

The strategy of professional forecasting*

Marco Ottaviani^{a†}, Peter Norman Sørensen^b

^a *London Business School, Regent's Park, London NW1 4SA, England*

^b *University of Copenhagen, Studiestræde 6, DK-1455 Copenhagen K, Denmark*

August 2005

Abstract

We develop and compare two theories of professional forecasters' strategic behavior. The first theory, *reputational cheap talk*, posits that forecasters endeavor to convince the market that they are well informed. The market evaluates their forecasting talent on the basis of the forecasts and the realized state. If the market expects forecasters to report their posterior expectations honestly, then forecasts are shaded toward the prior mean. With correct market expectations, equilibrium forecasts are imprecise but not shaded. The second theory posits that forecasters compete in a *forecasting contest* with pre-specified rules. In a winner-take-all contest, equilibrium forecasts are excessively differentiated.

Keywords: Forecasting; Reputation; Cheap talk; Contest; Exaggeration

JEL Classification: D82; G29

*We thank Dean Croushore for essential feedback in the early stages of this project; Owen Lamont for generously providing part of the data we use for Fig. 1 and for very useful suggestions; Marco Battaglini for excellent discussion of the paper; Eric Zitzewitz for extremely helpful feedback; an anonymous referee for many constructive suggestions; Andrew Chesher, Christian Dustmann, Douglas Gale, Peter R. Hansen, Harrison Hong, Boyan Jovanovic, Jon Levin, John Morgan, Kent Osband, Mike Peters, Rafael Repullo, Ilya Segal, Daniel Seidmann, Shiva Shivakumar, Hyun Shin, Raman Uppal, Pietro Veronesi, Gianluca Violante, Xavier Vives, and Jeffrey Zwiebel for useful input; and seminar audiences at Aarhus, AEA (New Orleans), Bank of Italy, Bergen, Berlin, Bonn, Bristol, Brussels, CEMFI, Cambridge, Copenhagen, Danish Economic Society 2000, EEA (Stockholm), Erice, ECB, EUI, Gerzensee, Helsinki, INSEAD, Kellogg, Koç University, LBS, LSE, Mannheim, Munich, Paris, Pompeu Fabra, Princeton, Royal Holloway, Southampton, Stanford, Stern, Surrey, UCL, and Western Ontario for comments.

[†]Corresponding author. Tel.: +44 20 7706 6987; fax: +44 20 7402 0718. E-mail address: mottaviani@london.edu (M. Ottaviani).

1. Introduction

Professional forecasts guide market participants and inform their expectations about future economic conditions. Given the importance of this role and the potential rewards of accurate forecasting, we might expect that professional forecasters maximize their accuracy by truthfully releasing all their information. As Keane and Runkle (1998) report, “since financial analysts’ livelihoods depend on the accuracy of their forecasts . . . , we can safely argue that these numbers accurately measure the analysts’ expectations.”

However, many commentators argue that forecasters might *strategically* misreport their information, even when they are not interested in manipulating the investment decisions of their target audience. For example, Croushore (1997) suggests that “some [survey] participants might shade their forecasts more toward the consensus (to avoid unfavorable publicity when wrong), while others might make unusually bold forecasts, hoping to stand out from the crowd.” The importance of the microeconomic incentives of forecasters and analysts is stressed by a number of recent empirical studies, such as Ehrbeck and Waldmann (1996), Graham (1999), Hong, Kubik, and Solomon (2000), Lamont (2002), Welch (2000), and Zitzewitz (2001a).

This paper develops a theoretical framework for analyzing the strategic behavior of professional forecasters. In the context of a simple model, we address how reputation and competition affect the reporting incentives of forecasters. Moreover, we study the reaction of market prices to the release of strategic forecasts.

The basic ingredients of our model are best introduced by Fig. 1, which depicts *Business Week Investment Outlook*’s yearly GNP growth forecasts for the period 1972 to 2004. The forecast data come from a survey of professional forecasters that the magazine *Business Week* runs at the end of each year. Because there is substantial dispersion in the individual forecasts, our model assumes that forecasters are privately informed. In addition, there is substantial variation in the forecast dispersion across years. Hence, we use both the quality of private information and the precision of the prior distribution as parameters in our model. The figure also shows, for the same period, the realized GNP growth rates released by the *Bureau of Economic Analysis*. The market uses these realizations to evaluate the quality of the forecasters’ information.¹

¹A perennial problem in evaluating forecasts is that data on realized values are revised over time. Fig. 1 uses the latest revisions available from the *Bureau of Economic Analysis*. See Section 7.1 for more on this matter.




Fig. 1. The triangles represent *Business Week Investment Outlook's* individual forecasts of annual real GNP growth rate for the period 1972 to 2004. The connected circles represent the realizations, obtained from the *Bureau of Economic Analysis*.

We formulate two theories of strategic forecasting and contrast them with the benchmark case of nonstrategic forecasting. To facilitate comparison, we adopt a unified, tractable statistical model. The *state* has a normal prior distribution and the forecasters' *signals* are normally distributed around the state. After the forecasters simultaneously release their forecasts, the state is publicly observed. To isolate the effect of the *professional* objectives of forecasters who predict the future evolution of economic or financial variables, we assume that the forecasters cannot affect the distribution of the state variable and that they do not care about the investment decisions taken on the basis of their forecasts.

The benchmark case we consider is that of a forecaster rewarded according to the *absolute* accuracy of her prediction. A nonstrategic forecaster honestly reports the posterior expectation of the state, which is a weighted average of the signal and the mean of the prior distribution. Hence, when the state turns out to be above the prior mean, the honest forecast tends to be lower than the realized state.² Being equal to conditional expecta-

²The popular press often takes this empirically documented negative correlation of the forecast errors with the realized state as evidence of herd behavior. The academic literature avoids this misconception and focuses instead on the correlation between forecasts and their errors.

tions, honest forecasts are “rational” in Muth’s (1961) sense and are uncorrelated with their forecast errors. The two theories of strategic forecasting that we develop here have different implications for the correlation of forecasts and forecast errors.

First, according to our theory of *reputational cheap talk*, forecasters wish to foster their reputation for being well informed, while the market uses forecasts and the realized state to evaluate forecasting talent. In the baseline model, we further assume that forecasters do not have private information about their own talent prior to receiving their private signals.³ Contrary to naive intuition, we show that honest forecasting cannot occur in equilibrium. If the market naively expects honest forecasting, the forecaster has an incentive to bias the forecast toward the prior mean in order to be perceived to be better informed. More generally, we show that if the market can invert from forecasts to signals, the best *reputational deviation* for the forecaster is to pretend to have received a signal equal to the posterior expectation of the state. We find, therefore, that while concern for reputation drives forecasters to herd on the prior belief, this incentive is self defeating. In equilibrium, the market must have rational expectations about the behavior of forecasters. As a consequence, equilibrium forecasts cannot be perfectly inverted, i.e., there is no fully revealing equilibrium. We conclude that *reputational equilibrium* forecasts are not shaded, but rather are systematically less precise than if forecasters were not strategic. Paradoxically, the desire of analysts to be perceived as good forecasters turns them into poor forecasters.⁴

Our second theory of strategic behavior posits that forecasters compete in a *forecasting contest* with pre-specified rules. For instance, in competitions such as the semi-annual *Wall Street Journal (WSJ) Forecasting Survey* (macroeconomic variables), *WSJ All Star Analysts* (earnings), and *WSJ Best on the Street* (stock picking) forecasters are often ranked by their *relative* accuracy. We find that reporting the best predictor of the state (the posterior expectation conditional on the signal observed) is not an equilibrium in the tournament. With an infinite number of forecasters, the equilibrium strikes a balance between two contrasting forces. On the one hand, an individual forecaster has an incentive to report the honest forecast, which is most likely to be on the mark. On the other hand, a forecaster gains from moving away from the prior mean because the farther the state is from the prior mean, the lower the number of forecasters that correctly guess the state. In equilibrium, forecasters differentiate their predictions from those of competitors by putting greater weight on their private signals than they would in an honest report of

³We discuss the importance of this assumption in Section 4.4.

⁴Ironically, *The Economist* (“Dustmen as Economic Gurus,” June 3, 1995) reports the surprisingly good performance of a sample of London garbagemen in forecasting key economic variables.

the posterior expectations. Yet, in principle, rational market participants can invert the equilibrium strategy to recover the forecasters' information and thereby construct accurate expectations about the state.

The presence of strategic behavior raises questions about the interpretation and use of professional forecasts to test the predictions of theories on how agents' decisions depend on expectations. Empirical studies often use professional forecasts as proxies for unobservable market expectations. However, to the extent they are taken at face value, strategic forecasts do not reflect these expectations.

The two theories we develop have implications for both the dispersion of the forecasts and the correlation between forecasts and forecast errors. In the symmetric equilibrium of our forecasting contest, the forecast error is positively correlated with the forecast, but the correlation is negative in the reputational deviation. The reputational equilibrium forecast is uncorrelated with its error, but is not efficient. According to recent empirical work (Zitzewitz, 2001a; Bernhardt, Campello, and Kutsoati, 2005), forecast errors exhibit a strong positive correlation with forecasts, consistent with our contest theory.

The paper is organized as follows. Section 2 sets up the baseline model. Section 3 discusses the benchmark case of honest forecasting. Section 4 introduces the reputational cheap talk theory. Section 5 develops the forecasting contest theory. Section 6 compares the empirical implications of the different theories. Section 7 extends the baseline model in two directions that are relevant for empirical work. Section 8 derives implications about the stock price reaction to the arrival of a forecast. Section 9 concludes.

2. Model

Our baseline model considers n forecasters who simultaneously issue forecasts on an uncertain state of the world. We assume that it is common prior belief that the state x is normally distributed with mean μ and precision ν , i.e., $x \sim N(\mu, 1/\nu)$. Each forecaster i observes the private signal $s_i = x + \varepsilon_i$. Conditional on state x , signals s_i are independently normally distributed with mean x and precision τ_i , i.e., $s_i|x \sim N(x, 1/\tau_i)$. Forecaster i 's observation of signal s_i leads to a normal posterior belief on the state with mean $E(x|s_i) = (\tau_i s_i + \nu \mu) / (\tau_i + \nu)$ and precision $\tau_i + \nu$ (DeGroot, 1970). We denote the density of this posterior distribution by $q_i(x|s_i)$.

Because actual forecasts are typically dispersed, we allow forecasters to have private information about the state. If forecasters share a common prior, observe the same public information, and honestly report their expectations, they should make identical forecasts.⁵ To avoid introducing a bias against honest forecasting, we posit that forecasters

⁵Forecast dispersion could also be due to heterogeneous prior beliefs or to different models. According

are endowed with some private information about the state.⁶ Indeed, the presence of heterogeneous private information provides a rationale for the market to reward forecaster accuracy.

