\sigma_x^2}\right)\int_{-\infty}^{\mathrm{F}_i^L}\left(\theta-\mathrm{F}_i^L\right)\left(\theta-\mathbb{E}\left[\theta|x_i\right]\right) f\left(\theta|x_i\right)\mathrm{d}\theta,\end{aligned}$$

where $\mathrm{P}\left(\theta < \mathrm{F}_i^L \middle| x_i\right)$ denotes the conditional probability of negative forecast error. Using the fact that $\mathrm{F}_i^L < \mathbb{E}\left[\theta|x_i\right]$, we can prove that

$$\frac{1-2\alpha}{\alpha}\left(\frac{1}{\sigma_\theta^2}+\frac{1}{\sigma_x^2}\right)\int_{-\infty}^{\mathrm{F}_i^L}\left(\theta-\mathrm{F}_i^L\right)\left(\theta-\mathbb{E}\left[\theta|x_i\right]\right) f\left(\theta|x_i\right)\mathrm{d}\theta > 0.$$

Combined with the fact that $\frac{\mathrm{d}\mathbb{E}[\theta|x_i]}{\mathrm{d}x_i} > 0$, it is straight-forward to see that

$$\frac{\mathrm{dF}_i^L}{\mathrm{d}x_i} > 0.$$

∎