UNIVERSITY OF COPENHAGEN

The Strategy of Professional Forecasting

Sørensen, Peter Norman; Ottaviani, Marco

Publication date:
2001

Document version
Early version, also known as pre-print

Citation for published version (APA):
Sørensen, P. N., & Ottaviani, M. (2001). *The Strategy of Professional Forecasting*. Department of Economics, University of Copenhagen.

Download date: 07. apr.. 2020

DISCUSSION PAPERS
Department of Economics
University of Copenhagen

01-09

The Strategy of Professional Forecasting

Marco Ottaviani
Peter Norman Sørensen

Studiestræde 6, DK-1455 Copenhagen K., Denmark
Tel. +45 35 32 30 82 - Fax +45 35 32 30 00
http://www.econ.ku.dk

The Strategy of Professional Forecasting*

Marco Ottaviani† **Peter Norman Sørensen**‡

October 2001

Abstract

This paper develops and compares two theories of strategic behavior of professional forecasters. The first theory posits that forecasters compete in a *forecasting contest* with pre-specified rules. In equilibrium of a winner-take-all contest, forecasts are excessively differentiated. According to the alternative *reputational cheap talk* theory, forecasters aim at convincing 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 has naive views on forecasters' behavior, forecasts are biased toward the prior mean. Otherwise, equilibrium forecasts are unbiased but imprecise.

Keywords: Forecasting, contest, reputation, cheap talk.

JEL Classification: D82 (Asymmetric and Private Information), D83 (Search, Learning, and Information), G20 (Financial Institutions and Services), J30 (Wages, Compensation, and Labor Costs).

*We thank Dean Croushore for essential feedback at a number of crucial stages of this project, Owen Lamont for generously giving us the *Business Week Investment Outlook* data along with very helpful suggestions, Marco Battaglini for valuable discussion of the paper, Andrew Chesher, Christian Dustmann, Daniel Seidmann, Hyun Shin, and Eric Zitzewitz for important input, and seminar audiences at AEA 2001 Meetings (New Orleans), Bonn, Brussels, Cambridge, Copenhagen, Danish Economic Society 2000, ENTER 2001 (Mannheim), ESRC-PIPPS (Bristol), ESSET 2000 (Gerzensee), EUI (Florence), Helsinki, LBS, LSE, Munich, Pompeu Fabra, Princeton, Royal Holloway, Southampton, Surrey, and UCL for comments.

†Marco Ottaviani, Department of Economics and ELSE, University College London, and Economics Subject Area, London Business School, Sussex Place, Regent's Park, London NW1 4SA, England. Phone: +44-20-7262-5050 ext. 3778. Fax: +44-20-7402-0718. E-mail: `mottaviani@london.edu`. Web: http://www.london.edu/mottaviani.

‡Institute of Economics, University of Copenhagen, Studiestræde 6, DK–1455 Copenhagen K, Denmark. Phone: +45–3532–3056. Fax: +45–3532–3000. E-mail: `peter.sorensen@econ.ku.dk`. Web: http://www.econ.ku.dk/sorensen.

"Much has been written about the doubtful accuracy of economists' predictions. ... they are better at predicting the direction than the actual magnitude of events. ... This is disappointing, but it does not mean that economics is not a science." ('Economics', Encyclopædia Britannica Online).

1. Introduction

Even if economists have a bad name as forecasters, those who are successful are rewarded by markets and governments alike. Notably, Alan Greenspan and Laurence Meyer practiced economic forecasting before becoming the current Chairman of the Board of Governors of the Federal Reserve Bank and one of its other four current members respectively. As first argued by Lamont (1995), the microeconomic incentives of macroeconomic forecasters appear empirically relevant in explaining their behavior.

At the end of every year, Business Week asks a number of economists to predict output, inflation, and unemployment for the subsequent year. Figure 1 plots the yearly GNP growth forecasts for the period 1972-1993 from the *Business Week Investment Outlook*, as collected by Lamont (1995). Notice that the individual forecasts are rather dispersed, especially in years such as 1974 (in the aftermath of the oil shock), 1982-3, and 1991. Yet, average forecasts are less volatile than the realizations. This suggests the possibility that forecasters strategically herd. In particular, forecasters often fall short of the mark in years with extreme growth rates. The data also reveals that forecasts are more dispersed in those years in which the consensus forecast (a misnomer used for the average of the individual forecasts) is more inaccurate. Regressing the standard deviation of the forecasts on the absolute error of the consensus forecast, we find a coefficient of .145 with standard error .063. Thus, there is a significant negative relationship between accuracy and dispersion.

The existence of anonymous surveys (starting in 1946 with the Livingston Survey) presupposes a belief that forecasters might not report their best predictions if anonymity were not preserved. While the identity of the forecasters belonging to the panel is typically available, anonymous surveys do not reveal which individual made which forecasts. A possible rationale for preserving anonymity is that it can guarantee honest forecasting. For example, Croushore (1997) argues that in non-anonymous surveys "some 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." When reporting instead to an anonymous survey, forecasters would have no clear reason to bias or shade their private information.[1]

In the absence of private information, honest forecasting would be at odds with observed heterogeneity of forecasts. Forecasters who share a common prior and are given the same (public) information but no private information should make identical forecasts if they

[1]By comparing anonymous with non-anonymous surveys, one can evaluate the usefulness of anonymity and propose methods to improve these surveys (see Section 6).

honestly report their expectations.[2] In their classic study on the rationality of forecasts using data from the NBER-ASA survey of professional forecasters (later to be called the Survey of Professional Forecasters), Keane and Runkle (1990) found that differences in individual forecasts cannot be explained by publicly available information.[3] They inferred that differences in forecasts are due to asymmetric information, but this conclusion rests on their maintained assumption of honest forecasting. The observed dispersion might also indicate strategic behavior. In order not to introduce a bias against honest forecasting, we posit that forecasters are endowed with some private information about the state. After all, the presence of heterogeneous private information is the very reason for the market to reward forecasters according to their (absolute or relative) accuracy.