This model abstracts from the strategic incentives relevant to partisan forecasters, whose payoff instead depends on the investment decisions made on the basis of their forecasts.⁷ Our baseline model can be applied to situations in which the state cannot be affected by the forecasts and yet can be meaningfully forecasted and later observed. In Section 8.2, we extend the model to consider the additional strategic incentives that are present when forecasters are concerned with the effect of their forecasts on the state.

3. Honest forecasting

Forecasters are presumed honest, unless proven strategic. As Keane and Runkle (1990) argue, “professional forecasters . . . have an economic incentive to be accurate. Because these professionals report to the survey the same forecasts that they sell on the market, their survey responses provide a reasonably accurate measure of their expectations.”

Our benchmark forecast is the honest report of the Bayesian posterior expectation,

$$h_i(s_i) = E(x|s_i) = \left(\frac{\tau_i}{\tau_i + \nu}\right) s_i + \left(\frac{\nu}{\tau_i + \nu}\right) \mu, \quad (1)$$

as most empirical investigations assume. In the normal model, this posterior expectation minimizes the mean of any symmetric function of the forecast error, such as the mean squared error (Bhattacharya and Pfleiderer, 1985).

The honest forecast can explain several features of the data. First, forecast dispersion is a consequence of private information.

Second, as Fig. 1 illustrates, forecasts tend to be less volatile than realizations. This fact does not indicate the presence of herd behavior. The realization $x = h_i + (x - h_i)$

to industry participants, while forecasters seem to have access to the same pool of public data, they interpret the data differently depending on their model. Indeed, Kandel and Pearson (1995) and Kandel and Zilberfarb (1999) find empirical support for heterogeneous processing of public information. Note that the private information of the forecasters could be due to their private knowledge of the model they use to process the available public information.

⁶In their classic study on the rationality of forecasts using data from the NBER-ASA survey of professional forecasters (which is now called the Survey of Professional Forecasters), Keane and Runkle (1990) find that differences in individual forecasts cannot be explained by publicly available information. They infer that differences in forecasts are due to asymmetric information, but this conclusion rests on the maintained assumption of honest forecasting. The observed forecast dispersion might also be the outcome of strategic behavior.

⁷For empirical evidence of equity analyst *partisan bias* see, for example, Michaely and Womack (1999), Hong and Kubik (2003), and references therein. For theoretical investigations, see Crawford and Sobel (1982) and Morgan and Stocken (2003).

is necessarily more volatile than the honest forecast h_i when the forecasters' information is noisy, $V(x) = V(h_i) + V(x - h_i) > V(h_i)$. Equivalently, the realization x is negatively correlated with the forecast error $h_i - x = \nu(\mu - x) / (\tau_i + \nu) + \tau_i \varepsilon_i / (\tau_i + \nu)$. In contrast, the shock $h_i - E(h_i|x) = \tau_i \varepsilon_i / (\tau_i + \nu)$, is uncorrelated with x and with the shocks of other forecasters. The expected forecast $E(h_i|x)$ is commonly estimated by the consensus forecast, which is equal to the unweighted average of forecasts $\bar{f} = \sum_{i=1}^n f_i/n$ across all forecasters.

Third, forecasts are more dispersed and less accurate in years characterized by relatively little public information. This observation is consistent with a finding that Zarnowitz and Lambros (1987) report with respect to the ASA-NBER survey of professional forecasters. In addition to point forecasts, this survey initially asked forecasters to report probability distributions. Zarnowitz and Lambros document that forecast dispersion is positively correlated with a measure of forecast uncertainty. Likewise, in the data of our Fig. 1, when regressing the standard deviation of the forecasts on the absolute error of the consensus forecast, we find a coefficient of 0.105 with standard error 0.058. There is a negative correlation between accuracy and dispersion.

Finally, the honest forecast $h_i(s_i)$ has the key statistical property that it is uncorrelated with its forecast error $h_i(s_i) - x$, since

$$E\{E(x|s_i)[E(x|s_i) - x]\} = E\{E(x|s_i)E[E(x|s_i) - x|s_i]\} = 0. \quad (2)$$

According to this orthogonality property, the honest forecast does not carry information about its forecast error. While orthogonality may seem a necessary property of rational forecasts, such is not the case. For example, asymmetric scoring rules generally result in forecasts that violate the orthogonality property, as Granger (1969) and Zellner (1986) note.⁸ Instead, in our model we maintain symmetry and examine whether strategic factors lead rational players to make nonorthogonal forecasts.

While we use a Bayesian framework, we briefly consider the benchmark case of *classical* forecasting. In this case, forecaster i 's maximum likelihood estimator (MLE), which maximizes $g_i(s_i|x_i)$ over x_i , is s_i . The maximum likelihood forecast violates the orthogonality property, since $E[s_i(s_i - x)] = E[(x + \varepsilon_i)\varepsilon_i] = 1/\tau_i > 0$. The MLE can also be seen as resulting from Bayesian updating when the prior distribution on the state is the improper uniform distribution on the real line. However, forecasters typically share some pre-existing information about the variable to be predicted, as confirmed by a number of

⁸Granger (1999) defines generalized forecast errors for any given loss function and notes that these errors satisfy orthogonality.

empirical studies (e.g., Welch, 2000). The presence of prior information drives the results of our two theories.

4. Reputational cheap talk

Forecasters are subject to the informal (or subjective) evaluation of financial and labor markets. Forecasters who appear to have access to better information can command higher compensation and enjoy more favorable career prospects. The theory we develop in this section captures the fact that the market uses all the publicly available information to evaluate the forecaster talent. This information typically includes the forecasts and the ex post realization of the state. In turn, the forecasters want the market to believe that they are highly talented.

Reputational forecasting is a game of cheap talk (Crawford and Sobel, 1982), since a forecaster's payoff depends on the forecast released only indirectly, through the evaluation performed by the market. Our reputational cheap talk game builds on elements first introduced by Holmström (1999) and Scharfstein and Stein (1990). While Scharfstein and Stein consider a dynamic game in which better informed forecasters have conditionally more correlated signals, in this paper we focus on the static game with conditionally independent signals. Our model (further developed by Ottaviani and Sørensen, 2005) offers a new theory of herd behavior that does not rely on the presence of multiple forecasters. In the context of the normal model we analyze, we are able to obtain a sharp characterization of the deviation incentives and of the distributional properties of the equilibrium forecasts.

4.1. Model

We extend the information structure of Section 2 by introducing the parameter $t_i > 0$, which represents the unknown talent of forecaster i . All talents t_i and the state x are statistically independent. The market and the forecasters share both the common prior on the state, $x \sim N(\mu, 1/\nu)$, and the nondegenerate prior belief $p_i(t_i)$ on the talent. Conditional on x and t_i , signal s_i is generated by a symmetric location experiment with scale parameter t_i . The signal's density is $\tilde{g}(s_i|x, t_i) = t_i \hat{g}(t_i|s_i - x)/2$, in which \hat{g} is a density on $[0, \infty)$. We retain the assumption that $s_i|x \sim N(x, 1/\tau)$, so the primitives \hat{g} and p are restricted to satisfy $g_i(s_i|x) = \int_0^\infty \tilde{g}(s_i|x, t_i) p_i(t_i) dt_i$. Greater talent is associated with smaller signal errors, i.e., the likelihood ratio $\tilde{g}(s_i|x, t_i) / \tilde{g}(s_i|x, t'_i)$ is increasing in $|s_i - x|$ when $t_i < t'_i$.⁹

⁹As noted by Lehmann (1955, Example 3.3), this property is equivalent to the log-concavity in a of $\hat{g}(e^a)$. Equivalently, the elasticity $\varepsilon \hat{g}'(\varepsilon) / \hat{g}(\varepsilon)$ is decreasing in ε . The example $\hat{g}(\varepsilon) = K_1 \exp(-\varepsilon^4/12)$

We simplify the game by eliminating all payoff interactions among the forecasters. In particular, we assume that the private signals s_i are independent conditional on x and t_1, \dots, t_n , so that forecaster i cannot signal anything to the market about t_j for $j \neq i$. In addition, we assume that forecaster i 's payoff function does not depend on the posterior beliefs about t_j for $j \neq i$. Hence, we can focus on a single forecaster's problem in isolation and remove the subscript i .

The reputational cheap talk game proceeds as follows. First, the forecaster observes the private signal s and issues the forecast (or message) m . Second, the market observes the true state x and then uses (m, x) to update the belief $p(t)$ about the forecaster's talent.

To update the beliefs about the forecaster's talent, the market formulates a *conjecture* about the forecaster's strategy mapping s into m and derives the message distribution denoted by $\varphi(m|x, t)$. The market uses Bayes' rule to calculate the posterior reputation $p(t|m, x) = \varphi(m|t, x)p(t)/\varphi(m|x)$, where $\varphi(m|x) = \int_0^\infty \varphi(m|t, x)p(t) dt$. We assume that the forecaster's payoff from reputation $p(t|m, x)$ is given by the expected value expression $W(m|x) = \int_0^\infty u(t)p(t|m, x) dt$. Since an expert with greater talent t receives more accurate signals, we assume that u is strictly increasing. When reporting the message m , the forecaster does not yet know the state x , but believes it to be distributed according to $q(x|s)$. The forecaster then chooses the message m that maximizes the expected value $U(m|s) = \int_{-\infty}^\infty W(m|x)q(x|s) dx$.

4.2. Deviation

We first show that truthtelling cannot be an equilibrium. Consider what happens when the market conjectures that the forecaster reports the honest $m = h(s)$. As Fig. 2 depicts, a forecaster with signal $s > \mu$ believes that the state is normally distributed with mean $E(x|s)$, a weighted average of the prior mean μ and the signal s . If the forecaster were to honestly report $h(s) = E(x|s)$, the market would invert this strategy and infer the true signal $s = h^{-1}(h(s)) > E(x|s)$. Below, we show formally that a signal s closer to the state x is better news about the talent t and thus results in a higher forecaster payoff. To minimize the average distance between the signal inferred by the market and the best predictor of the state, the forecaster wishes to be perceived as having signal $\hat{s} = E(x|s)$. We conclude that if the market believes that the forecast truthfully reflects the forecaster's posterior expectation, then the forecaster deviates by reporting $d(s) = h(h(s)) = E[x|\hat{s} = E(x|s)]$. In this deviation, forecasters are biased towards the prior mean.

and $p_i(t_i) = K_{2i}\tau_i^{3/2}t_i^{-4} \exp(-3\tau_i^2 t_i^{-4}/4)$ satisfies all our assumptions. In this example, $s_i|x, t_i$ has an exponential power distribution (Box and Tiao, 1973, page 517).




Fig. 2. Illustration of the incentive to deviate from honest forecasting in the reputational cheap talk model. A forecaster with signal s believes that the state x is normally distributed around the posterior mean $E[x|s]$, strictly between s and the prior mean μ . A forecaster who observes signal s optimally deviates by reporting signal $E[x|s]$.

