Forecasters' Objectives and Strategies[*]

Iván Marinovic[†]   Marco Ottaviani[‡]   Peter Norman Sørensen[§]

Chapter Prepared for the Handbook of Economic Forecasting, Volume 2

July 2011 (revised June 2012)

Abstract

This chapter develops a unified modeling framework for analyzing the strategic behavior of forecasters. The theoretical model encompasses reputational objectives, competition for the best accuracy, and bias. Also drawing from the extensive literature on analysts, we review the empirical evidence on strategic forecasting and illustrate how our model can be structurally estimated.

Keywords: Reputational cheap talk, forecasting contest, herding, exaggeration, bias.

[†]Graduate School of Business, Stanford University; imvial@stanford.edu.

[‡]Bocconi University and Kellogg School of Management, Northwestern University; marco.ottaviani@unibocconi.it.

[§]University of Copenhagen; peter.sorensen@econ.ku.dk.

Dans ses écrits, un sage Italien

Dit que le mieux est l'ennemi du bien

In his writings, an Italian sage

says that the best is the enemy of the good

Voltaire, La Bégueule, Contes, 1772

1 Introduction

Forecasting as a scientific enterprise has a long history, arguably dating back to the Chaldeans, the Mesopotamian astronomers of Assyria and Babylonia. Throughout a number of centuries, the Chaldeans collected a vast amount of astronomical observations, historical events, and market outcomes with the objective of using correlations to discover regularities and forecast future variables such as the water level of the river Euphrates, the prices of agricultural commodities, and political events.[1]

Since those early days, forecasting has developed to become a thriving industry in areas ranging from meteorology to economics. Given the important role of forecasters as information providers to decision makers, a lot of attention has also been devoted to the evaluation of their services. In fact, forecasters who attain outstanding reputation face remarkable career prospects. For example, a number of recent key members of the Board of Governors of the Federal Reserve Bank had a successful career as economic forecasters.[2]

Market participants, the popular press, as well as researchers actively monitor the accuracy of forecasters. In a disarming piece published in the first volume of Econometrica, Alfred Cowles (1933) was one of the first scholars to document the poor track record of stocks recommendations and stock price predictions released by a number of prominent pundits and professional agencies.[3]

---

[1] See, for example, Slotsky (1997).

[2] Until he was appointed at the Federal Reserve Board, for most of his career Alan Greenspan was Chairman and President of Townsend-Greenspan & Co., Inc., an economic consulting firm in New York City that offered forecasts and research to large businesses and financial institutions. Similarly, before becoming a member of the Fed Board, Laurence Meyer was President of Laurence H. Meyer and Associates, a St. Louis-based economic consulting firm specializing in macroeconomic forecasting and policy analysis.

[3] Denton (1985)—a hidden gem—provides an entertaining and insightful theory of the role of (un)professional forecasters for the stability of speculative markets.

But is this evaluation sufficient to ensure the performance and truthfulness of forecasters? After all, forecasters are economic agents who make strategic choices. To be able to interpret the content of forecasts, it is essential to understand the role of the economic incentives forecasters face.

On the one hand, forecasters who wish to be perceived as well informed may be reluctant to release information that could be considered inaccurate, when contrasted with other sources of public information. Forecasters would then shade their forecasts more toward the established consensus on specific indicators to avoid unfavorable publicity when wrong.

On the other hand, the fact that only the most accurate forecaster obtains a disproportionate fraction of public attention induces a convex incentive scheme whereby the payoff of being the best is significantly higher than that of being the second best. Forecasters might then exaggerate their true predictions, on the off chance of getting it right so as to be able to stand out from the crowd of competing forecasters. By exaggerating their private information, forecasters reduce the probability of winning but they increase their visibility conditional on winning. Being the single winner always entails more glory than sharing the prize with other fellows.

How do these incentives shape the forecasters' behavior? In particular how can we interpret professional forecasts, and how informative should we expect them to be? This chapter develops a simple framework to analyze the statistical properties of strategic forecasts that result from a combination of the reputational objective with the contest objective. We find that conservatism or exaggeration arises in equilibrium depending on whether reputational incentives or contest incentives prevail, while truthful forecasting results in a knife-edge case.

It is worth remarking from the outset that our model can be seen as providing micro-foundations for the specification of asymmetric loss functions. As noted by Granger (1969) and Zellner (1986), forecasters with asymmetric loss functions optimally issue forecasts that violate the orthogonality property. Elliott, Komunjer, and Timmermann (2005) develop a methodology for estimating the parameters of the loss function given the forecasts. The approach we develop here is similar in spirit.

There is also a parallel literature in statistics and game theory on calibration and expert testing with a different focus from ours. The expert testing literature, pioneered by Foster

and Vohra (1998) and recently surveyed by Olszewski (2012), studies the following question in an infinite horizon setting. Is there a test that satisfies: (a) a strategic forecaster who ignores the true data generating process cannot pass the test and (b) a forecaster who knows the data generating process passes the test with arbitrarily large probability no matter what the true data generating process may be? To answer this question, Foster and Vohra (1998) first examine Dawid's (1982) calibration test.[4] By choosing a suitably complex mixed strategy, they prove that an ignorant forecaster can attain virtually perfect calibration.[5] Compared to the expert testing literature, our model streamlines the dynamic structure but allows forecasters to have noisy information of varying precision about a continuous state, as is natural for applications. In the reputational cheap talk model, the evaluator is also a player and performs the ex post optimal test—the focus is then on the characterization of the (lack of) incentives for truthful reporting for a forecaster with information and on the resulting equilibrium of the game without commitment on the test performed. Forecasting contest theory, instead, posits a relatively simple comparative test across forecasters, justified on positive grounds.

The chapter is structured as follows. Section 2 introduces the model. Section 3 reviews the literature developments that led to the two theories encompassed by the model. Section 4 derives the equilibrium depending on the weight forecasters assign to the reputational payoff relative to the contest payoff and characterizes the bias, dispersion, and orthogonality properties of equilibrium forecasts. Section 5 shows how the model can be estimated. Section 6 discusses a number of extensions and robustness checks. Section 7 discusses the possibility of comparing anonymous surveys of professional forecasters with non-anonymous surveys in order to test for strategic reporting. Section 8 concludes.

---

[4]Suppose nature chooses a probability of rain every period and a forecaster, similarly, tries to predict this probability. Dawid (1982) proposes a calibration rule to evaluate the forecaster. A forecaster is well calibrated if for every day that she predicted rain with probability x% the empirical frequency is also x%. Calibration is a minimal statistical requirement of good forecasting. Predicting the long-term frequency every day would be enough for the forecaster to pass the calibration test.

[5]Relying on the minimax theorem, Sandroni (2003) shows that this negative result holds for all finite horizon tests. A number of papers have examined the robustness of this result showing that the manipulability of tests disappears when one restricts the set of stochastic functions used by nature, or when the tester is Bayesian (see Al-Najjar et al. (2010)), or in mechanism design-like contexts. Finally, Fortnow and Vohra (2007) prove positive results based on algorithmic complexity of theories. Roughly, there is a test that cannot be manipulated if the expert is restricted to submit (algorithmically) simple forecasts.
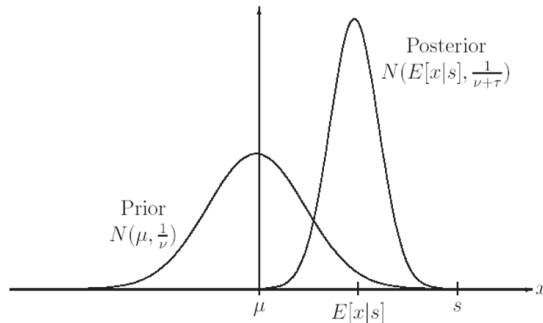
Figure 1: Posterior Distribution

2 Model with Mixed Reputational and Contest Payoffs

We model strategic forecasting through a game played by a continuum of forecasters indexed by $i \in [0, 1]$. The forecasters share a common prior about the uncertain state of the world, $x \sim N(\mu, 1/\nu)$, as represented in Figure 1. Initially, each forecaster i privately observes a signal $s_i \sim N(x, 1/\tau)$, independently distributed conditional on x.[6] Each forecaster then simultaneously and independently releases the forecast f_i. Finally, the state x is realized and the forecasters are rewarded on the basis of the evaluation made by the market.

To encompass the two main theories of strategic forecasting developed by Ottaviani and Sørensen (2006c) (referred to as OS from now on) we posit that the forecasters' objective function tractably combines a reputational component, u_r, with a contest component, u_c, through a geometric average:

$$U = u_r^{\delta} u_c^{1-\delta}.$$

For future reference define $\theta \equiv \frac{\delta}{1-\delta}$ as the intensity of reputational concerns, so that $1/\theta$ represents the intensity of competition.

2.1 Reputational Signaling Payoff

The *expected reputational signaling payoff* for an individual issuing a forecast f after observing signal s is

$$u_r = \exp\left\{ -E\left[\hat{E}\left(x - s | f, x\right)^2 | s \right] \right\}, \tag{1}$$

---

[6]To interpret independence with a continuum of forecasters, signals are actually distributed in the population as posited. Any one forecaster constitutes a null set, so the distribution is unchanged by removing this one forecaster, as is the point of view in a Bayesian Nash equilibrium.

where E represents the forecaster's expectation and $\hat{E}$ represents the market's expectation. According to this formulation, the forecaster's reputational loss increases in the market's assessment of the mean squared error between the observed state x and the market's inference on the unobserved signal s given the observed forecast f. As explained by OS (see the model formulation, Proposition 1, and its proof in the appendix), the motivation for this assumption is based on the forecaster's reputational incentive to appear to have precise information.