Proposition 1 *If the market conjectures honest forecasting $h(s)$, the forecaster shades the forecast towards the prior mean by reporting*

$$d(s) = h(h(s)) = \left(\frac{\tau}{\tau + \nu}\right)^2 s + \left[1 - \left(\frac{\tau}{\tau + \nu}\right)^2\right] \mu. \quad (3)$$

Another interpretation of the conservative deviation is based on the following logic. A forecaster who receives a signal s above the prior mean μ concludes that the average forecast error $E(\varepsilon|s) = \nu(s - \mu) / (\tau + \nu)$ is positive. The forecaster then optimally deviates by removing this expected error from the true signal and thereby pretending to have signal $\hat{s} = s - E(\varepsilon|s) = E(x|s)$.

The incentive to deviate relies on the simultaneous presence of private and public information. Indeed, if the signal is perfectly informative (i.e., in the limit as $\tau \rightarrow \infty$), the posterior expectation puts zero weight on the prior belief, $E(x|s) = s$, and $x = s$ with probability one. Similarly, in the absence of prior information (i.e., in the limit as $\nu \rightarrow 0$), the posterior expectation is again equal to the signal. In these two extreme cases, there is no incentive to deviate, $d(s) = s$, so truth-telling is an equilibrium.

According to Proposition 1, sophisticated forecasters who are taken at face value report conservative forecasts in order to fool the market into believing that they have more

accurate signals. This characterization of the deviation incentive is relevant for several reasons. Understanding the pressure on forecasters to deviate from honesty shows us why truth-telling is impossible and sheds light on out-of-equilibrium forces. If the forecaster has mixed incentives, caring about both her reputation and forecast accuracy, the incentive to deviate from honesty can induce a conservative bias in equilibrium. Finally, deviation incentives persist when the market is not fully rational.

4.3. Equilibrium

According to Proposition 1, honest forecasting is incompatible with equilibrium. As in any cheap talk game, ruling out truth-telling implies that there is no fully separating equilibrium. By definition, in a fully separating equilibrium, the strategy that maps signals into forecasts can be inverted. As before, the market then infers the signal, in which case the forecaster with signal s wishes to deviate to the forecast that corresponds to signal $s' = E(x|s)$, which is different from s whenever $s \neq \mu$.

There exists an equilibrium with complete pooling, as is common in cheap talk games. In such a “babbling” equilibrium, the forecaster issues the same message m regardless of the signal received. Any message the market receives is then interpreted as carrying no information about the signal. More generally, in equilibrium only part of the forecaster’s information is conveyed to the market. Equilibrium forecasting necessarily involves some degree of pooling of signals into messages.

Due to the cheap talk nature of the game, the actual language used to send equilibrium messages is indeterminate. However, the market can easily translate message m into the best estimate of the state that incorporates the information contained in that message, namely, $E(x|m)$. In this natural language, the forecaster is effectively communicating $E(x|m)$ to the market. Since it is a conditional expectation of x , this forecast is uncorrelated with its error. In this sense, the reputational equilibrium forecast satisfies the orthogonality property. We conclude, therefore,

Proposition 2 (Coarseness in Reputational Equilibrium) *There is no reputational cheap talk equilibrium in which information is fully revealed. Any equilibrium can be defined with a language such that the forecast has the orthogonality property.*

We now show by example that there exists a partially informative equilibrium. The example involves a binary forecasting strategy in which the forecaster reports a “high” message m_H whenever the signal s weakly exceeds a threshold and a “low” message m_L otherwise. Our binary equilibrium strategy is symmetric, because the threshold signal equals the average signal μ . That is,

Proposition 3 (Binary Reputational Equilibrium) *There exists a symmetric binary reputational cheap talk equilibrium. This equilibrium is the unique equilibrium in binary strategies. If the forecasters use the natural language, we have $m_L = \mu - \sqrt{2\tau/\pi\nu(\nu + \tau)}$ and $m_H = \mu + \sqrt{2\tau/\pi\nu(\nu + \tau)}$.*

4.4. Discussion

The theory of reputational cheap talk relies on the fact that the market rationally uses all the information available ex post to evaluate the forecaster. Suppose that instead the market commits ex ante to evaluate the forecaster by comparing the forecast m with the realization x , according to the magnitude of the error, $|m - x|$. In this case, the forecaster's optimal strategy is to honestly report $m = E(x|s)$. This outcome is essentially the default case of honest forecasting as explained in Section 2. This outcome also obtains if the market (incorrectly) believes the forecaster's message to be equal to her signal, $m = s$.

We now relate our results to previous work in the agency literature. Brandenburger and Polak (1996) consider a privately informed agent who makes an investment decision with the aim of realizing the most favorable assessment by the stock market. In their model the market already assesses the profitability of the exogenously coarse (in fact, binary) decision before the true state is realized. Despite the differences, in both models agents have an incentive to deviate conservatively.

At a superficial level, the deviation result we obtain in Proposition 1 is reminiscent of Prendergast's (1993) "yes-men" effect, though it is driven by different forces and essentially goes in the opposite direction. While in Prendergast's model the agent does not move sufficiently away from her information about the principal's signal, in our model the agent does not move away from the prior mean. By identifying the principal's signal with the ex post (noisy) realization of the state, one can see that Prendergast's deviation report is biased toward the state, rather than the prior. In both models, in equilibrium the agent cannot transmit all her information.

Our result on the direction of the deviation incentives relies on the assumption that forecasters do not have superior information about their forecasting talent compared to the market. This assumption is questionable in a dynamic setting since dishonest forecasters would learn about their precision faster than the market. If the forecaster has private information about her forecasting talent, the model becomes one of two-dimensional signaling (see Trueman, 1994, for a first analysis). The addition of this second dimension of private information introduces a new signalling incentive. Intuitively, since the posterior

expectations of more talented forecasters are more variable around the prior mean, the attempt to signal talent generates an incentive to exaggerate.¹⁰

Prendergast and Stole (1996) isolate this exaggeration tendency in a managerial reputational signaling model in which the market that does not observe the state. In our model, in contrast, when evaluating forecasters the market has access to additional information about the state, such as its ex post realization or the contemporaneous forecasts of others. The addition of such ex post information introduces the new conservatism effect we isolate in Proposition 1 for the case in which forecasters do not know their talent.¹¹ Overall, concerns for absolute accuracy drive forecasters to be conservative if they do not know their ability, but to exaggerate if they do know it sufficiently well.¹²

Rather than performing direct tests of reputational cheap talk, most of the existing empirical literature provides indirect evidence of reputational concerns based on heterogeneity across forecasters. Lamont (2002) finds that older forecasters tend to deviate more from the consensus. Chevalier and Ellison (1999) find that older mutual fund managers follow bolder investment strategies. Hong, Kubik, and Solomon (2000) conclude that the lower accuracy of older stock analysts is due to the fact that they move earlier. In Section 6 we return to alternative methodologies for detecting strategic behavior in forecasts.

5. Forecasting contest

Our first theory posits that the market optimally evaluates the forecasting talent based on all the information that is available ex post, but is unable to commit ex ante on a particular evaluation rule. However, in many instances there are different mechanisms in place for rewarding successful forecasts. For example, forecasters often participate in contests with prizes allocated to the best performers. Even in the absence of monetary prizes, the publicity effect for the winner can be large. In the much publicized semi-annual Wall Street Journal forecasting contest, the most accurate forecaster over the previous six months is typically rewarded with a write-up.¹³

¹⁰As Lim (2002) argues, a forecaster who knows her own ability also has an incentive to underreact to new public information.

¹¹For an example of reputational signaling with bounded rationality, see Zitzewitz's (2001b) model, wherein the market evaluates forecast quality using a simple econometric technique. In his model, forecasters have information on their own ability and thus have an incentive to exaggerate.

¹²As also Avery and Chevalier (1999) suggest, young managers with little private information about their own ability should be conservative, whereas older managers would exaggerate. This contrasts with Prendergast and Stole's (1996) prediction of impetuous youngsters and jaded old-timers when the same manager makes repeated observable decisions with an unobserved but constant state.

¹³See Stekler (1987) for an early study of the relative accuracy of forecasts. A number of rankings are available on-line. For example, Validea (www.validea.com), BigTipper.com (www.bigtipper.com) and BulldogResearch.com (www.bulldogresearch.com) track stock recommendations made by Wall Street pro-

Note that a rational observer should take into account not just whether a forecaster has won a tournament, but also the absolute size of the forecast error. If the processing power of the observing public were unlimited, the business media could simply publish a list of forecasts and realized outcomes, instead of creating rankings and contests. However, the rank-order information is particularly salient for the problem of evaluating the quality of forecasts. This salience may stem from the behavioral notion (e.g., Kahneman, 1973) that the public pays limited attention to forecasters.¹⁴ That is, it is easier for people to keep in mind who is an “all star” analyst, or who came first in a contest, than specific details about forecast accuracy. As a result, forecasters are concerned about their relative accuracy.

In this section, we consider a symmetric simultaneous winner-take-all contest with a large number of forecasters. In this forecasting contest, the participating forecasters simultaneously submit individual forecasts based on their private information. A prize is awarded to the forecaster whose forecast is closest to the realized state. As opposed to what happens in our reputational model, in a forecasting contest the market commits ex ante to a particular reward scheme.

There is remarkably little previous work on the behavior forecasters in contests. Steele and Zidek (1980) are the first to study a sequential forecasting contest between two privately informed forecasters. They assume away game-theoretic considerations by supposing truthful reporting by the first forecaster. After observing the first forecast, the second guesser faces a simple decision problem and has a clear advantage. Indeed, Bernhardt and Kutsoati (2004) confirm empirically that financial analysts who release late earnings forecasts tend to overshoot the consensus forecast in the direction of their private information. Laster, Bennett, and Geoum (1999) study a winner-take-all simultaneous forecasting contest in which all forecasters share the same public information. Here, we allow forecasters to have private information on the state.¹⁵ In a forecasting contest, a forecaster’s payoff is the probability that the realized state is closer to her forecast than to any other forecast. This probability is equivalent to the market share or fraction of votes to be maximized in Hotelling’s (1929) pure location game.¹⁶

fessionals and then rank the analysts based on the performance of their selections. Forecasting contests are run also for noneconomic variables (see e.g., the *National Collegiate Weather Forecasting Contest*).

¹⁴For a broader discussion of limited attention and its implications for finance, see, e.g., Hirshleifer and Teoh (2003).

¹⁵There is no clear reason to reward accurate forecasters in the absence of heterogeneous private information.

¹⁶An extensive literature in economics and political science analyzes versions of this game without private information. As is well known (see, e.g., Osborne and Pitchik, 1986), equilibria in this classic game crucially depend on the number of players and often involve mixed strategies. A forecasting contest

5.1. Model

The forecasting contest proceeds as follows. First, forecasters observe their private signals s_i of common precision τ and simultaneously submit their forecasts c_i . Once the true state x is publicly observed, the forecaster whose forecast c_i turns out to be closest to x wins a prize proportional to the total number of forecasters participating in the contest. We consider the limit of the game as the number of forecasters tends to infinity.

The expected payoff of forecaster i who observes signal s_i and reports forecast c_i is

$$U(c_i|s_i) = \frac{q(c_i|s_i)}{\gamma(c_i|c_i)}, \quad (4)$$

where $\gamma(c|x)$ is the density of the forecasts released by all forecasters conditional on state x , with full support on the real line. Issuing forecast c_i , forecaster i wins only if $x = c_i$, which occurs with chance $q(c_i|s_i)$. Conditional on being on the mark, the prize is divided among all the winning forecasters, and their density computed at $x = c_i$ is equal to $\gamma(c_i|c_i)$.

5.2. Deviation

We now show that honest forecasting is not an equilibrium. Let us consider a single forecaster with signal s competing against forecasters who are all reporting their honest forecasts. Without loss of generality, we focus on $s > \mu$ as depicted in Fig. 3. What is the best reply for such a forecaster?

According to (4), the best forecast maximizes the ratio between the probability of winning the first prize and the number of forecasters with whom this prize is shared. First, the probability of winning conditional on signal s is equal to the posterior belief on the state $x|s$, that is the normal distribution centered at $E(x|s)$ and depicted on the right in Fig. 3. Second, the curve on the left in Fig. 3 depicts $\gamma(x|x)$, the denominator of the ratio maximized by the forecaster. This represents how the mass of forecasters with *correct* forecasts changes as a function of the state x . The shape of $\gamma(x|x)$ depends on the weight the other forecasters assign to their private signals. Since the other forecasters put a positive weight on the prior mean, $\gamma(x|x)$ is bell-shaped around μ .¹⁷

is a version of Hotelling's simultaneous location game in which the forecasters (firms or politicians) have private information on the distribution of the state (location of consumers or voters).

¹⁷When the other forecasters put zero weight on the prior (e.g., because they are perfectly informed), $\gamma(x|x)$ is constantly equal to one. More generally, $\gamma(x|x)$ is *not* a probability density function of x .

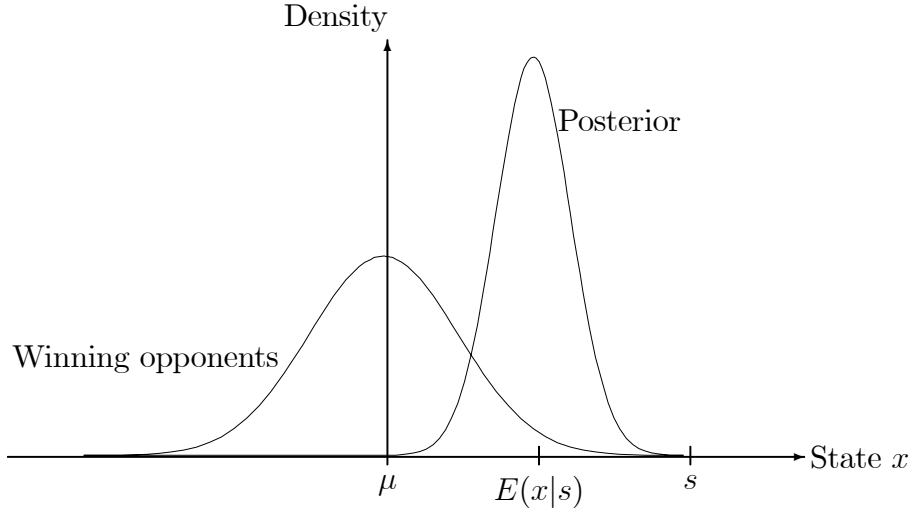


Fig. 3. Illustration of the incentive to deviate from honest forecasting in the forecasting contest model. The forecaster’s belief density is maximized at the posterior mean $E[x|s]$. The further the winning forecast is from the prior mean μ , the fewer the winning opponents. The forecaster deviates from $E[x|s]$ to a forecast further from μ .

Fig. 3 illustrates that the probability of winning is flat at the honest $E(x|s)$, while the frequency of correct forecasts is decreasing in the distance $|x - \mu|$. At the honest $E(x|s)$, it is then optimal for the forecaster to move away from the prior mean μ toward the private signal s , because the second-order loss resulting from a lower probability of winning is more than compensated by the first-order gain due to reduced competition.

Proposition 4 (Exaggeration in Contest Deviation) *If all other forecasters use the honest strategy $h(s)$, the best response in the contest for forecaster i is to exaggerate.*

The optimal deviation forecast is a weighted average of s_i and μ , but the weight on μ is lower than in the honest forecast. The contest deviation forecast is then positively correlated with its error: when x is above μ , the forecast is too high on average.

5.3. Equilibrium

Having established that honest forecasting is incompatible with equilibrium, we now turn to characterizing the symmetric Bayes-Nash equilibrium. At such an equilibrium, each forecaster for every signal s_i best replies to her conjecture about the opponents’ distribution $\gamma(c|x)$, and this conjecture is correctly derived from the opponents’ strategies.

Proposition 5 (Exaggeration in Contest Equilibrium) *For any values of $\nu > 0$ and $\tau > 0$, the contest has a unique symmetric linear equilibrium with strategy $c(s) = Cs + (1 - C)\mu$ with $C \in (0, 1)$. Forecasters put more weight on their private information than according to the honest conditional expectation: $C = (\sqrt{\tau^2 + 4\nu\tau} - \tau) / 2\nu > \tau / (\nu + \tau)$.*

Forecasters differentiate themselves from their competitors by putting excessive weight on their signals. As in the honest forecast, the weight on the signal is increasing in τ and decreasing in ν . This weight is larger than in the honest forecast, so the contest gives an incentive to move away from μ .¹⁸ The symmetric equilibrium strikes a balance: opponents disperse themselves to such an extent that forecaster i is happy to reply precisely with the same dispersion. The equilibrium forecast is positively correlated with its error.

5.4. Discussion

In the absence of private information ($\tau = 0$), the only symmetric equilibrium is in mixed strategies as in Laster, Bennett, and Geoum (1999) and Osborne and Pitchik (1986), who find that with infinitely many symmetrically informed players the distribution of equilibrium locations replicates the common prior distribution about x . The addition of private information has the desirable effect of inducing a symmetric location equilibrium in pure rather than mixed strategies.

The contest equilibrium satisfies $C < 1$, so the forecast is not as extreme as the maximum likelihood estimate (MLE). However, the MLE results in the contest when the prior on the state x is improper, i.e., uniform on the real line. If the opponents forecast $c = s$, their forecasts are normally distributed around x , and the term $\gamma(c|c)$ is constant in c . Forecaster i 's best reply will then be $c_i = s_i$, since the constant term $\gamma(c|c)$ does not distort the forecaster's problem. Truth-telling by all forecasters is then an equilibrium in the absence of public information. Thus, the contest distortion depends on the presence of prior information that anchors the forecasts of the opponents around μ . The tendency of opponents to cluster around the prior mean drives forecasters away from it.

6. Empirical testing

We now discuss how to empirically test our theories. In Section 6.1 we revisit the orthogonality test presented in Eq. (2) and discuss the implications of recent empirical

¹⁸Even though a best reply to honesty does not exist for all parameter values, the equilibrium in linear strategies exists for all parameter values. Intuitively, with increased weight on their signal, the opponents are less concentrated around μ , mitigating the incentive to move away from μ .

work in light of our theories. In Section 6.2 we consider the effect of the strategic incentives on distributional properties of the forecasts.¹⁹

6.1. Orthogonality

Except for the reputational equilibrium forecasts, we find linear forecasting rules of the form $f_i(s_i) = F_i s_i + (1 - F_i) \mu$ for some constant weight F_i between zero and one. The conditional distribution of the linear forecast is then

$$f_i|x \sim N(F_i x + (1 - F_i) \mu, F_i^2 / \tau_i). \quad (5)$$

We have already noted that under honesty the forecast is uncorrelated with its error $f_i - x$ when $F_i = \tau_i / (\nu + \tau_i)$. If forecasters report honestly, their error cannot be predicted from the forecast. In contrast, the contest and reputational deviation forecasts fail to inherit this property; thus, once the forecast has been released, the sign of its error can be predicted.

Proposition 6 *For the linear forecasting rules of the form (5), the correlation of the forecast and its error has the same sign as $F_i - \tau_i / (\tau_i + \nu)$. The correlation is positive in the contest and negative in the reputational deviation. The reputational equilibrium forecast satisfies orthogonality, but is not efficient.*

A typical empirical test for the hypothesis that the forecasts are conditional expectations ($E(x|I_i)$ for some information set I_i) is based on regressing the realized forecast error on the forecasts. Most studies report a positive correlation of the forecast and its error, consistent with the prediction of our contest theory. For example, Batchelor and Dua (1992) find that forecasters put too little weight on the forecasts previously released by other forecasters (or, equivalently in our model, on the prior mean). However, Keane and Runkle (1990, 1998) question the statistical significance of any bias, as they note that the tests are not as powerful as is usually assumed once the correlation among forecast errors across forecasters is properly taken into account (see also Section 7.1 below).

When deriving the correlation of the forecast and its error, we treat the prior mean μ as a parameter. Empirical work, however, must control for the prior mean. Recently,

¹⁹To directly test for the implications of strategic behavior, we could compare nonanonymous with anonymous forecasting data. In anonymous surveys the authorship of individual forecasts is not disclosed. The mere existence of anonymous surveys (starting in 1946 with the Livingston Survey) pre-supposes a belief that anonymous forecasts are more honest. As Croushore (1993) reports, “This anonymity is designed to encourage people to provide their best forecasts, without fearing the consequences of making forecasts errors. In this way, an economist can feel comfortable in forecasting what she really believes will happen.” Yet, Stark’s (1997) empirical analysis of the anonymous Survey of Professional Forecasters confirms that forecasters in anonymous surveys face similar incentives as in nonanonymous surveys.

Zitzewitz (2001a) and Bernhardt and Kutsoati (2004) perfect the orthogonality methodology to fully account for the presence of prior information and the correlation of forecast errors. They rewrite $f_i = F_i s_i + (1 - F_i) \mu$ as $f_i - \mu = F_i (s_i - \mu) = (F_i/H_i) (h_i - \mu)$, where h_i is the honest forecast and $H_i = \tau_i / (\tau_i + \nu)$ is the honest weight on the signal, and condition on all publicly available information at the moment of forecasting.

These findings point clearly toward the presence of exaggeration rather than herding in the earning forecasts released by I/B/E/S analysts. Exaggeration is in line with the equilibrium of our forecasting contest, but is inconsistent with the deviation or the equilibrium of our reputational cheap talk model. As Zitzewitz (2001b) argues, the observed exaggeration is also consistent with an alternative version of the reputational signaling model in which forecasters are privately informed about their own ability, as in Trueman's (1994) model, and are evaluated according to an econometric technique.