In their model of reputational forecasting, OS posit that the market employs the forecast to optimally learn about the true underlying precision t of the forecaster's private information. This precision is ex ante unknown to both the market and the forecaster, being distributed according to a prior $p(t)$, with expectation equal to the known parameter τ, the expected precision introduced above. The true precision represents the quality of the forecaster's information because forecasters with higher t observe signal realizations s that are closer to the true state, x; intuitively, a forecaster with maximal precision observes a signal exactly equal to the state. The market updates the prior $p(t)$ into the posterior $p(t|f, x)$.

To this end, the market formulates a conjecture about the forecaster's strategy $f(s)$ that maps signals s into forecasts f, and uses this conjecture and the reported forecast f to recover the signal s the forecaster has privately observed. In the equilibrium we construct below, the forecasting strategy is strictly increasing. The market recovers the signal $s = f^{-1}(f(s))$ by inverting this strictly increasing function. The market's belief will be correct in equilibrium, but out of equilibrium, a forecaster with private signal s who deviates from $f(s)$ to some other $f(s')$ will be treated by the market as if in possession of signal s'.

The market then confronts this recovered signal with the realized state. When t parametrizes the precision of information in a location experiment, the closer the recovered signal is to the realized state, the more precise the forecaster is thought to be. The reputational loss is then an increasing function of the inferred signal error $|x - s|$, as shown by OS. This construction justifies the reputational objective (1).

2.2   Forecasting Contest Payoff

The expected contest payoff of a forecaster who observes signal s and reports forecast f is

$$u_c \propto \frac{\phi\left(f|s\right)}{\psi\left(f|f\right)}, \tag{2}$$

where:

- $\phi\left(f|s\right)$ is the forecaster's probability belief that the state is equal to the forecast $x = f$ conditional on signal s and

- $\psi\left(f|f\right)$ is the fraction of other forecasters issuing forecast f, conditional on the realization of the state being equal to that forecast, $x = f$.[7]

As explained by OS, in a winner-take-all forecasting contest there is one prize whose size is proportional to the number of forecasters.[8] The prize is shared among the forecasters whose forecast turns out to be closer to the realized state. With a continuum of forecasters who each uses a forecasting strategy that is strictly increasing in the signal they observe, all signals in the real line of the support of the normal distribution are observed so that all forecasts on the real line are made.

In this setting with a continuum of forecasters, each forecaster wins only when the forecast is spot on the mark and then shares the prizes with all the forecasters who are also spot on the mark. Each forecaster's payoff is then equal to the ratio of:

- the probability that the state is equal to the forecast reported—the numerator in (2) is the continuum analogue of the probability of winning with a forecast f conditional on the signal being s; and

- the number of accurate forecasts made by the other competing forecasters—the denominator is the continuum analogue of the mass of winners who issue forecast f conditional on the state being $x = f$.

Intuitively, in a forecasting contest the forecaster benefits when making an accurate forecast and when few of the competing forecasters are accurate. As shown by Otta-

---

[7]Note that $\psi\left(f|f\right)$ is not a probability density of the forecasts made by the other forecasters because the conditioning event f is contemporaneously changed.

[8]Equivalently, there is a number of prizes proportional to the number of forecasters.

viani and Sørensen (2005), the contest payoff (2) arises in the limit of a winner-take-all forecasting contest as the number of forecasters tends to infinity.

3   Development of Reputational and Contest Theories

Before analyzing the model with mixed payoff in the next section, in this section we zoom in on the reputational signal theory ($\theta = 0$) and the forecasting contest theory ($\theta = 1$) in isolation. For each of these two theories we present the developments of the literature that led to their formulation. We also summarize the intuitive logic that underlies the main results.

3.1   Reputational Cheap Talk

A basic premise of the reputational theory is that forecasters who are perceived to have more accurate information obtain an improved compensation for their forecasting services. The key question is then: Why would the reputational incentive lead a forecaster to manipulate the information reported? At first blush it might seem that a forecaster's reputation for accuracy is maximized through truthful (or honest) forecasting whereby the forecaster reports the posterior expectation about the state conditional on the signal privately observed. After all, the conditional expectation is the most accurate forecast, expected to minimize the mean squared error between the state x and the forecast f. However, the answer to the question is more subtle.

The possibility of deviation from honest reporting has been put forward in the Bayesian statistics literature. For example, Bayarri and DeGroot (1989) consider a decision maker who asks a group of experts to individually report a predictive distribution for the observable variable x. Each expert is assigned a weight which is then updated based on the observation of x. In this setting, the expert has an incentive to choose the reported distribution in order to maximize the updated weight obtained, rather than predicting honestly. Bayarri and DeGroot, then, characterize the incentive to deviate from honest reporting in the context of an example—however, they do not investigate the optimal reaction by the market who performs the evaluation, and the resulting equilibrium.

OS provide a simple mechanism that drives forecasters to generally deviate from truthful reporting. A key ingredient is that everyone, market participants and forecasters alike,

shares knowledge of some public information, whether CNN headlines, Financial Times articles, or others. This common information creates baseline expectations regarding several economic variables. Importantly, these baseline expectations usually point in the right direction. But professional forecasters should be able to improve upon what is commonly known because they have access to additional information about the economy's state. Thus, forecasters combine what everybody knows (public information) with what they alone know (private information) to generate their personal predictions. As such, forecasters' personal predictions typically fall between public and private beliefs. Personal predictions deviate from public expectations on the basis of private information, but are not as exaggerated as the private information might be because they take public information into account. If forecasters were not strategic, they would honestly report their personal predictions, i.e., their conditional expectations $f = E(x|s)$. Market participants would then be right to take these forecasts at face value.

To understand the mechanism that drives forecasters to deviate from truthful reporting, suppose with OS that the evaluator conjectures that the forecaster truthfully reports the conditional expectation. If the market believes that forecasters are honest and rewards them based on their reputation for accuracy, will forecasters be content to truthfully report their personal predictions? A forecaster's payoff depends on the reputation about the quality of information t, which is updated on the basis of the realized state and the signal the market infers the forecaster had. Given that the truthful forecast $f = E(x|s)$ is strictly increasing in s (being a weighted average of the prior μ and the signal s in our normal learning model) the market infers the signal $s = f^{-1}(E(x|s)) = [(\tau + \nu)f - \nu\mu]/\tau$, and then uses it in combination with the realized state x to update the prior on the forecaster's talent, $p(t)$.

However, and this is the crux of the argument, the forecaster maximizes her expected reputation by pretending to have observed a signal equal to $\hat{s} = E(x|s)$, which is the closest possible signal to the state the market will observe. The optimal deviation for the forecaster consists in reporting a conservative forecast $E[x|\hat{s} = E(x|s)] < E(x|s)$, so that the inferred signal would actually be $E(x|s)$. In other words, a forecaster can be perceived as more accurate by eliminating the predictable error $s - E[x|s] > 0$.[9] As a result of

---

[9]Given that the market conjectures truthful forecasting $f = E[x|s]$ and infers signal s, by sticking to truthful forecasting the forecaster suffers an average error equal to $s - E[x|s] > 0$.

this deviation, the issued forecast would be more conservative than the truthful forecast, assigning a larger weight to the public information μ and a smaller weight to the private signal s. As illustrated in Figure 1, however, the forecaster has an incentive to deviate.[10]

Taking into account all the information they have, forecasters expect their personal predictions to be correct. To convince the market that their private information is accurate forecasters would like the market to believe that their private information is located at their personal prediction. In other words, if forecasters can convince the market that their predictions are based fully on private information, they would be considered even better informed than they really are.

The incentive to deviate from truthful reporting arises because of the interplay between public and private information. Indeed, the forecaster's incentive to deviate originates from the fact that the market has an incentive to filter out the prior μ from the forecast to estimate the forecaster's signal s. In turn, the forecaster wants to use the prior to better forecast the state x. The conflict of interest between the forecaster and the market lies here—the forecaster wants to weigh the public information more than the market wants. Consequently, forecasters have an incentive to confirm the original belief of the market by making predictions closer to the prior consensus than their expectation. If so, the market's original belief that forecasters report honestly their conditional expectations is not consistent with the actual behavior of the forecasters.

The next natural step is to consider the Bayes Nash equilibrium of this game. This is a game of "cheap talk" because the forecast enters the forecaster's payoff only through the market's inference of the forecaster's signal. Truthful forecasting (and, more generally, a fully separating equilibrium) is not sustainable in this cheap talk game. For this reputational cheap talk game, OS showcase a partially separating equilibrium whereby the forecaster can credibly communicate the direction of her signal, but not its intensity. Thus, the incentive toward conservativeness does not persist in equilibrium, which is only partially separating. As a result, only partial learning about the state of the world (as well as about the forecaster's true precision) takes place in equilibrium.

If the market is fully rational, it will be able to anticipate that forecasters are distorting their predictions to pretend to be more informed than they really are. As a consequence,

---

[10]Ottaviani and Sørensen (2006b) provide a general characterization of this deviation incentive, valid beyond the normal information model.

the market can only trust forecasters to communicate part of the information they have. Paradoxically, the desire of analysts to be perceived as good forecasters turns them into poor forecasters. In line with this, *The Economist* magazine reports the "surprisingly good performance of a sample of London garbage men in forecasting key economic variables." Presumably the garbage men were free of reputation-focused incentives.

The impossibility of full revelation in a reputational cheap talk equilibrium is reminiscent of Crawford and Sobel's (1982) result for games of cheap talk in which sender and receiver have exogenously different objective functions. As explained by Ottaviani and Sørensen (2006b), in our reputational setting the divergence in the objective functions arises endogenously depending on the informational assumptions. For example, if there is diffuse (or no) public information, the forecaster has no incentive to deviate from truthful forecasting. But in a realistic setting in which both private and public information are relevant for forecasting, the information that can be credibly transmitted in a cheap talk equilibrium is endogenously coarse.