6.2. Forecast variability

We now turn to the statistical properties of an individual's forecast f . We derive the conditional variance of the forecast under our different theories.

The linear forecasts' weights F on the signal implied by the theories all depend on τ only through the relative signal precision $\rho \equiv \tau/\nu$, the precision of private information relative to the prior precision. Apart from a common scaling factor equal to the variance of the prior distribution, $1/\nu$, all variances below can be written as a function of ρ .

Proposition 7 *The conditional variances are $V(h|x) = (1/\nu)\rho/(1 + \rho)^2$ for the honest forecast, $V(d|x) = (1/\nu)\rho^3/(1 + \rho)^4$ for the reputational deviation forecast, and $V(c|x) = (1/\nu) \left(2 + \rho - \sqrt{\rho^2 + 4\rho}\right) / 2$ for the contest equilibrium forecast. The reputational binary equilibrium forecast satisfies $V(r|x) \leq (1/\nu) 2\rho/\pi (1 + \rho) = V(r|x = \mu)$.*




Fig. 4. Conditional forecast variances as functions of the relative precision ρ , fixing $\nu = 1$. The solid line shows honesty's $V(h|x)$, the dots the reputational deviation's $V(d|x)$, the line with full dots the reputational equilibrium's upper bound $V(r|x = \mu)$, and line with empty dots the contest equilibrium's $V(c|x)$.

As Fig. 4 illustrates, the honest forecast has vanishing variance both when the forecasters are very poorly informed (and thus forecast near μ) and when they are very well informed (and thus forecast the realized x well). Similar properties hold for the reputational forecasts.²⁰ In general, herding or exaggeration can be inferred from forecast dispersion only after controlling for the quality of the forecaster's information. Zitzewitz (2001a) also emphasizes this point.

In the limit as the private signals become uninformative ($\rho \rightarrow 0$), the contest makes a markedly different prediction. In this limit, the distribution of equilibrium forecast locations converges to the common prior distribution about the state, with conditional variance equal to $1/\nu > 0$. This result is consistent with Osborne and Pitchik's (1986) and Laster, Bennett, and Geoum's (1999) findings, obtained for the limit case of Hotelling's location game with a large number of players. In cases with imprecise private signals, empirical tests among the theories could build on this finding.

²⁰When the private signals become very informative, $\rho \rightarrow \infty$, then for any fixed $x > \mu$, the reputational forecast becomes concentrated near $\mu + \sqrt{2/\nu}$ and has vanishing variance. This fact is not evident from the upper bound $V(r|x = \mu)$ displayed in Fig. 4. This upper bound corresponds to the variance conditional on the particular signal realization $s = \mu$.

7. Theoretical extensions

We now extend the baseline model in two directions relevant for empirical work. In Section 7.1 we allow for common errors in the signals and in Section 7.2 we consider simultaneous forecasting of multiple variables.

7.1. Common error

As Keane and Runkle (1998) stress, the significant positive correlation among the residuals in the orthogonality regression indicates the presence of a common error in the forecasts. Since the forecasts are released well in advance, there are often unpredictable changes to the variables after the forecasts are submitted. In addition, realizations are often observed with noise. To check the robustness of our results, in this section we modify the model by allowing for the presence of common errors.

Suppose that the market does not observe x but only the imperfect signal $y = x + \varepsilon_0$ before evaluating the forecasters. Thus, the forecaster's rewards are defined in terms of the new target variable y rather than x . Assume that $x, \varepsilon_0, \varepsilon_1, \dots, \varepsilon_n$ are independent normal variables, and let τ_0 denote the precision of ε_0 . The signal is $s_i = x + \varepsilon_i = y + \varepsilon_i - \varepsilon_0$. Conditional on y , the relevant signal errors $\varepsilon_1 - \varepsilon_0, \dots, \varepsilon_n - \varepsilon_0$ are now correlated as suggested by Keane and Runkle. Yet, the honest forecast of y is the same as the honest forecast of x , i.e., $E(y|s_i) = E(x + \varepsilon_0|s_i) = E(x|s_i)$.

Proposition 8 *Propositions 1–4 continue to hold with common error, with suitable modification of the closed-form expressions. If the noise in the common error is sufficiently small (i.e., τ_0/ν is large compared to τ/ν), there is a linear equilibrium with exaggeration in the forecasting contest.*

7.2. Multiple dimensions

In some applications, forecasters are evaluated on the basis of several contemporaneous forecasts or their entire record of past forecasts. To simplify the presentation, in our baseline model above we focus on the one-dimensional case in which forecasters are evaluated on the basis of a single forecast of one variable. We now extend our normal learning model to multivariate settings in which the evaluation is made on the basis of a number of different variables.

We extend the model by letting $\tilde{g}(s|x, t) = t\hat{g}(t||x - s||)/2$, where the state x and the signal s are multivariate and the talent t is univariate. Under the assumption that signals

closer to the state are better news about the talent, Proposition 1 and 3 continue to hold. In addition, $\gamma(x|x)$ is a multivariate bell-shaped function centered on μ and the posterior $q(x|s)$ is a multivariate bell-shaped density around $E(x|s)$. Proposition 4 then continues to hold. In conclusion, the strategic distortions we identify above hold more generally for multidimensional states, signals, and forecasts.

8. Impact of forecasts

In this section, we discuss the impact of forecasts on market expectations and prices in a financial setting. Our theories are general with respect to the interpretation of the forecasted state variable x . For concreteness, we apply the model to forecasting stock prices, but the results are also valid for forecasts of exchange rates (e.g., Frankel and Froot, 1987) and other macroeconomic variables (e.g., Romer and Romer, 2000).

We first extend our baseline model to represent the pricing of an asset with fundamental value V , which is partly explained by the forecasted state $x = V + \varepsilon_x$. The value has prior mean μ and precision ν_V . The timing is as follows. First, as in our baseline model, the forecasters simultaneously report the forecasts f_1, f_2, \dots, f_n based on their private signals s_1, s_2, \dots, s_n , with $s_i = x + \varepsilon_i$. Second, the market observes the realization of the public signal $y = V + \varepsilon_y$ and determines the asset's price as the conditional expectation, $P = E(V|y, f_1, \dots, f_n)$. Third, the state is realized, $x = V + \varepsilon_x$. We assume that $V, \varepsilon_y, \varepsilon_x, \varepsilon_1, \dots, \varepsilon_n$ are independent normal variables, that all of the error terms have mean zero, and that ε_i has precision τ_i . The precision of x then satisfies $1/\nu = 1/\nu_V + 1/\tau_x$.

We address two issues. First, we study how market prices react to the release of forecasts (Section 8.1). Second, we analyze strategic forecasting in a modified version of the model in which the forecasters target the market price P (Section 8.2).

8.1. Price reaction

A perfectly rational market should adjust for the biases induced by strategic forecasting. In all of the linear forecasts, there is a one-to-one mapping between the forecasts and the private signals, so the market can fully recover the forecasters' private signals. Since all variables are normally distributed, P is a weighted average of $E(V|y)$ and $E(x|f_1, \dots, f_n)$. Likewise, $E(x|f_1, \dots, f_n)$ is a weighted average of the honest forecasts $E(x|f_1), \dots, E(x|f_n)$.

If rational, the market correctly adjusts for the strategic distortions present in the forecasts. If the forecast function is invertible, a rational market is able to recover the signal and form the correct posterior expectation. Hence, the predictions for asset price

reactions are identical regardless of the particular forecast function used, provided that the forecast is invertible. Only the reputational equilibrium forecast is not invertible; its statistical properties are given in Proposition 3.

Consider an outside observer (such as an empiricist) who analyzes the price reaction to a forecast by incorrectly assuming that this forecast is honest rather than strategic. Such an observer would wrongly conclude that the market overreacts to conservative reputational deviation forecasts, and underreact to exaggerated contest forecasts.²¹

8.2. Forecasting prices

We now turn to a variant of the model that is relevant to the study of financial analysts. Suppose that the target variable of the forecasts is not x but rather the asset price P , and that the forecasters are evaluated on the basis of their ability to predict P . Thus, P plays the role of x in the payoff functions of the forecasters. For simplicity, we again make the symmetry assumption that all forecasters have equally precise signals, i.e., $\tau_i = \tau$ for all $i = 1, \dots, n$.

In this setting, honest forecasting of P is more complicated because the forecasts directly influence the distribution of P . As long as the forecasters' strategies are fully invertible, the symmetry assumption implies that $P = \alpha_y y + \alpha_s \bar{s} + \alpha_\mu \mu$, where \bar{s} is the average of the forecasters' signals, which the market extracts by inverting the forecast strategy. The parameters $\alpha_y, \alpha_s, \alpha_\mu \in (0, 1)$ sum to one, and are determined by the precision of the prior belief about V relative to the precision of the public signal y and of the average private signal \bar{s} .

Proposition 9 *There exists a rational expectations equilibrium for honest forecasting of P in linear strategies, $h(s) = Hs + (1 - H)\mu$, in which*

$$H = (\alpha_y + \alpha_s) \frac{\tau}{\tau + \nu} + \frac{\alpha_s}{n} \left(\frac{\nu}{\tau + \nu} \right) \in (0, 1). \quad (6)$$

Relative to this benchmark, the reputational and limit contest theories have similar properties to those reported in Proposition 8.

The equilibrium in honest strategies is identical to the one in Eq. (1) when $\alpha_s = 0$ and $\alpha_y = 1$. Otherwise, two effects pull the forecast away from the honest forecast of x . First, P depends positively on the forecasts of the opponents. This yields a beauty contest effect that encourages forecasters to attach greater weight to their private signal about x .

²¹We refer to Gleason and Lee (2003) for a recent empirical analysis of price reactions to forecast revisions.

Second, the estimate P is systematically closer to the prior mean μ than the signal x about which the forecasters have information, thus there is an incentive to attach less weight to the private signal. The same effects influence our two strategic theories, but relative to the honest benchmark the properties are the same as before.

9. Conclusion

In this paper we formulate and contrast two distinct theories of strategic forecasting within the normal model. In the process we advance our understanding of the forces that might drive informed agents to deviate from honest reporting of their conditional expectations. Misreporting results from the subtle interaction of the private information available to each individual with the public prior information available to the market and commonly shared by all agents.

Our first theory posits that the market has all the information contained in both the forecasts and the realization of the state and uses it to ex post optimally evaluate the forecasters. We assume that on average better informed forecasters observe signals closer to the state and that forecasters who are reputed to be better informed have a higher payoff. We show that forecasters wish to appear to have received a signal equal to the posterior expectation of the state conditional on the signal actually received. In the presence of public information, the observed signal is necessarily different from the posterior expectation. If the market naively believes that forecasters are honest, forecasters then shade their forecasts toward the prior mean. If the market is fully aware of the forecasters' strategic incentives, equilibrium forecasts are imprecise but not shaded.²²

Our second theory posits that competition for best accuracy takes place with pre-specified rules. The evaluation in a forecasting contest is ex post optimal when the market can only observe the accuracy ranking, possibly due to the market's limited attention. Since the forecasters share the same public information, competition is highest when the state turns out to be equal to the prior mean. At the posterior expectation, a small deviation away from the prior mean results in a first-order gain due to reduced competition and a second-order loss due to lower probability of winning. Equilibrium forecasts in a winner-take-all contest are then excessively differentiated relative to the corresponding conditional expectations.