Differently from the cheap talk case considered by OS and also considered in this section, the setting with mixed objectives introduced in Section 2 corresponds to a signaling game in which the forecast also directly affects the forecaster's objective through the contest payoff, in addition to the indirect effect through the inferred signal $s = f^{-1}(f)$. As we will see in the next section, provided the reputational incentive is not too dominant, a fully separating equilibrium results in our full model, and the deviation incentive has a direct impact on the equilibrium strategy. At a theoretical level, the mechanism that turns cheap talk into costly talk is similar to the one investigated by Kartik, Ottaviani, and Squintani (2007) in the context of Crawford and Sobel's (1982) model of partisan advice.

The characterization of the incentive toward a conservative deviation we have highlighted above has been first derived by OS. This result is distinct from the claim by Scharfstein and Stein (1990) that reputational concerns induce herding behavior. Adding some interesting structure to a model formulated by Holmström (1999) in a paper that circulated since 1982, Scharfstein and Stein (1990) consider a streamlined sequential reputational cheap talk model in which two agents (corresponding to our forecasters) make investment decisions one after the other. These agents (like our forecasters) are only interested in the inference that is made by a principal (the market in our setting) about the

11

quality of their information. Scharfstein and Stein (1990) argue that the second agent in (a reputational cheap talk) equilibrium will decide solely based on past choices, disregarding the information privately held, provided that better informed agents observe signals that are positively correlated conditional on the state. Essentially, under these conditions the cheap talk equilibrium is completely pooling. Thus, in Scharfstein and Stein (1990) herding results because a second forecaster's incentive to imitate a first forecaster who reported earlier destroys the second forecaster's ability to report an informative forecast. Our analysis, instead, focuses on the reporting incentives of a single forecaster in a setting with a continuous signal, highlighting that truthful forecasting is incompatible with cheap talk equilibrium and that some information can always be transmitted in the normal model (and, more generally, in a location model).[11]

Empirically, there is some evidence of a negative relation between the degree of herding and analyst experience (e.g., Hong, Kubik, and Solomon 2000 and Clement and Tse 2005). These findings are consistent with reputational herding given that younger analysts still have to earn their reputation. There is also evidence that past performance has a negative impact on the propensity to herd (Stickel 1990 and Graham 1999). Also, the uncertainty about the environment is found to be positively related to herding (Olsen 1996). Finally, and relatedly, the forecast horizon has been shown to have a positive impact on herding (Krishnan, Lim, and Zhou 2006). As the forecast horizon decreases, more information becomes available, thereby reducing information uncertainty and herding behavior.

3.2 Contest Theory

Forecasting contests are often run among meteorologists (for example, consider the National Collegiate Weather Forecasting Contest) and economists (see the Wall Street Journal semi-annual forecasting survey) by assigning prizes to the forecasters who is closest to the mark. Forecasters in these contests are rewarded depending on their relative accuracy level. Businesses (such as Corning) also use forecasting contests as a design for prediction

---

[11]As shown by Ottaviani and Sørensen (2000), Scharfstein and Stein's (1990) assumption that better informed agents observe signals that are conditionally more positively correlated is not necessary to obtain herding. Instead, in such a sequential setting, herding can also be obtained when the signals observed by the better informed agents are independent conditional on the state through a mechanism similar to statistical herding that has been highlighted by Banerjee (1992) and Bikhchandani, Hirschleifer, and Welch (1992). This can be the case also when Scharfstein and Stein's (2000) stronger definition of herding is adopted whereby the second agent takes the same action (be it investment or forecast) as the first agent whatever the action of the first agent is, as shown by Ottaviani and Sørensen (2006a, Section 6).

markets with the aim of collecting decision-relevant information from inside or outside experts.

In a pioneering piece, Francis Galton (1907) reports an entertaining account of an ox-weighing competition in Plymouth. Francis Galton collected the guesses provided by participants in the contest on stamped and numbered cards, and found the median of the individual estimates to be within 1% of the correct weight. Galton justified the interest in the contest, "a small matter" in itself, on the grounds of the result's relevance to the assessment of the trustworthiness of democratic judgements.

In the applied probability literature, Steele and Zidek (1980) analyze a simple sequential forecasting contest between two forecasters who guess the value of a variable (such as the weight of a party participant). The second guesser not only possesses private information on the variable whose value is to be predicted, but also observes the first mover's forecast. Abstracting away from strategic problems, this work characterizes the second guesser's advantage in the winning probability over the first guesser.[12]

Ottaviani and Sørensen (2005) and OS analyze forecasting contests in which forecasts are issued simultaneously. As discussed above, in a limit winner-take-all forecasting contest a forecaster wants to maximize the expected prize from participating in the contest, which is the ratio of the probability of winning to the density of the winning forecasts. Whenever the prior belief is not completely uninformative as in the case with improper prior, OS show that the denominator $\psi(f|f)$ decreases in the distance between the forecast (equal to the realized state) and its prior mean μ. In fact, ψ is centered around μ because every forecaster assigns a positive weight to the common prior. The probability of winning the contest is maximized at $E(x|s)$. However, at that point, the posterior is flat, while the number of winning opponents is decreasing. Then, at $E(x|s)$, it is optimal to deviate to issuing a forecast f which is closer to s than $E(x|s)$ because the first-order reduction in the expected number of winners with whom the prize must be shared more than compensates the second-order reduction in the probability of winning.

Intuitively, forecasters have an incentive to distance themselves from market consensus on the off chance of being right when few other forecasters are also right. Indeed when forecasters merely repeat what everybody else is already saying, they stand to gain little, even when they are right. Competition for the best accuracy record induces a tendency to

---

[12]This "second guessing" problem is also analyzed by Pittenger (1980) and Hwang and Zidek (1982).

exaggerate, rather than to be conservative. The incentive to deviate goes in the opposite direction compared to the reputational objective.

Unlike in the case of reputational cheap talk, the contest payoff induces a direct link between a forecaster's payoff and the forecast. Thus, this is essentially a signaling game similar to an auction. Because the link between forecast and payoff is direct, the incentive to deviate persists in the equilibrium of the pure forecasting contest, unlike in the pure reputational cheap talk setting. OS show that there is a unique symmetric linear equilibrium, in which exaggeration takes place. In the next section we generalize this linear equilibrium for the mixed model with both reputational and contest incentives.

Just like in the reputational theory, the incentive to deviate from honest reporting originates from the availability of both public and private information. In fact, if there is no public information, $\psi(f|f)$ is flat, and it is optimal to issue $f = E(x|s)$. Instead, if the private signals are uninformative, with infinitely many symmetrically informed players, the distribution of the equilibrium locations replicates the common prior distribution, as shown formally by Osborne and Pitchik (1986). As a matter of fact, in the absence of private information, the forecasting game is identical to Hotelling's (1929) location game. Indeed, Laster, Bennet, and Geoum (1999) obtain Hotelling's infinite player result in a simultaneous forecasting contest in which forecasters have perfectly correlated information or, equivalently, they do not have private information but only observe the common prior μ.

In sum, if forecasters care about their relative accuracy, like when they compete for ranking in contests, they will tend to exaggerate predictions in the direction of their private information. In contrast, the reputational cheap talk theory applies if forecasters are motivated by their absolute reputation for accuracy. Forecasters would then be expected to align their predictions more closely with publicly available information. To determine which theory better explains the data we turn to the model with mixed incentives.

4   Equilibrium with Mixed Incentives

Having privately observed the signal s, a forecaster chooses a forecast f that maximizes U. The market, upon observing the forecast and the outcome realization updates its expectations using Bayes' rule. The structure of the game is common knowledge. We look for a perfect Bayesian (Nash) equilibrium. Since forecasters are symmetric we focus on

symmetric equilibria. The following proposition establishes the existence of a linear equilibrium.

Proposition 1 *If $\theta < \frac{\tau+v}{2v}\tau \equiv \theta^+$, there exists a unique symmetric equilibrium in increasing linear strategies,*

$$f = [1 - \alpha(\theta)]\mu + \alpha(\theta)s,$$

where

$$\alpha(\theta) = \frac{\sqrt{\tau(4v+\tau) - \frac{8\theta v^2}{\tau+v}} - \tau}{2v}. \tag{3}$$

In the limit, as θ approaches θ^+ the equilibrium becomes uninformative: $\lim_{\theta \to \theta^+} \alpha(\theta) = 0$.

Proof. We conjecture an equilibrium in which forecasters use linear strategies

$$f(s) = (1 - \alpha)\mu + \alpha s$$

with $\alpha > 0$. After seeing signal s_i, a forecaster chooses forecast f_i to maximize

$$\frac{\ln U_i}{1 - \delta} = -\theta E\left\{ [x - f^{-1}(f_i)]^2 \,|s_i \right\} + \left(-\frac{[f_i - E(x|s_i)]^2}{2/\tau(x|s_i)} + \frac{[f_i - E(f|x = f_i)]^2}{2/\tau(f|x = f_i)} \right),$$

where $\tau(x|s_i)$ denotes the precision of x conditional on signal s_i and $\tau(f|x = f_i)$ denotes the precision in the distribution of others' forecasts conditional on realized state $x = f_i$. Without loss of generality we assume that $\mu = 0$. Differentiating with respect to f_i we get the first order condition

$$-\theta 2\frac{-1}{\alpha}\left(\frac{\tau s_i}{\tau + v} - \frac{f_i}{\alpha} \right) - (\tau + v)\left(f_i - \frac{\tau s_i}{\tau + v} \right) + \frac{\tau}{\alpha^2}(1 - \alpha)^2 f_i = 0,$$

where we used that $\tau(x|s_i) = \tau + v$ and $\tau(f|x = f_i) = \frac{\tau}{\alpha^2}$. Solving the above equation yields the linear relation

$$\tau\frac{2\theta + \alpha\tau + \alpha v}{\tau + v}\alpha s_i = \left(2\theta + v\alpha^2 + 2\alpha\tau - \tau \right) f_i,$$

and by equating coefficients with the conjectured strategy $f = \alpha s$, we arrive at the quadratic equation

$$\tau\frac{2\theta + \alpha\tau + \alpha v}{\tau + v} = 2\theta + v\alpha^2 + 2\alpha\tau - \tau. \tag{4}$$

Solving this equation, note first that the right-hand side is positive for any $\alpha > 0$. This implies that the second order condition for the quadratic optimization problem of forecaster

15

i is satisfied. Next, the quadratic equation (4) has a unique positive solution if and only if its left-hand side exceeds the right-hand side at $\alpha = 0$. This is equivalent to $\theta < \frac{\tau+v}{2v}\tau \equiv \theta^+$. The positive solution to equation (4) is then given by (3). ∎

Several features of the equilibrium are worth noting. The existence of a linear equilibrium depends on reputational concerns being sufficiently low. When θ is above θ^+ the equilibrium ceases to be fully revealing and forecasting becomes coarse, as in OS's "pure" reputational cheap talk model. This result suggests that reputational concerns paradoxically may jeopardize the possibility of informative forecasts. When reputational concerns are overwhelming, fully revealing a signal may be too costly for the forecaster, especially when the signal is too far from the prior expectation. At the other extreme, when reputational concerns are negligible, the equilibrium converges to OS's contest. Thus, when reputational concerns are low the existence of a fully revealing equilibrium is guaranteed.