In both the reputational and the contest theories the incentive to deviate from honesty is driven by the property that private signals are unimodally centered around the state

²²In reality, competition among forecasters combines elements of both theories. For example, *Institutional Investors* ranks analysts based on the opinions of large institutional investors. See Ottaviani and Sørensen (2005) for results on relative reputational concerns.

and the fact that public information is available at the moment of forecasting. In the absence of public information, honest forecasting is an equilibrium in both models. In reality, public information is pervasive, and its presence as well as the correlation across forecasts pose major challenges for empirical work. Having dealt with these issues, recent empirical studies find a significant and strong exaggeration in the forecasts. These findings are in line with the equilibrium of our forecasting contest.

Appendix

Proof of Proposition 1. By observing $m = h(s)$ and x , the market infers the realized signal $\hat{s} = h^{-1}(m)$ and error $\hat{\varepsilon} = \hat{s} - x$. The updated reputation is then $p(t|m, x) = \tilde{g}(\hat{s}|x, t) p(t) / g(\hat{s}|x)$. This posterior reputation inherits intuitive properties from the assumptions imposed on \tilde{g} . First, the posterior reputation depends on m and x only through the absolute size of the error $|\hat{\varepsilon}|$. Second, a small realized absolute error is good news about the forecaster's talent: for any $t < t'$, the likelihood ratio

$$\frac{p(t|m, x)}{p(t'|m, x)} = \frac{\tilde{g}(\hat{s}|x, t) p(t)}{\tilde{g}(\hat{s}|x, t') p(t')} \quad (7)$$

is increasing in $|\hat{\varepsilon}|$. This second property and the fact that v is strictly increasing imply (see Milgrom, 1981) that $W(m|x)$ is a strictly decreasing function of the inferred absolute error $|\hat{\varepsilon}|$.

We now consider the best response of a forecaster with signal s . The posterior distribution on x is normal with mean $h(s)$ and variance $1/(\nu + \tau)$. The inferred forecast error $\hat{\varepsilon} = h^{-1}(m) - x$ is then normally distributed with mean $h^{-1}(m) - h(s)$ and variance $1/(\nu + \tau)$. The best reply maximizes the expected value of W , or equivalently, minimizes a symmetric loss function of the error $h^{-1}(m) - x$. The forecaster then chooses m such that the error has mean zero by setting $h^{-1}(m) = h(s)$. \square

Lemma 1 *If a density $\hat{g}(\cdot)$ satisfies the property*

$$\hat{g}(t'\varepsilon) \hat{g}(t\varepsilon') < \hat{g}(t\varepsilon) \hat{g}(t'\varepsilon') \text{ for } \varepsilon' > \varepsilon \geq 0 \text{ and } t' > t, \quad (8)$$

its countercumulative distribution satisfies it as well:

$$\left[1 - \hat{G}(t'\varepsilon)\right] \left[1 - \hat{G}(t\varepsilon')\right] < \left[1 - \hat{G}(t\varepsilon)\right] \left[1 - \hat{G}(t'\varepsilon')\right] \text{ for } \varepsilon' > \varepsilon \geq 0 \text{ and } t' > t. \quad (9)$$

Proof. Integrating (8) for $\varepsilon'' > \varepsilon'$, we obtain

$$t'\hat{g}(t'\varepsilon) \left[1 - \hat{G}(t\varepsilon')\right] < t\hat{g}(t\varepsilon) \left[1 - \hat{G}(t'\varepsilon')\right] \quad (10)$$

for $\varepsilon' > \varepsilon$. The left-hand side (LHS) and the right-hand side (RHS) of (9) are equal for $\varepsilon' = \varepsilon$. By (10) we know that the derivative of the LHS is larger than the derivative of the RHS of (9). We conclude that (9) holds. \square

Proof of Proposition 3. To support this equilibrium, we also need to specify the market's beliefs following out-of-equilibrium messages. When receiving a message different from m_L and m_H , the market assumes that the forecaster possesses a signal below the

threshold, which results in the same posterior reputation as message m_L . These beliefs satisfy the requirements of a perfect Bayesian equilibrium.

We assume that the market conjectures a binary strategy with threshold signal \check{s} . We find $\varphi(m_H|x, t) = \int_{\check{s}}^{\infty} \tilde{g}(s|x, t) ds = \int_{\check{s}}^{\infty} t\hat{g}(t|s-x)/2ds$. This equation reduces to $\left[1 - \hat{G}(t|\check{s}-x)\right]/2$ when $\check{s} > x$ and to $\left[1 + \hat{G}(t|\check{s}-x)\right]/2$ when $\check{s} < x$, where \hat{G} is the distribution function that corresponds to the density \hat{g} . Therefore, we have $\varphi(m_H|x, t) = 1 - \varphi(m_H|2\check{s}-x, t) = \varphi(m_L|2\check{s}-x, t)$, which implies the symmetry property $W(m_H|x) = W(m_L|2\check{s}-x)$.

When $\check{s} > x$, it follows from Lemma 1 that message m_H is worse news about the talent than the observation that $s \geq x$. Thus, we have $W(m_H|x) < W(m_H|\check{s})$ for all $x < \check{s}$. Symmetrically, we have $W(m_H|x) > W(m_H|\check{s})$ for all $x > \check{s}$. These inequalities and symmetry imply that $W(m_H|x) > W(m_L|x)$ when $x > \check{s}$.

We now show that when $\check{s} = \mu$, the forecaster does not wish to deviate from the putative equilibrium strategy. By symmetry, it suffices to assume that $s \geq \mu$ and to verify that $U(m_H|s) \geq U(m_L|s)$. Using the symmetry of W , $U(m_H|s) - U(m_L|s)$ equals

$$\int_{-\infty}^{\infty} [W(m_H|x) - W(m_L|x)] q(x|s) dx = \int_{\mu}^{\infty} [W(m_H|x) - W(m_L|x)] [q(x|s) - q(2\mu - x|s)] dx. \quad (11)$$

Since $q(x|s)$ is the density of the symmetric normal distribution with a mean weakly above μ , we have $q(x|s) \geq q(2\mu - x|s)$ when $x \geq \mu$. We have already shown that $W(m_H|x) > W(m_L|x)$ when $x > \mu$, so the integrand of the last integral in (11) is everywhere non-negative, implying that (11) is non-negative, so that $U(m_H|s) \geq U(m_L|s)$ as desired.

To establish uniqueness of this equilibrium when binary strategies are used, we now show that when $\check{s} \neq \mu$, the forecaster wishes to deviate from the binary strategy. Without loss of generality, we focus on the case $\check{s} > \mu$. We show that there exists a signal $s > \check{s}$ such that $U(m_H|s) < U(m_L|s)$. As above, we have $U(m_H|s) - U(m_L|s) = \int_{\check{s}}^{\infty} [W(m_H|x) - W(m_L|x)] [q(x|s) - q(2\check{s} - x|s)] dx$. Also, we again have $W(m_H|x) > W(m_H|\check{s})$ for $x > \check{s}$ and $q(x|s) = q(2\check{s} - x|s)$ at $x = \check{s}$. By properties of the normal distribution, we obtain $q(x|s) < q(2\check{s} - x|s)$ for $x > \check{s}$, provided $E(x|s) = (\tau s + \nu\mu) / (\tau + \nu) < \check{s}$, which is certainly true for s slightly greater than \check{s} .

We derive the analytical expressions for the equilibrium messages resulting with the natural language. By applying the result that $E(y|y > 0) = \sigma\sqrt{2/\pi}$ for a normal variable $y \sim N(0, \sigma^2)$ (Johnson and Kotz, 1970), we find that $m_H = E(x|s \geq \mu)$ is equal to $E[E(x|s) | s \geq \mu] = E\left(\frac{\tau s + \nu\mu}{\tau + \nu} | s \geq \mu\right) = \mu + \frac{\tau}{\tau + \nu} E(s - \mu | s > \mu) = \mu + \sqrt{2\tau / (\pi\nu(\tau + \nu))}$. By symmetry, we have $m_L = \mu - \sqrt{2\tau / (\pi\nu(\tau + \nu))}$. \square

Proof of Proposition 4. As we derive in the text, $x|s_i$ is normally distributed

with mean $(\tau s_i + \nu \mu) / (\tau + \nu)$ and precision $\tau + \nu$. Let us suppose that all opponents use the linear strategy $c(s) = As + (1 - A)\mu$, with $A \in (0, 1]$. Then $c|x$ is normal with $E(c|x) = Ax + (1 - A)\mu$ and $V(c|x) = A^2/\tau$. Disregarding an irrelevant constant term and using the density of the normal distribution, we find

$$\log \gamma(c|c) = -\frac{\tau \{c - [Ac + (1 - A)\mu]\}^2}{2A^2} = -\frac{\tau(1 - A)^2(c - \mu)^2}{2A^2}. \quad (12)$$

This is a quadratic and concave function of c with its peak at μ . The forecaster maximizes $\log q_i(c_i|s_i) - \log \gamma(c_i|c_i)$, the difference of two quadratic and concave functions. The objective function is concave when the first concave term prevails, i.e., for $\tau + \nu > \tau(1 - A)^2/A^2$.

When $\tau + \nu > \tau(1 - A)^2/A^2$, the forecaster has a unique best reply $c_i = Bs_i + (1 - B)\mu$, with $B = \tau / [\tau + \nu - \tau(1 - A)^2/A^2] \in [\tau / (\tau + \nu), +\infty)$. When instead $\tau + \nu < \tau(1 - A)^2/A^2$, there is no best response because the incentive to move away from μ is so strong that forecaster i wishes to go to the extremes $\pm\infty$. In the knife-edge case $\tau + \nu = \tau(1 - A)^2/A^2$, the objective function is linear: whenever $s_i \neq \mu$, there is again no best reply, as the forecaster wishes to go to one of the extremes.

In particular, in the honest case, $A = \tau / (\tau + \nu)$, it is optimal to reply with $B = \tau^2 / (\tau^2 + \tau\nu - \nu^2)$, provided that $\nu/\tau < (1 + \sqrt{5})/2$. We conclude that the best reply is to exaggerate against truthtelling opponents. \square

Proof of Proposition 5. $C = 0$ is not compatible with a symmetric equilibrium, since in this case the opponents' forecasts are all equal to $c = \mu$, thus all replies other than μ yield forecaster i a higher payoff. We assume then that the forecasters use linear strategies of the form $c(s) = Cs + (1 - C)\mu$, with $C \in (0, 1]$. As we show in the proof of Proposition 4, forecaster i 's best reply is linear with weight $\tau / [\tau + \nu - \tau(1 - C)^2/C^2]$ on the signal, provided that $\tau + \nu > \tau(1 - C)^2/C^2$.

The fixed-point condition for a symmetric Nash equilibrium is that this linear strategy be equal to the one posited, i.e., $(1 - C)\tau = C^2\nu$. Inserting the values $C = 0, \tau / (\tau + \nu), 1$ in this quadratic equation, we conclude that there is only one positive solution, and that this solution belongs to the interval $(\tau / (\tau + \nu), 1)$. The second-order condition for the forecaster's optimization requires $\tau + \nu > \tau(1 - C)^2/C^2$. Using $(1 - C)\tau = C^2\nu$, this condition reduces to $\tau > -\nu C$, which is satisfied by the positive solution for C . The solution of the quadratic equation is $C = (\sqrt{\tau^2 + 4\nu\tau} - \tau) / 2\nu$. \square

Proof of Proposition 6. For the linear forecasting rules of the form (5) the correlation of the forecast and its error is

$$E[(f_i - x)f_i] = E\{[F_i\varepsilon_i + (1 - F_i)(\mu - x)][F_i(x + \varepsilon_i) + (1 - F_i)\mu]\} = F_i \left(\frac{F_i}{\tau_i} - \frac{1 - F_i}{\nu} \right), \quad (13)$$

and so has the same sign as $F_i - \tau_i / (\tau_i + \nu)$. \square

Proof of Proposition 7. The linear forecasts $f_i = F_i(x + \varepsilon_i) + (1 - F_i)\mu$ have conditional variance F_i^2/τ_i . Substituting the respective expressions for F_i yields the results.

The reputational equilibrium forecast r_i is binomially distributed. Given x , the chance of $r_i = m_H$ is $1 - \Phi(\sqrt{\tau_i}(\mu - x))$, where Φ is the distribution function of the standard normal distribution. Then r_i has mean $E(r_i|x) = \mu + [1 - 2\Phi(\sqrt{\tau_i}(\mu - x))] \sqrt{2\tau_i/[\pi\nu(\tau_i + \nu)]}$ and variance $V(r_i|x) = 4[1 - \Phi(\sqrt{\tau_i}(\mu - x))] \Phi(\sqrt{\tau_i}(\mu - x)) 2\tau_i/[\pi\nu(\tau_i + \nu)]$. Since $\Phi(\sqrt{\tau_i}(\mu - x)) \in [0, 1]$ for any x , we have $4[1 - \Phi(\sqrt{\tau_i}(\mu - x))] \Phi(\sqrt{\tau_i}(\mu - x)) \leq 1$, where the bound is tight, being achieved for $x = \mu$. We conclude that $V(r_i|x) \leq 2\tau_i/[\pi\nu(\tau_i + \nu)] = (1/\nu) 2\rho_i/[\pi(1 + \rho_i)]$. \square

Proof of Proposition 8. We consider first reputational cheap talk. For the deviation analysis, suppose that all opponents $j \neq i$ use a fully separating strategy. Besides the forecast of forecaster i , the evaluator observes n independent signals about x , namely, y and every s_j with $j \neq i$. From the well-known updating of beliefs on a normal state, this is equivalent to the observation of just one more precise signal about x . So, without loss of generality, we can imagine that y itself contains all the evaluator's external information on x . Since $y = x + \varepsilon_0$, where ε_0 is independent of ε_i and t_i , x is a sufficient statistic for y when predicting t_i , so the law of iterated expectations gives $p_i(t_i|m_i, y) = E[p_i(t_i|m_i, x) | y]$. Then $U_i(m_i|s_i) = \int_{-\infty}^{\infty} [\int_0^{\infty} u(t)p_i(t_i|m_i, y) dt] q_i(y|s_i) dy$ can be written

$$\int_{-\infty}^{\infty} \left[\int_0^{\infty} u(t) \int_{-\infty}^{\infty} p_i(t_i|m_i, x) q(x|y) dx dt \right] q_i(y|s_i) dy = \int_{-\infty}^{\infty} \left[\int_0^{\infty} u(t) p_i(t_i|m_i, x) dt \right] \left[\int_{-\infty}^{\infty} q(x|y) q_i(y|s_i) dy \right] dx. \quad (14)$$

This resembles the original expression for $U_i(m_i|s_i)$, except that $q_i(x|s_i)$ has been replaced by the average $\int_{-\infty}^{\infty} q(x|y) q_i(y|s_i) dx$ of the evaluator's beliefs. Both $q(x|y)$ and $q_i(y|s_i)$ are normal densities, and their product can be rewritten as $A_0 \exp\{-A_1 [y - A_2(x, s_i)]^2 - A_3 [x - A_4(s_i)]^2\}$, where A_0, A_1, A_3 are constants not depending on x, y, s_i , the constant A_2 depends on x, s_i only, and $A_4(s_i) = [\tau_0\tau_i s_i + (\nu + \tau_0 + \tau_i)\nu\mu]/[\tau_0\tau_i + (\nu + \tau_0 + \tau_i)\nu]$. We then find $\int_{-\infty}^{\infty} q(x|y) q_i(y|s_i) dy = A_5 \exp\{-A_3 [x - A_4(s_i)]^2\}$, where A_5 is independent of x, s_i . The forecaster's objective function (14) is of the same form as previously, where $\int_{-\infty}^{\infty} q(x|y) q_i(y|s_i) dx$ is a normal density with mean $A_4(s_i)$ strictly between μ and s_i . As in Proposition 1, forecaster i will deviate from any fully separating strategy m_i by issuing the conservative $m_i(A_4(s_i)) \neq m_i(s_i)$ for any $s_i \neq \mu$.

For the binary reputational equilibrium, suppose that all opponents apply the binary strategy with threshold μ . Again, (14) gives $U_i(m_i|s_i) = E\{W_i(m_i|x) E[q(x|y, m_{-i})|s_i]\}$, where $W_i(m_i|x)$ is precisely the same as in the proof of Proposition 3. That proof carries

over to this new situation, since $E[q(x|y, m_{-i})|s_i] \geq E[q(2\mu - x|y, m_{-i})|s_i]$ when $x \geq \mu$ and $s_i \geq \mu$. The latter fact follows from the fact that when $s_i \geq \mu$, all opponents are more likely to issue messages m_{-i} favorable to x .

Second, we turn to the forecasting contest. Suppose that all opponents use the linear strategy $\hat{m}(s) = Cs + (1 - C)\mu$, where $C \in (0, 1]$. The hypothetical observation of $y = c_i$ and of signal s_i gives two independent sources of information about x . Then $x|c_i, s_i \sim N((\nu\mu + \tau_0 c_i + \tau s_i) / (\nu + \tau_0 + \tau), 1 / (\nu + \tau_0 + \tau))$. Conditionally on $y = c_i$ and s_i , the message $\hat{m}(s_j) = Cx + C\varepsilon_j + (1 - C)\mu$ is normally distributed with mean $C(\nu\mu + \tau_0 c_i + \tau s_i) / (\nu + \tau_0 + \tau) + (1 - C)\mu$ and variance $C^2(\nu + \tau_0 + 2\tau) / [(\nu + \tau_0 + \tau)\tau]$. The density of correct opponent guesses is

$$\gamma(c_i|c_i, s_i) = \sqrt{\frac{(\nu + \tau_0 + \tau)\tau}{(\nu + \tau_0 + 2\tau)C^2 2\pi}} \exp\left[-\frac{[\nu + (1 - C)\tau_0 + \tau]^2 \tau}{2(\nu + \tau_0 + 2\tau)(\nu + \tau_0 + \tau)C^2} \left(c_i - \frac{[\nu + (1 - C)(\tau_0 + \tau)]\mu + C\tau s_i}{\nu + (1 - C)\tau_0 + \tau}\right)^2\right], \quad (15)$$

which is centered between μ and s_i . Nevertheless, when $C < 1$ this center remains closer to μ than the honest estimate $h(s_i)$ since the weight on s_i is smaller: $C\tau / (\nu + (1 - C)\tau_0 + \tau) < \tau / (\tau + \nu)$. Provided there exists a best response, this response is therefore biased away from μ , by the same logic as before.

Recall that $y|s_i \sim N((\nu\mu + \tau s_i) / (\nu + \tau), (\nu + \tau_0 + \tau) / ((\nu + \tau)\tau_0))$. The objective function $\log U_i(c_i|s_i) = \log q_i(c_i|s_i) - \log \gamma(c_i|c_i, s_i)$ is quadratic in the choice variable c_i . The first-order condition characterizing the unique maximizer is

$$\frac{[\nu + (1 - C)\tau_0 + \tau]^2 \tau}{(\nu + \tau_0 + 2\tau)C^2} \left(c_i - \frac{[\nu + (1 - C)(\tau_0 + \tau)]\mu + C\tau s_i}{\nu + (1 - C)\tau_0 + \tau}\right) = \tau_0(\nu + \tau) \left(c_i - \frac{\nu\mu + \tau s_i}{\nu + \tau}\right). \quad (16)$$

Gathering terms, (16) can be rewritten as $c_i = Ks_i + (1 - K)\mu$. The equilibrium fixed-point condition requires that the weight on s_i be equal to C , i.e.,

$$\frac{\tau_0}{\tau}(\nu + \tau)(\nu + \tau_0 + 2\tau)C^2 - [\nu + (1 - C)\tau_0 + \tau]^2 = \tau_0(\nu + \tau_0 + 2\tau)C - (\nu + (1 - C)\tau_0 + \tau)\tau. \quad (17)$$

The total coefficient on C^2 on the LHS is positive. At $C = 0$, the RHS exceeds the LHS. At $C = 1$ the opposite is true. The unique solution $C \in (0, 1)$ defines an equilibrium if it satisfies the second-order condition. The second-order condition requires that the LHS be positive, or equivalently the RHS be positive, i.e., $C > \tau(\nu + \tau_0 + \tau) / [\tau_0(\nu + \tau_0 + 3\tau)]$. This condition can be checked by inserting $\tau(\nu + \tau_0 + \tau) / [\tau_0(\nu + \tau_0 + 3\tau)]$ for C in Eq. (17) and verifying that the RHS exceeds the LHS. This criterion for equilibrium existence is then

$$\frac{(\nu + \tau)(\nu + \tau_0 + 2\tau)\tau(\nu + \tau_0 + \tau)^2}{\tau_0(\nu + \tau_0 + 3\tau)^2} - \left(\frac{(\nu + \tau_0 + \tau)(\nu + \tau_0 + 2\tau)}{\nu + \tau_0 + 3\tau}\right)^2 < \frac{(\nu + \tau_0 + 2\tau)\tau(\nu + \tau_0 + \tau)}{(\nu + \tau_0 + 3\tau)} - \frac{(\nu + \tau_0 + \tau)(\nu + \tau_0 + 2\tau)\tau}{\nu + \tau_0 + 3\tau}. \quad (18)$$