In the remainder of this section we investigate the statistical properties of the equilibrium forecasts: bias (Section 4.1), dispersion (Section 4.2), and orthogonality (Section 4.3)

4.1 Forecast Bias

Forecasts are a weighted average of public and private information but forecasts may be biased relative to forecasters' posterior beliefs. We now study how this bias depends on the incentives of forecasters.

Definition 2 *The* (conditional) forecast bias *is defined as*

$$b \equiv f - E\left(x|s\right) = \left(s - \mu\right)\left[\alpha\left(\theta\right) - \frac{\tau}{\tau + v}\right],$$

while the average forecast bias *is defined as* $E\left(b\right)$.

The bias is thus defined relative to a forecaster's posterior expectation

$$E\left(x|s\right) = \frac{v\mu + \tau s}{\tau + v}.$$

A forecasts that is conditionally biased, is inefficient in the sense that the forecast does not minimize the mean squared error. In essence such forecasts do not exploit in full the available information.

16

Note that equilibrium forecasts are unbiased *on average*, $E(b) = 0$. That is forecasters do not have a systematic tendency to be over-optimistic or over-pessimistic, as would be the case if they wished to influence market beliefs in a particular direction, as in the extension analyzed in Section 6.4.[13]

Yet, even though forecasts are unbiased on average, they are generically biased given any realization of the signal. Indeed, forecasters may under-react to their private information, shading their forecasts toward public information. In this case, *conservatism* (often referred to as herding) would result. In the presence of conservatism, forecasts cluster around the consensus forecast, which is defined to be the average forecast across all forecasters. Conversely, *exaggeration* (often referred to as anti-herding) results when forecasters over-react to their private information generating forecasts that are excessively dispersed.

Corollary 1 *(i) Forecasts are (conditionally) unbiased if and only if $\theta = \overline{\theta} \equiv \frac{1}{2}\frac{\tau v}{\tau + v}$, where $\overline{\theta} < \theta^+$. If $\theta < \overline{\theta}$ forecasters exaggerate by assigning excessive weight to their private information. By contrast, if $\theta > \overline{\theta}$, forecasters are conservative by shading their forecast toward the prior. (ii) The extent of conservatism increases in the reputational concern: $\frac{\partial \alpha(\theta)}{\partial \theta} \leq 0$.*

According to Corollary 1, conservatism is associated with reputational concerns. Forecasters are aware that their signals are noisy but would like to conceal this fact from the market to avoid being perceived as imprecise. A forecast that strongly deviates from the consensus has a high chance of resulting in a low level of reputation ex post. So the forecaster may choose to inefficiently dismiss the private information possessed when this information strongly differs from the consensus. The forecast issued is then more aligned with the consensus compared with the forecaster's conditional expectation. For example, when forecasters observe private signals that are favorable relative to the consensus, they conclude that they may have realized a positive error. The less accurate the forecaster, the higher the error associated with a more extreme signal, and vice versa. If the state of nature were not observable and forecasters only cared about their reputation, they would

---

[13]For example, when releasing earnings forecasts, managers may benefit from the market believing that the firm has good prospects. Also, incumbent politicians may benefit from the public's belief that the economy's future is outstanding. Finally, a security analyst may generate more investment banking business for his brokerage house when he issues overly-optimistic forecasts on companies that are likely to raise capital in the future.

always report their prior so as to conceal the presence of noise. Of course, the fact that the market will observe the state of nature later and will use it as a basis to evaluate forecasters, mitigates the tendency to shade forecasts toward public information, but does not fully eliminate it.[14]

By contrast, the contest generates a tendency to exaggerate. To understand this phenomenon, consider the trade-off that forecasters face in the contest. A forecaster who naively reported the posterior expectation would maximize the chances of winning the contest prize. However, conditional on winning, the forecaster would earn a small prize because the prize would be shared among too many rival forecasters. In fact, if everyone was naive, a small deviation from naive forecasting would not only entail a second-order reduction in the chances of winning but also a first-order increase in the size of the prize (as the prize would be shared with less people).

Forecasts are (conditionally) unbiased despite the presence of distorted economic incentives for a knife-edge configuration of parameters. While the reputational concern generates a tendency to herd, the contest induces a tendency to anti-herd. These two countervailing incentives exactly offset each other at $\theta = \overline{\theta}$.

Regardless of the forecast bias, it is worth noting that an outside observer who has complete knowledge of the model would be able to compute the correct $E(x|s)$, where the signal is obtained by inverting the equilibrium reporting strategy, $s = f^{-1}(f)$.

4.2 Forecast Dispersion

Incentives also affect the dispersion of forecasts. In particular, greater competition leads to more differentiation among forecasters and results in greater dispersion in the distribution of forecasts.

Corollary 2 *The volatility of forecasts,*

$$var(f) \equiv \alpha^2 \frac{\tau + v}{\tau v}, \tag{5}$$

decreases in θ: $\frac{\partial var(f)}{\partial \theta} \leq 0.$

Proof. The result follows immediately from the observation that $\frac{\partial \alpha}{\partial \theta} \leq 0$. ∎

---

[14]Conservatism would also arise if forecasters were underconfident about their precision.

This result suggests that stronger reputational concerns lead to lower forecast dispersion, other things being equal. Based on this intuition, empirical studies sometimes interpret the lack of forecast dispersion as evidence of herding. For instance, Guedj and Bouchaud (2004) assert that low forecast dispersion coupled with high forecast error dispersion is a strong indication of herding. However, this relation is not necessarily true. In general, even in the absence of herding, one may observe great forecast error volatility and little forecast volatility if the signal precision τ is small as compared with the underlying uncertainty in the environment v^{-1}.

$$
\begin{aligned}
var\left(f-x\right) &= \frac{(\alpha-1)^2}{v} + \frac{\alpha^2}{\tau} \\
&= var\left(f\right) + \frac{1-2\alpha}{v}
\end{aligned}
$$

The idea that stronger reputational concerns are associated with lower forecast variability suggests that as the population of forecasters becomes older, and their reputational concerns weaker, one should observe greater variability in the distribution of forecasts.

4.3   Forecast Orthogonality

Another empirical implication of our analysis suggests that forecasts are generally not orthogonal to their errors, as would be the case if they were unbiased. The covariance between errors and forecasts (denoted by ρ) is negative when reputational concerns are high and positive when contest incentives are high.

Corollary 3 *The covariance between forecasts and forecast errors*

$$
\rho\left(\theta\right) \equiv cov\left(f-x,f\right) = \alpha\frac{v\alpha - \tau\left(1-\alpha\right)}{v\tau} \tag{6}
$$

is positive (negative) if $\theta \leq (>) \overline{\theta}$, as defined in Corollary 1. Furthermore, $\rho\left(\theta\right)$ is U-shaped in θ.

Proof. Since $E\left(f-x\right) = 0$, then

$$
\begin{aligned}
\rho\left(\theta\right) &= E\left(f-x\right)f = E\left(f^2\right) - E\left(fx\right) \\
&= \alpha\frac{v\alpha + \tau\alpha - \tau}{v\tau}.
\end{aligned}
$$

19

It is easy to verify that $\rho'(\hat{\theta}) = 0$ is uniquely solved by $\hat{\theta} = \frac{1}{8}\tau\frac{2\tau^2+5v\tau+4v^2}{v(v+\tau)}$. Furthermore, $\rho(\theta) = 0$ has two solutions at $\theta = \bar{\theta}$ and $\theta = \theta^+$ and $\rho(\hat{\theta}) = -\frac{1}{4v}\frac{\tau}{v+\tau}$ is a minimum. ∎

The sign of the correlation between forecasts and forecast errors has been used empirically to test whether or not security analysts herd or anti-herd. In the context of financial analysts, the evidence regarding herding (underweighing of private information) is mixed. Chen and Jiang (2005) test the orthogonality between forecast errors and forecast surprises (i.e., $E(f) - E(x)$), which should hold under truthful Bayesian forecasting. Proxying prior expectations using the consensus forecast, they find evidence of anti-herding (or overweighting of private information). Bernhardt, Campello, and Kutsoati (2006) test the prediction that, for a Bayesian analyst, the probability of a positive forecast error should equal $1/2$ regardless of whether an analyst's signal is higher than prior expectations or not. They also find evidence that analysts anti-herd. By contrast, in the context of recommendations, Jegadeesh and Kim (2010) find that recommendations that lie further away from the consensus induce stronger market reactions, consistent with herding.

5   Estimation

In the following we discuss a simple structural approach to estimating the model of Section 4 and then we implement the estimation using data from the *Business Week Investment Outlook*'s yearly *GNP* growth forecasts.