For small τ_0 this condition fails since $(\nu + \tau)^3(\nu + 2\tau)\tau^2 / (\nu + 3\tau)^2 > 0$. For large τ_0 , this condition holds since the coefficient on τ_0^2 is $-1 < 0$. \square

Proof of Proposition 9. First, we consider honesty. Supposing that all n forecasters use a linear rule $f = Hs + (1 - H)\mu$ with $H \in (0, 1)$, inversion gives $s = \mu + (f - \mu)/H$. Note that $y = V + \varepsilon_y = x + \varepsilon_y - \varepsilon_x$. Then $P - \alpha_s s_i/n$ is equal to

$$\alpha_y (x + \varepsilon_y - \varepsilon_x) + \frac{\alpha_s}{n} \sum_{k \neq i} (x + \varepsilon_k) + \alpha_\mu \mu = \frac{\alpha_y n + \alpha_s (n-1)}{n} x + \frac{\alpha_y n (\varepsilon_y - \varepsilon_x) + \alpha_s \sum_{k \neq i} \varepsilon_k}{n} + \alpha_\mu \mu. \quad (19)$$

Given signal s_i and forecast f_i , the price P is normally distributed with mean

$$\frac{\alpha_s}{n} \left[\mu + \frac{1}{H} (f_i - \mu) \right] + \frac{\alpha_y n + \alpha_s (n-1)}{n} E(x|s_i) + \alpha_\mu \mu. \quad (20)$$

Honest forecasting implies that this mean equals f_i , i.e.,

$$f_i = \frac{\alpha_y n + \alpha_s (n-1)}{n - \alpha_s/H} E(x|s_i) + \frac{\alpha_s (1 - 1/H) + \alpha_\mu n}{n - \alpha_s/H} \mu. \quad (21)$$

The condition for a symmetric equilibrium requires

$$H = \frac{\alpha_y n + \alpha_s (n-1)}{n - \alpha_s/H} \frac{\tau}{\tau + \nu}, \quad (22)$$

which is solved by the expression given in equation (6).

For the reputational cheap talk theory, the market now observes (y, f_1, \dots, f_n) . The analysis is identical to the one given in Proposition 8. The closed-form solutions for the natural language in predicting P are modified in a style similar to the above modification of the honest strategy.

In the forecasting contest, suppose that all forecasters use a linear strategy $f = Cs + (1 - C)\mu$. With an infinite number of forecasters, their average signal is at the mean, so that $P = \alpha_y y + \alpha_s x + \alpha_\mu \mu = (\alpha_y + \alpha_s)x + \alpha_y(\varepsilon_y - \varepsilon_x) + \alpha_\mu \mu$. Note here that $\alpha_s < 1$ corresponds to the relative precision of x versus y and the prior on V . Now, we have $x = [P - \alpha_y(\varepsilon_y - \varepsilon_x) - \alpha_\mu \mu] / (\alpha_y + \alpha_s)$. Given P , the forecasts are normally distributed with mean $C(P - \alpha_\mu \mu) / (\alpha_y + \alpha_s) + (1 - C)\mu$ and variance $C^2 [\alpha_y^2 (1/\tau_y + 1/\tau_x) + 1/\tau]$. This defines the density γ of opponents' forecasts. Similarly, given signal s_i , P is normally distributed with mean $(\alpha_y + \alpha_s) E(x|s_i) + \alpha_\mu \mu$ and variance $(\alpha_y + \alpha_s)^2 / (\tau + \nu) + \alpha_y^2 (1/\tau_y + 1/\tau_x)$. This defines the posterior belief density q . To find the best response, each forecaster maximizes $U(f_i|s_i) = q(f_i|s_i) / \gamma(f_i|f_i)$. The remaining analysis of the game is along the lines of Proposition 8. \square

References

- Avery, C., Chevalier, J., 1999. Herding over the career. *Economics Letters* 63, 327–333.
- Batchelor, R., Dua, P., 1992. Conservatism and consensus-seeking among economic forecasters. *Journal of Forecasting* 11, 169–181.
- Bernhardt, D., Campello, M., Kutsoati, E., 2005. Who herds? *Journal of Financial Economics* forthcoming.
- Bernhardt, D., Kutsoati, E., 2004. Analyst compensation and forecast bias. Unpublished working paper. University of Illinois and Tufts University.
- Bhattacharya, S., Pfleiderer, P., 1985. Delegated portfolio management. *Journal of Economic Theory* 36, 1–25.
- Box, G.E., Tiao, G.C., 1973. *Bayesian Inference in Statistical Analysis*. Addison-Wesley, Reading.
- Brandenburger, A., Polak, B., 1996. When managers cover their posteriors: making the decisions the market wants to see. *RAND Journal of Economics* 27, 523–541.
- Chevalier, J., Ellison, G., 1999. Career concerns of mutual fund managers. *Quarterly Journal of Economics* 114, 389–432.
- Crawford, V., Sobel, J., 1982. Strategic information transmission. *Econometrica* 50, 1431–1452.
- Croushore, D., 1993. Introducing: the Survey of Professional Forecasters. *Federal Reserve Bank of Philadelphia Business Review*, November/December, 3–13.
- Croushore, D., 1997. The Livingston Survey: still useful after all these years. *Federal Reserve Bank of Philadelphia Business Review*, March/April, 15–26.
- DeGroot, M., 1970. *Optimal Statistical Decisions*. McGraw Hill, New York.
- Ehrbeck, T., Waldmann, R., 1996. Why are professional forecasters biased? Agency versus behavioral explanations. *Quarterly Journal of Economics* 111, 21–40.
- Frankel, J.A., Froot, K.A., 1987. Using survey data to test standard propositions regarding exchange rate expectations. *American Economic Review* 77, 133–153.
- Gleason, C.A., Lee, C.M.C., 2003. Analyst forecast revisions and market price discovery. *Accounting Review* 78, 193–225.
- Graham, J., 1999. Herding among investment newsletters: theory and evidence. *Journal of Finance* 54, 231–268.
- Granger, C.W.J., 1969. Prediction with a generalized cost of error function. *Operational Research Quarterly* 20, 199–207.
- Granger, C.W.J., 1999. Outline of forecast theory using generalized cost functions. *Spanish Economic Review* 1, 161–173.

- Hirshleifer, D., Teoh, S.H., 2003. Limited attention, financial reporting, and disclosure. *Journal of Accounting and Economics* 36, 337–386.
- Holmström, B., 1999. Managerial incentive problems: a dynamic perspective. *Review of Economic Studies* 66, 169–182.
- Hong, H., Kubik, J.D., 2003. Analyzing the analysts: career concerns and biased earnings forecasts. *Journal of Finance* 58, 313–351.
- Hong, H., Kubik, J.D., Solomon, A., 2000. Security analysts' career concerns and herding of earning forecasts. *RAND Journal of Economics* 31, 121–144.
- Hotelling, H., 1929. Stability in competition. *Economic Journal* 39, 41–57.
- Johnson, N.L., Kotz, S., 1970. *Distributions in Statistics: Continuous Univariate Distributions*, Vol. 1. Wiley, New York.
- Kahneman, D., 1973. *Attention and Effort*. Prentice-Hall, Englewood Cliffs, NJ.
- Kandel, E., Pearson, N.D., 1995. Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy* 103, 831–872.
- Kandel, E., Zilberfarb, B.Z., 1999. Differential interpretation of information in inflation forecasts. *Review of Economics and Statistics* 81, 217–226.
- Keane, M.P., Runkle, D.E., 1990. Testing the rationality of price forecasts: new evidence from panel data. *American Economic Review* 80, 714–735.
- Keane, M.P., Runkle, D.E., 1998. Are financial analysts' forecasts of corporate profits rational? *Journal of Political Economy* 106, 768–805.
- Lamont, O., 2002. Macroeconomic forecasts and microeconomic forecasters. *Journal of Economic Behavior and Organization* 48, 265–280.
- Laster, D., Bennett, P., Geoum, I.S., 1999. Rational bias in macroeconomic forecasts. *Quarterly Journal of Economics* 114, 293–318.
- Lehmann, E.L., 1955. Ordered families of distributions. *Annals of Mathematical Statistics* 26, 399–419.
- Lim, S.S., 2002. Analysts' reputational concerns and underreaction to public news: theory and evidence. Unpublished working paper. Fisher College of Business, Ohio State University.
- Michaely, R., Womack, K.L., 1999. Conflict of interest and credibility of underwriter analyst recommendations. *Review of Financial Studies* 12, 654–686.
- Milgrom, P., 1981. Good news and bad news: representation theorems and applications. *Bell Journal of Economics* 12, 380–391.
- Morgan, J., Stocken, P., 2003. An analysis of stock recommendations. *RAND Journal of Economics* 34, 183–203.

- Muth, J.F., 1961. Rational expectations and the theory of stock price movements. *Econometrica* 29, 315–335.
- Osborne, M., Pitchik, C., 1986. The nature of equilibrium in a location model. *International Economic Review* 27, 223–237.
- Ottaviani, M., Sørensen, P.N., 2005, Reputational cheap talk. *RAND Journal of Economics* forthcoming.
- Prendergast, C., 1993. A theory of ‘yes men’. *American Economic Review* 83, 757–770.
- Prendergast, C., Stole, L., 1996. Impetuous youngsters and jaded oldtimers: acquiring a reputation for learning. *Journal of Political Economy* 104, 1105–1134.
- Romer, C.D., Romer, D.H., 2000. Federal Reserve information and the behavior of interest rates. *American Economic Review* 90, 429–457.
- Scharfstein, D., Stein, J., 1990. Herd behavior and investment. *American Economic Review* 80, 465–479.
- Stark, T., 1997. Macroeconomic forecasts and microeconomic forecasters in the Survey of Professional Forecasters. Working Paper 97-10. Federal Reserve Bank of Philadelphia.
- Steele, J.M., Zidek J., 1980. Optimal strategies for second guessers. *Journal of the American Statistical Association* 75, 596–601.
- Stekler, H.O., 1987. Who forecasts better? *Journal of Business and Economics Statistics* 5, 155–158.
- Trueman, B., 1994. Analyst forecasts and herding behavior. *Review of Financial Studies* 7, 97–124.
- Welch, I., 2000. Herding among security analysts. *Journal of Financial Economics* 58, 369–396.
- Zarnowitz, V., Lambros, L.A., 1987. Consensus and uncertainty in economic predictions. *Journal of Political Economy* 95, 591–621.
- Zellner, A., 1986. Biased predictors, rationality and the evaluation of forecasts. *Economics Letters* 21, 45–48.
- Zitzewitz, E., 2001a. Measuring herding and exaggeration by equity analysts. Unpublished working paper. Graduate School of Business, Stanford University.
- Zitzewitz, E., 2001b. Opinion-producing agents: career concerns and exaggeration. Unpublished working paper. Graduate School of Business, Stanford University.