Let x_t be *GNP* in year $t \in \{1, 2, ..., T\}$ and f_{it} be the year-ahead forecast of analyst i for x_t. From Proposition 1 the equilibrium forecasting strategy is

$$f_{it} = (1 - \alpha)\mu + \alpha s_{it}$$

where α is given by (3) so that θ is equal to

$$\theta = \frac{1}{2}(\tau + v)\frac{\tau(1 - \alpha) - \alpha^2 v}{v}. \tag{7}$$

The model has three independent parameters, $\varphi \equiv (\tau, v, \alpha)$, which we can estimate using the Method of Moments (MM).

To estimate φ we need at least three moment conditions for identification. Let

$$\overline{\mathbf{g}} = \left(\widehat{var}(x_t), \widehat{var}(f_{it}), \widehat{Cov(x_t, f_{it})}\right)$$

be a consistent estimator of

$$\gamma\left(\varphi\right) = \left(var\left(x_t\right), var(f_{it}), cov(x_t, f_{it})\right).$$

Recall that

$$var\left(x_t\right) = \tau^{-1},$$

$$var\left(f_{it}\right) = \alpha^2 \left(\frac{1}{v} + \frac{1}{\tau}\right),$$

and

$$cov(x_t, f_{it}) = \frac{\alpha}{\tau}.$$

Hence the MM estimator of φ is defined as the solution to the following system of equations

$$\overline{\mathbf{g}} - \gamma\left(\hat{\varphi}\right) = 0,$$

which gives

$$\hat{v} = \widehat{var(x_t)}^{-1},$$

$$\hat{\tau} = \left(\frac{\widehat{var(f_{it})}}{[\frac{\widehat{cov(x_t, f_{it})}}{\widehat{var(x_t)}}]2} - \widehat{var(x_t)}\right)^{-1},$$

$$\hat{\alpha} = \frac{\widehat{cov(x_t, f_{it})}}{\widehat{var(x_t)}}.$$

Using (7) the following estimator of θ is easily derived,

$$\hat{\theta} = (\hat{\tau} + \hat{v})\frac{\hat{\tau}\left(4\hat{v} + \hat{\tau}\right) - \left(2\hat{v}\hat{\alpha} + \hat{\tau}\right)^2}{2\hat{v}^2}.$$

The asymptotic distribution of $\hat{\varphi}$ is obtained by applying the Delta Method and invoking the Lindeberg-Levy Central Limit Theorem (see Cameron and Trivedi, 2005 page 174):

$$\sqrt{n}\left(\hat{\varphi} - \varphi\right) \rightarrow_D N(0, G^{-1}\Sigma G'^{-1}),$$

where G is the 3×3 Jacobian matrix of the vector function $\gamma\left(\varphi\right)$, and Σ is the asymptotic variance-covariance matrix of $\overline{\mathbf{g}}$.
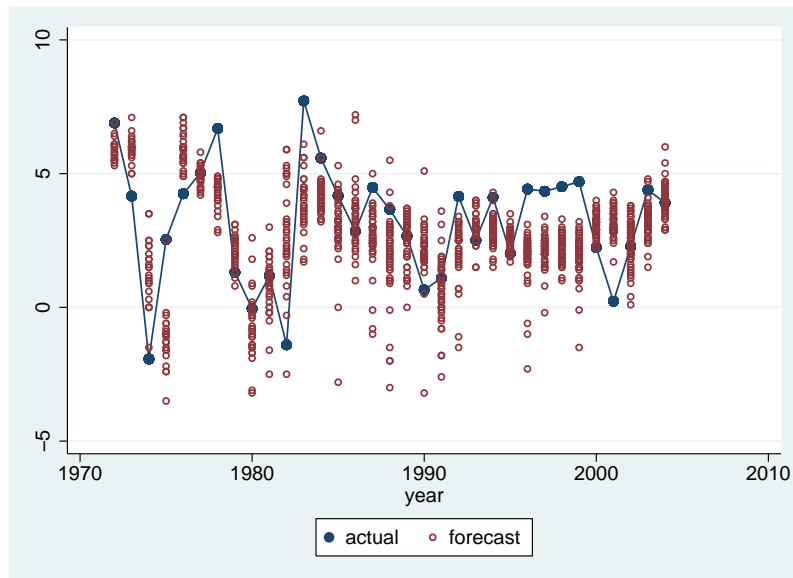
Figure 2: The red circles represent Business Week Investment Outlook's individual forecasts of annual real GNP growth rate for the period from 1972 to 2004. The connected blue balls represent the realizations, obtained from the Bureau of Economic Analysis.

5.1 Data

For illustrative purposes we use data from the *Business Week Investment Outlook*'s yearly *GNP* growth forecasts for the period from 1972 to 2004. The forecast data is obtained from a survey of professional forecasters that the magazine *Business Week* publishes at the end of each year. For the fundamentals we use the latest revisions of *GNP* growth rates released by the *Bureau of Economic Analysis*.[15] The number of forecasters in our sample rose steadily over time from around 20 in 1972 to more than 60 in 2004.

On average, *forecast* is smaller than actual, which corroborates Lahiri and Teigland (2002) finding that forecasters are too conservative in the sense that they assign excessive probability to very low rates of growth in *GNP*. We also find that *forecast* is less disperse than actual. The cross sectional standard deviation of *forecast* is, on average, 0.84 (which is more than 50% of the overall standard deviation, 1.624) confirming the idea that forecasters use private information to form their forecasts (see Figure 2).[16]

---

[15]OS also use these data to produce a graph similar to Figure 2. Part of the Business Week data have been collected by Lamont (2002) and part by OS.

[16]Unfortunately we do not know the identity of individual forecasters to be able to estimate the within-forecaster standard deviation during the sample period. However, if we randomly assign identities and then use those assigned identities to compute the within forecaster standard deviation, the average standard deviation across forecasters is 1.432.

Table 1: Summary statistics.

Variable	Mean	Std. Dev.	n
actual	3.192	2.206	33
forecast	2.697	1.624	1318
ferror (f-x)	-0.499	2.009	1318

The forecast error is negatively correlated with *actual*. The positive correlation between *forecast* and *ferror* already suggests that the forecasting strategy is not Bayesian because the orthogonality between the forecast and the forecast error is violated.
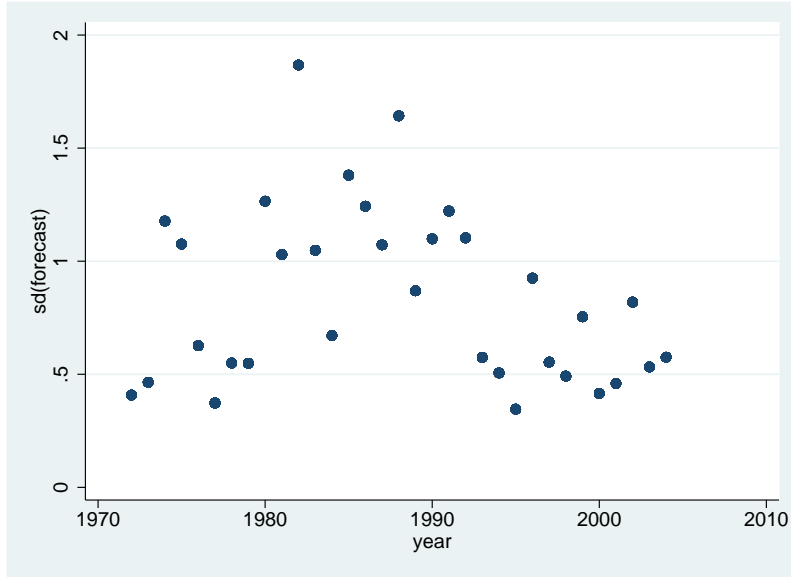

Figure 3: Standard deviation of *forecast*.

5.2   Results

Table 3 reports the MM estimator of $\hat{\varphi}$, with standard errors in parentheses, while Table 4 reports the estimated variance-covariance matrix of **g**.

Table 2: Cross-correlation table.

Variable	actual	forecast	ferror
actual	1.000		
forecast	0.401	1.000	
ferror	-0.672	0.409	1.000

Table 3: MM estimates of φ.

$\hat{v}$	$\hat{\tau}$	$\hat{\alpha}$
.20566633	.03143623	.26808195
(.01042794)	(.00004777)	(.0119226)

These estimates imply $\hat{\delta} = 0.0047204$ with standard error 2.96414×10^{-3}. The estimated variance of the noise in forecasters' signals is about 6.5 times the estimated variance of x_t, suggesting that forecasters face substantial uncertainty about fundamentals. The value of $\hat{\delta}$ reveals that forecasters are mostly driven by competitive incentives as opposed to reputation concerns. In fact $\hat{\delta}$ is 1.6 times the value of its estimated standard deviation, so it is not possible to reject the hypothesis that forecasters are exclusively driven by forecasting incentives, at a 5% significance level.[17]

Table 4: MM estimate of the var-covar matrix Σ.

Variables	g_1	g_2	g_3
g_1	33.48133		
g_2	3.5163562	23.637073	
g_3	8.4808813	8.7078661	12.373894

According to the estimates in Table 3, the weight a Bayesian forecaster would assign to his private information s is equal to

$$\hat{\alpha}^{BAYES} = \frac{\hat{\tau}}{\hat{\tau} + \hat{v}} = 0.1327.$$

The weight that forecasters use in our sample is $\hat{\alpha} \approx 0.27$, twice as high as the Bayesian weight. The difference $\hat{\alpha} - \hat{\alpha}^{BAYES}$ is statistically significant at a 1% significance level.[18] This evidence goes against the notion that forecasters *herd* and corroborates the idea that contest/competitive incentives dominate the forecasters' choices.

Notice that our *exaggeration* evidence is consistent with a consensus forecast that is much *smoother* that the actual *GNP*, as can be seen in Figure 4. The lack of variability of the consensus forecast relative to the actual *GNP* may reflect the poor quality of private information and is also driven by the fact that the consensus forecast overweighs public information by construction.

---

[17]The estimated standard deviation of $\hat{\delta}$ was computed using the Delta method.

[18]It should be noted that the assessment of significance depends on the implicit i.i.d. assumption.
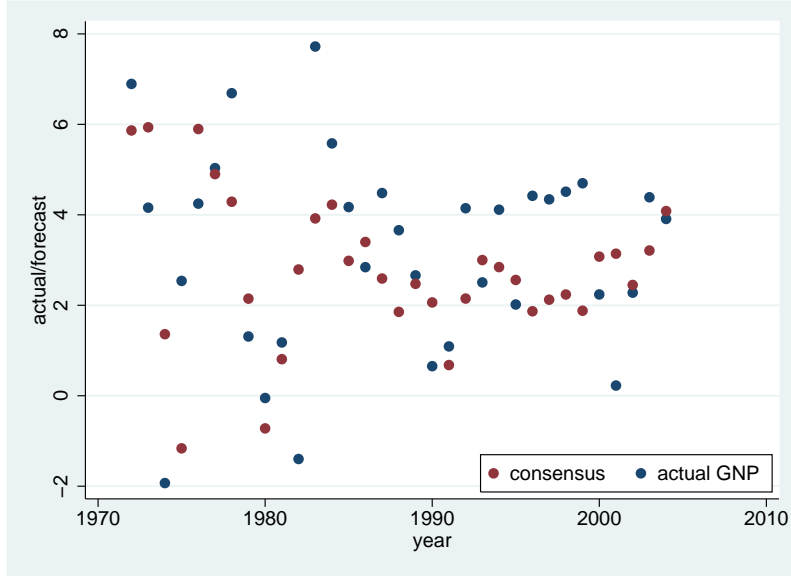
Figure 4: The consensus forecast.

To understand this point note that the unweighted average forecast $\bar{f} = \sum_{i=1}^{n} f_i/n$ (also known as the consensus forecast) is not the efficient forecast given the n signals. For example, suppose that all forecasts are truthful, $f_i = E(x|s_i)$. When the individual errors ε_i are statistically independent, the correlation of the consensus forecast with its error is negative, since the forecast tends to be too low when x exceeds μ and too high when x is below μ. The weight that the conditional expectation $E[(x|s_1, \ldots, s_n) = (\nu\mu + \tau_1 s_1 + \cdots + \tau_n s_n)/(\nu + \tau_1 + \cdots + \tau_n)$ attaches to s_i is τ_i/τ_j times the weight to s_j. In the consensus forecast, the ratio of weights is instead $\tau_i(\tau_j + \nu)/\tau_j(\tau_i + \nu)$, so that too much weight is given to the least precise signals. Even when all forecasters are equally precise, the weight accorded to the prior mean μ is too large and the consensus forecast fails to inherit the orthogonality property from the individual forecasts. In this case the consensus honest forecast is $\bar{f} = (n\nu\mu + \tau \sum_{i=1}^{n} s_i)/(n\nu + n\tau) = (n\nu\mu + n\tau x + \tau \sum_{i=1}^{n} \varepsilon_i)/(n\nu + n\tau)$ and the error is $\bar{f} - x = (n\nu(\mu - x) + \tau \sum_{i=1}^{n} \varepsilon_i)/(n\nu + n\tau)$, so that the covariance is always negative: $E[\bar{f}(\bar{f}-x)] = -(n-1)\tau/n(\nu + \tau)^2 < 0$ for $n > 1$. See also Kim, Lim, and Shaw (2001) and Crowe (2010) for methods to adjust this overweighing of the prior when averaging forecasts. As shown by Marinovic, Ottaviani, and Sørensen (2011), when forecasters underweigh sufficiently their private information relative to common public information (for example because of beauty contests concerns à la Morris and Shin 2002), an increase in the number of forecasters can actually lead to a reduction of the informativeness of the

consensus forecast.

5.3   Limitations

Our approach assumes that the data is *identically and independently distributed*, perhaps not the most realistic assumption for the *GNP* series. Moreover, we omit dynamic considerations, and implicitly assume that the forecasting strategy is stationary. In reality, the forecasting strategy is likely to be modified by the *history*, as an individual forecaster's reputation evolves.[19] These limitations however don't change the qualitative evidence that, in our sample, forecasters overweigh private information (relative to the Bayesian forecast), in line with our contest theory.

Also, we assume that forecasters hold common priors thus omitting the possibility that forecasters hold heterogenous priors; see Section 6.3 for further discussion. Finally, we have ruled out behavioral explanations, such as overconfidence, as the drivers of our forecasting data, but our evidence is certainly consistent with overconfidence about the precision of private information; see for example Gervais and Odean (2001) or Griffin and Tversky (1992).

6   Robustness and Extensions

We now turn to a discussion of the robustness of our results in light of a number of natural extensions. Section 6.1 explores the possibility that the information all forecasters observe contains a common error component. Section 6.2 considers the case where forecasters are informed about their own ability. Section 6.3 discusses the effect of relaxing the assumption that the prior about the state is common to all participants to allow for heterogeneous priors. Section 6.4 extends the model to deal with partisan forecasters, who benefit from biased forecasts.

---

[19]Clement and Tse (1995) find that (1) boldness likelihood increases with the analyst's prior accuracy, brokerage size, and experience and declines with the number of industries the analyst follows, consistent with theory linking boldness with career concerns and ability; (2) bold forecasts are more accurate than herding forecasts; and (3) herding forecast revisions are more strongly associated with analysts' earnings forecast errors (actual earnings – forecast) than are bold forecast revisions.

6.1 Common Error

Suppose that a shock ε_0 occurs to the variable of interest between the time when forecasts are published, t, and $t + 1$, so that the realization is $y = x + \varepsilon_0$ rather than x as in our baseline. For each agent i, the observed signal can then be written as $s_i = x + \varepsilon_i = y + \varepsilon_i - \varepsilon_0$, and the forecast errors $\varepsilon_1 - \varepsilon_0$, $\varepsilon_2 - \varepsilon_0$, ..., $\varepsilon_n - \varepsilon_0$ are positively correlated because of their common component. As suggested by Keane and Runkle (1998), this mechanism could explain the significant positive correlation among the residuals in the orthogonality regression. OS (Proposition 8) argue that the introduction of a common error does not qualitatively alter the results of both the reputational cheap talk and the forecasting contest models, if the noise in the realization y is sufficiently small. Partly correcting this argument, Lichtendahl, Grushka-Cockayne, and Pfeifer (2012) obtain a closed-form solution for the linear equilibrium resulting when the forecasters observe conditionally correlated signals.

6.2 Information about Own Precision

In our reputational model the forecaster is not privately informed about the signal's true underlying precision, t. The advantage of this formulation is that the forecaster's information is then one-dimensional. What would happen if, instead, each forecaster privately observed not only the signal s but also the signal's precision t?

Private knowledge of the precision induces the forecaster to issue a forecast further away from the mean of the prior distribution of x. Intuitively, more extreme predictions indicate more precise information, as noticed by Trueman (1994) and further analyzed by Ottaviani and Sørensen (2006b). A force toward exaggeration then also arises in the reputational model.

The comparison of the magnitude of the conservatism and exaggeration effects depends on the precision of the prior distribution relative to the signal precision. This new effect is almost absent when the distribution of the state x is highly dispersed (ν is small), in which case a bias toward the middle results. When the state x is highly uncertain, there is little to gain from being far from the prior mean μ. When ν is large, instead, the signaling effect is important and a bias away from the prior mean results. Once this second effect is taken into account, our findings are reconciled with the results Zitzewitz (2001) obtains in a model that assumes managers to know their precision.

The incentive to assign excessive weight to the private signal is already present in Prendergast and Stole's (1996) reputational signaling model. In their model, however, the evaluation of the market is based exclusively on the action taken by a manager who knows the precision of her private information without conditioning on the ex post state realization. This tendency to exaggerate seems robust to the introduction of a small amount of ex post information on the state, in the same way as our conservatism finding is robust to small amounts of private information on the precision. Empirically, Ehrbeck and Waldmann (1996) find that changes in forecasts are positively correlated with forecast errors and that forecasters who make larger changes in forecasts have larger forecast errors. These findings cast doubt on the reputational explanation for forecast bias.

Lamont (2002) finds that older and more established forecasters tend to issue more extreme forecasts which turn out to be less accurate. As also suggested by Avery and Chevalier (1999), younger managers should have a tendency to be conservative, having little private information about their own ability; older managers instead exaggerate, being more confident about their ability. Notice the contrast with Prendergast and Stole's (1996) prediction of impetuous youngsters and jaded old-timers for cases in which the same manager privately informed about own ability makes repeated observable decisions with a *constant* state.

6.3   Heterogeneous Priors

In our model forecasters issue different forecasts on the basis of their different private information. It is natural to wonder where this information originates. After all, in most applications forecasters are exposed to similar public information. Nevertheless, it is reasonable that forecasters have different models to interpret this public information. As such, the different private signals can be seen as arising from the private information forecasters possess about their own models. More realistically, differences in the elaboration of information across forecasters can easily result by assuming heterogeneity of the priors about the state x and heterogeneous interpretation of information. Along these lines, Kandel and Zilberfarb (1999) suggest that a reason for the variability of forecasts is that forecasters have different opinions about the state x. We conjecture that an extension of our model in which forecasters have heterogeneous prior beliefs about the state would result in additional heterogeneity of forecasts. Thus, our estimate of exaggeration could

be upwardly biased.

Assuming that forecasters are Bayesian but not strategic, Lahiri and Sheng (2008) use data on the evolution of forecaster disagreement over time to analyze the relevance of three components: i) the initial heterogeneous priors; ii) the weight attached to the priors; iii) the heterogeneous interpretation of public information; they find that the first and the third components are important. In this vein, Patton and Timmermann (2010) offer an explanation for the excessive cross sectional dispersion of macroeconomic forecasts in terms of heterogeneous beliefs. Assuming that forecasters are non strategic, they find some evidence that dispersion among forecasters is highest at long horizons, where private information is of limited value, and lower at shorter forecast horizons. Given the mounting evidence that forecast dispersion is explained by differences in priors—or in models used by forecasters—it is a research priority to extend our framework in this direction.

6.4 Partisan Forecasting

Our baseline model predicts that equilibrium forecasts are on average unbiased. Neither reputational concerns nor contest prizes lead to systematic biases. In reality, forecasters in some contexts appear to be biased. Bias is particularly well documented in the literature on security analysts. The analysts' conflicts of interest has been attributed to a number of factors, such as the incentives to generate investment-banking business (see Michaely and Womack 1999), the desire to increase the brokerage commissions for the trading arms of the employing financial firms (Jackson 2005), and the need to gain access to internal information from the firms they cover (Lim 2001). This bias is consistent with the evidence that analysts who have been historically more optimistic are more successful in their careers, after controlling for accuracy (Hong and Kubik 2003).

In this section we extend the model to reward forecasters who display some bias by modifying the reputational component of the forecaster's utility as follows

$$\tilde{u}_r = \exp\left\{-E\left[\hat{E}\left(x + \beta - s | f, x\right)^2 | s\right]\right\}.$$

When $\beta > 0$, the market rewards forecasters who are overly optimistic, and vice versa.[20]

Proposition 3 *If $\theta < \theta^+$, there exists a unique linear equilibrium in which the forecaster's*

---

[20]This modeling of the bias is similar to Crawford and Sobel (1982), Morgan and Stocken (2004), and many others. For an alternative approach to modeling bias see, for example, Beyer and Guttman (2007).

strategy is given by

$$f = \frac{\beta \theta \tau}{(1 - \theta)(\tau + v)} + \alpha(\theta)\mu + (1 - \alpha(\theta))s.$$

The forecast is now biased and the magnitude of the bias increases in the size of the incentives β and the precision of the forecaster but decreases in the precision of public information. Note the complementarity between incentives and precision—the more precise the forecaster, the greater the impact of incentives in forecasters' bias.[21] Similarly, incentives and reputation are complements.[22]

7 Role of Anonymity

In order to test the different theories, it might be useful to compare non-anonymous with anonymous forecasting surveys. The *Survey of Professional Forecasters* of the Federal Reserve Bank of Philadelphia is arguably the most prominent anonymous survey of professional forecasters; see Stark (1997) for an analysis. Even though the name of the author of each forecast is not made public, each forecaster is identified by a code number. It is then possible to follow each individual forecaster over time. As reported by Croushore (1993):

> "This anonymity is designed to encourage people to provide their best forecasts, without fearing the consequences of making forecast errors. In this way, an economist can feel comfortable in forecasting what she really believes will happen . . . Also, the participants are more likely to take an extreme position that they believe in (for example, that the GDP will grow 5 per cent in 1994), without feeling pressure to conform to the consensus forecast. The negative side of providing anonymity, of course, is that forecasters can't claim credit for particularly good forecast performance, nor can they be held accountable for particularly bad forecasts. Some economists feel that without accountability,

---

[21]Fang and Yasuda (2009) find that personal reputation acts as discipline against conflicts of interest. Thus, their results suggest that bias and information quality are substitutes rather than complements as in our model.

[22]Exploiting the natural experiment provided by mergers of brokerage houses, Hong and Kacperczyk (2010) find that bias tends to decrease in the level of competition among analysts, consistent with our prediction.

forecasters may make less accurate predictions because there are fewer consequences to making poor forecasts."

When reporting to anonymous surveys, forecasters have no reason not to incorporate all available private information. Forecasters are typically kept among the survey panelists if their long-term accuracy is satisfactory. By effectively sheltering the forecasters from the short-term evaluation of the market, anonymity could reduce the scope for strategic behavior and induce honest forecasting. Under the assumption that forecasters report honestly in the anonymous surveys, one could test for the presence of strategic behavior in the forecasts publicly released in non-anonymous surveys.

A problem with the hypothesis of honest forecasting in anonymous surveys is that our theory does not predict behavior in this situation. According to industry experts, forecasters often seem to submit to the anonymous surveys the same forecasts they have already prepared for public (i.e. non-anonymous) release. There are two reasons for this. First, it might not be convenient for the forecasters to change their report, unless they have a strict incentive to do so. Second, the forecasters might be concerned that their strategic behavior could be uncovered by the editor of the anonymous survey.

We have computed the dispersion of the forecasts in the anonymous Survey of Professional Forecasters and found it even higher than in the non-anonymous Business Economic Outlook. This high dispersion suggests that more exaggeration might be present in the anonymous survey. This possibility needs more careful investigation. The composition of the forecasters' panel of the Survey of Professional Forecasters is now available to researchers, so it is possible to verify whether anonymous and non-anonymous releases of individual forecasters can be easily matched. Also, the joint hypothesis of honest reporting in anonymous surveys and strategic forecasting in non-anonymous surveys could be tested by pooling in a single regression all the forecasters belonging to both data sets.

8 Summary and Outlook

This chapter provides a strategic foundation for the forecasters' objectives through a framework that integrates the reputational theory with the contest theory of strategic forecasting. In general, other than in knife-edge situations, truthful reporting of conditional expectations is not an equilibrium when forecasters (a) possess both private and

public information, and (b) care about their reputation but also compete for the best accuracy record. While reputation induces forecasters to partially disregard their private information resulting in excessive agreement among forecasters, competition leads to the opposite—exaggeration of private information and excessive disagreement. Yet, the presence of distorted economic incentives is not sufficient to produce untruthful reporting. In fact, reputational concerns and competition form countervailing incentives which, under a knife edge condition, can perfectly offset one other inducing forecasters to truthfully report their conditional expectations.

Somewhat paradoxically, when reputational concerns are overwhelming, the informational content of forecasts deteriorates and only categorical information may be supplied. By contrast, the presence of strong competition among forecasters results in highly differentiated forecasts. Interestingly, a commitment to evaluate forecasters using only their relative performance (as opposed to their absolute performance) may generate more informative forecasts.

In spite of the progress in the area reviewed here, applied research on strategic forecasting is still in its infancy. We look forward to future work that finesses the approach and improves our interpretation of professional forecasts. A natural next step is to allow for ex-ante heterogeneity across forecasters.[23] For empirical work it would be particularly important to extend our model to allow for richer dynamics. A key challenge lies in finding a tractable and sufficiently general multi-period environment with learning about the precision as well as about the state.[24] It could also be useful to bridge the widening gap between the applied literature on strategic forecasting we reviewed here and the theoretical literature on expert testing we briefly discussed in the introduction.

---

[23]A number of empirical facts are emerging regarding heterogeneous behavior of forecasters with different characteristics. For example, Loh and Stulz (2010) note that recommendation changes by certain higher-status analysts tend to influence more stock prices. For another example, Evgeniou et al. (2010) find that low-skilled analysts provide significantly bolder forecasts as the environment becomes more uncertain.

[24]See Clarke and Subramanian (2006) for an interesting effort in this direction. They find that the forecasters that tend to be bolder are both the historical underperformers and the outperformers, relative to forecasters with middling performance.

References

Al-Najjar, Nabil I., Alvaro Sandroni, Rann Smorodinsky, and Jonathan Weinstein, "Testing Theories with Learnable and Predictive Representations," *Journal of Economic Theory*, 2010, *145* (6), 2203–2217.

Avery, Christopher N. and Judith A. Chevalier, "Herding Over the Career," *Economics Letters*, June 1999, *63* (3), 327–333.

Banerjee, Abhijit V., "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

Bayarri, M. J. and M. H. DeGroot, "Optimal Reporting of Predictions," *Journal of the American Statistical Association*, 1989, *84* (405), 214–222.

Bernhardt, Dan, Murillo Campello, and Edward Kutsoati, "Who Herds?," *Journal of Financial Economics*, 2006, *80* (3), 657–675.

Beyer, Anne and Ilan Guttman, "The Effect of Trading Volume on Analysts' Forecast Bias," *Accounting Review*, 2011, *86* (2), 451–481.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 1992, *100* (5), 992–1026.

Cameron, Colin A. and Pravin K. Trivedi, *Microeconometrics: Methods and Applications*, Cambridge University Press, May 2005.

Chen, Qi and Wei Jiang, "Analysts' Weighting of Private and Public Information," *Review of Financial Studies*, 2006, *19* (1), 319–355.

Chevalier, Judith and Glenn Ellison, "Career Concerns of Mutual Fund Managers*," *Quarterly Journal of Economics*, 1999, *114* (2), 389–432.

Clarke, Jonathan and Ajay Subramanian, "Dynamic Forecasting Behavior by Analysts: Theory and evidence," *Journal of Financial Economics*, 2006, *80* (1), 81–113.

Clement, Michael B. and Senyo Y. Tse, "Financial Analyst Characteristics and Herding Behavior in Forecasting," *Journal of Finance*, 2005, *60* (1), 307–341.

Cowles, Alfred, "Can Stock Market Forecasters Forecast?," *Econometrica*, 1933, *1* (3), 309–324.

Crawford, Vincent P. and Joel Sobel, "Strategic Information Transmission," *Econometrica*, 1982, *50* (6), 1431–1451.

Croushore, Dean, "Introducing the Survey of Professional Forecasters," *Federal Reserve Bank of Philadelphia Business Review*, 1993, (Nov), 3–15.

Dawid, A. Philip, "The Well-Calibrated Bayesian," *Journal of the American Statistical Association*, 1982, *77* (379), 605–610.

Denton, Frank Trevor, "The Effect of Professional Advice on the Stability of a Speculative Market," *Journal of Political Economy*, 1985, *93* (5), 977–993.

Ehrbeck, Tilman and Robert Waldmann, "Why Are Professional Forecasters Biased? Agency versus Behavioral Explanations," *Quarterly Journal of Economics*, 1996, *111* (1), 21–40.

Elliott, Graham, Ivana Komunjer, and Allan Timmermann, "Estimation and Testing of Forecast Rationality under Flexible Loss," *The Review of Economic Studies*, 2005, *72* (4), 1107–1125.

Evgeniou, Theodoros, Lily H. Fang, Robin M. Hogarth, and Natalia Karelaia, "Uncertainty, Skill and Analysts' Dynamic Forecasting Behavior," *INSEAD Working Paper No. 2010/50/DS/FIN*, 2010.

Fang, Lily and Ayako Yasuda, "The Effectiveness of Reputation as a Disciplinary Mechanism in Sell-Side Research," *Review of Financial Studies*, 2009, *22* (9), 3735–3777.

Fortnow, Lance and Rakesh V. Vohra, "The Complexity of Forecast Testing," *Econometrica*, 2009, *77* (1), 93–105.

Foster, Dean P. and Rakesh V. Vohra, "Asymptotic Calibration," *Biometrika*, 1998, *85* (2), 379–390.

Galton, Francis, "Vox Populi," *Nature*, 1907, *75*, 450–451.

Gervais, Simon and Terrance Odean, "Learning to be Overconfident," *Review of Financial Studies*, 2001, *14* (1), 1–27.

Graham, John R., "Herding Among Investment Newsletters: Theory and Evidence," *Journal of Finance*, 1999, *54* (1), 237–268.

Granger, Clive W. J., "Prediction with a Generalized Cost of Error Function," *Operations Research*, 1969, *20* (2), 199–207.

Griffin, Dale and Amos Tversky, "The Weighing of Evidence and the Determinants of Confidence," *Cognitive Psychology*, 1992, *24* (3), 411–435.

Guedj, Olivier and Jean-Philippe Bouchaud, "Experts' Earning Forecasats: Bias, Herding and Gossamer Information," *International Journal of Theoretical and Applied Finance*, 2005, *08* (07), 933–946.

Hirshleifer, David and Siew Hong Teoh, "Herd Behaviour and Cascading in Capital Markets: a Review and Synthesis," *European Financial Management*, 2003, *9* (1), 25–66.

Holmström, Bengt, "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 1999, *66* (1), 169–182.

Hong, Harrison and Jeffrey D. Kubik, "Analyzing the Analysts: Career Concerns and Biased Earnings Forecasts," *Journal of Finance*, 2003, *58* (1), 313–351.

___ **and Marcin Kacperczyk**, "Competition and Bias," *Quarterly Journal of Economics*, 2010, *125* (4), 1683–1725.

___ **, Jeffrey D. Kubik, and Amit Solomon**, "Security Analysts' Career Concerns and Herding of Earnings Forecasts," *RAND Journal of Economics*, 2000, *31* (1), 121–144.

Hotelling, Harold, "Stability in Competition," *Economic Journal*, 1929, *39* (153), pp. 41–57.

Hwang, Jiunn Tzon and James V. Zidek, "Limit Theorems for Out-Guesses with Mean-Guided Second Guessing," *Journal of Applied Probability*, 1982, *19* (2), 321–331.

Jackson, Andrew R., "Trade Generation, Reputation, and Sell-Side Analysts," *Journal of Finance*, 2005, *60* (2), 673–717.

Jegadeesh, Narasimhan and Woojin Kim, "Do Analysts Herd? An Analysis of Recommendations and Market Reactions," *Review of Financial Studies*, 2010, *23* (2), 901–937.

Kandel, Eugene and Ben-Zion Zilberfarb, "Differential Interpretation of Information in Inflation Forecasts," *Review of Economics and Statistics*, 1999, *81* (2), 217–226.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani, "Credulity, Lies, and Costly Talk," *Journal of Economic Theory*, 2007, *134* (1), 93–116.

Keane, Michael P. and David E. Runkle, "Are Financial Analysts' Forecasts of Corporate Profits Rational?," *Journal of Political Economy*, 1998, *106* (4), 768–805.

Kim, Oliver, Steve C. Lim, and Kenneth W. Shaw, "The Inefficiency of the Mean Analyst Forecast as a Summary Forecast of Earnings," *Journal of Accounting Research*, 2001, *39* (2), 329–335.

Krishnan, Murugappa, Steve C. Lim, and Zhou Ping, "Analysts' Herding Propensity: Theory and Evidence from Earnings Forecasts," *Working paper, SSRN*, 2011.

Lahiri, Kajal and Christie Teigland, "On the Normality of Probability Distributions of Inflation and GNP Forecasts," *International Journal of Forecasting*, 1987, *3* (2), 269–279.

__ **and Xuguang Sheng**, "Evolution of Forecast Disagreement in a Bayesian Learning Model," *Journal of Econometrics*, 2008, *144* (2), 325–340.

Lamont, Owen A, "Macroeconomic Forecasts and Microeconomic Forecasters," *Journal of Economic Behavior and Organization*, 2002, *48* (3), 265–280.

Laster, David, Paul Bennett, and In Sun Geoum, "Rational Bias in Macroeconomic Forecasts," *Quarterly Journal of Economics*, 1999, *114* (1), 293–318.

Lichtendahl, Casey, Yael Grushka-Cockayne, and Phil E. Pfeifer, "The Wisdom of Competitive Crowds," *Darden Business School Working Paper No. 1926330.*, 2011.

Lim, Terence, "Rationality and Analysts' Forecast Bias," *Journal of Finance*, 2001, *56* (1), 369–385.

Loh, Roger K. and René M. Stulz, "When Are Analyst Recommendation Changes Influential?," *Review of Financial Studies*, 2010, pp. 593–627.

Marinovic, Iván, Marco Ottaviani, and Peter Norman Sørensen, "Modeling Idea Markets: Between Beauty Contests and Prediction Markets," *Chapter 2 in Prediction Markets: Theory and Applications, edited by Leighton Vaughan Williams, Routledge*, 2011.

Michaely, Rony and Kent L. Womack, "Conflict of Interest and the Credibility of Underwriter Analyst Recommendations," *Review of Financial Studies*, 1999, *12* (4), 653–686.

Morgan, John and Phillip C. Stocken, "An Analysis of Stock Recommendations," *RAND Journal of Economics*, 2003, *34* (1), 183–2003.

Morris, Stephen and Hyun Song Shin, "Social Value of Public Information," *American Economic Review*, 2002, *92* (5), 1521–1534.

Olsen, Robert A., "Implications of Herding Behavior for Earnings Estimation, Risk Assessment, and Stock Returns," *Financial Analysts Journal*, 1996, *52* (4), 37–41.

Olszewski, Wojciech, "Calibration and Expert Testing," *Chapter prepared for the Handbook of Game Theory, Volume IV*, 2012.

Osborne, Martin J. and Carolyn Pitchik, "The Nature of Equilibrium in a Location Model," *International Economic Review*, 1986, *27* (1), 223–237.

Ottaviani, Marco and Peter Norman Sørensen, "Herd Behavior and Investment: Comment," *American Economic Review*, 2000, *90* (3), 695–704.

_ **and** _ , "Forecasting and Rank-Order Contests," 2005.

_ **and** _ , "Professional Advice," *Journal of Economic Theory*, 2006a, *126* (1), 120–142.

_ **and** _ , "Reputational Cheap Talk," *RAND Journal of Economics*, 2006b, *37* (1), 155–175.

_ **and** _ , "The Strategy of Professional Forecasting," *Journal of Financial Economics*, 2006c, *81* (2), 441–466.

Patton, Andrew J. and Allan Timmermann, "Why do Forecasters Disagree? Lessons from the Term Structure of Cross-sectional Dispersion," *Journal of Monetary Economics*, 2010, *57* (7), 803–820.

Pittenger, Donald B., "Some Problems in Forecasting Population for Government Planning Purposes," *American Statistician*, 1980, *34* (3), 135–139.

Prendergast, Canice and Lars Stole, "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, 1996, *104* (6), 1105–1134.

Sandroni, Alvaro, "The Reproducible Properties of Correct Forecasts," *International Journal of Game Theory*, 2003, *32*, 151–159.

Scharfstein, David S. and Jeremy C. Stein, "Herd Behavior and Investment," *American Economic Review*, 1990, *80* (3), 465–479.

__ **and** __ , "Herd Behavior and Investment: Reply," *American Economic Review*, June 2000, *90* (3), 705–706.

Slotsky, Alice Louise, *The Bourse of Babylon: Market Quotations in the Astronomical Diaries of Babylonia*, CDL Press (Bethesda, Md.), 1997.

Stark, Tom, "Macroeconomic Forecasts and Microeconomic Forecasters in the Survey of Professional Forecasters," Working Papers 97-10, Federal Reserve Bank of Philadelphia 1997.

Steele, J. Michael and James Zidek, "Optimal Strategies for Second Guessers," *Journal of the American Statistical Association*, 1980, *75* (371), 596–601.

Stickel, Scott E., "Predicting Individual Analyst Earnings Forecasts," *Journal of Accounting Research*, 1990, *28* (2), 409–417.

Trueman, Brett, "Analyst Forecasts and Herding Behavior," *Review of Financial Studies*, 1994, *7* (1), 97–124.

Zellner, Arnold, "Biased Predictors, Rationality and the Evaluation of Forecasts," *Economics Letters*, 1986, *21* (1), 45–48.

Zitzewitz, Eric, "Measuring Herding and Exaggeration by Equity Analysts and Other Opinion Sellers," Research Papers 1802, Stanford University, Graduate School of Business July 2001.