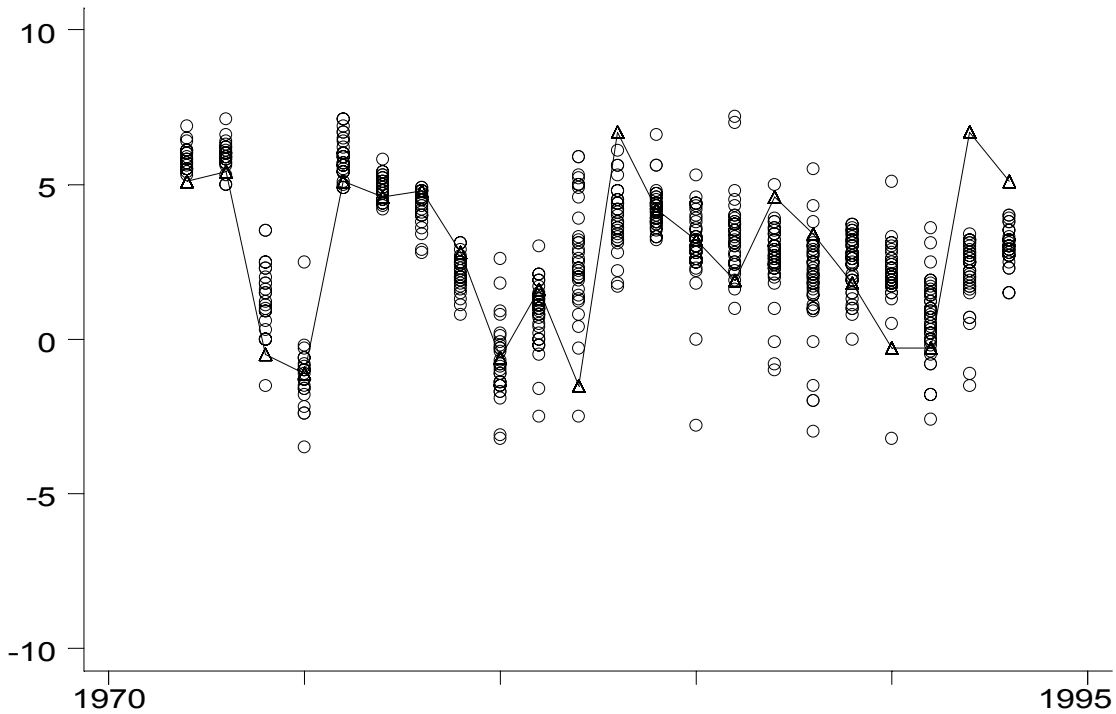

Figure 1. Comparison of individual forecasts of GNP growth from the Business Week survey with the realizations for the period 1972-1993. Horizontal axis: year forecasted for. Vertical axis: Annual real GNP growth in percentage points. Forecasts are represented by circles, ex-post realized values by connected triangles. Data source: Lamont (1995).

[2]Kandel and Pearson (1995) and Kandel and Zilberfarb (1999) find empirical support for heterogeneous processing of public information. Rather than having different private information, forecasters might have different models to process the same publicly available information. It is possible to interpret private information of the posterior belief as deriving from private information of the model (ranked by accuracy) used to process public information.

[3]Cf. Zarnowitz (1967) for an early evaluation of forecasts and Zarnowitz and Braun (1993) for a retrospective discussion of the NBER-ASA survey.

This paper proposes two theories of strategic behavior by professional forecasters and compares these with the benchmark case of non-strategic forecasting. In order to facilitate the comparison, these theories are developed in the context of a unified statistical model. Our model can be applied to situations where the state cannot be affected by policy and yet can be meaningfully forecast and later observed.[4] For reasons of tractability we use the normal location experiment: the state has a normal prior distribution and the signals of the forecasters are normally distributed around the true state. After the forecasters simultaneously release their forecasts (e.g. the predicted GNP growth rate), the state (e.g. the realized rate) is publicly observed.

Consider the benchmark case of a forecaster rewarded according to the absolute accuracy of the prediction. A non-strategic forecaster reports honestly the posterior expectation on the state, a weighted average of the signal with the prior mean. As forcefully argued by Keane and Runkle (1990), "... professional forecasters ... have an economic incentive to be accurate. Because these professionals report to the survey the same forecasts that they sell on to the market, their survey responses provide a reasonably accurate measure of their expectations." In their later paper, Keane and Runkle (1998) write: "... since financial analysts' livelihoods depend on the accuracy of their forecasts ... we can safely argue that these numbers accurately measure the analysts' expectations." Forecasters are presumed honest, unless proven strategic.

The benchmark model can already offer some explanations for the empirical observations. Forecasts are dispersed due to private information. Forecasts will be more dispersed and less accurate in years with relatively little public information, in accordance with the result of our regression.[5] More subtly, although honest forecast errors are uncorrelated with (i.e., unpredictable from) honest forecasts, they are in fact negatively correlated with the realized state. Ex post, when a high growth rate has been observed, there is a tendency for the honest forecasts to be too low. The popular press often takes this empirically documented negative correlation as evidence of herd behavior.[6] The academic empirical

[4]In a number of macroeconomic situations, forecasts can influence policy makers who act to control the very variable to be predicted (e.g. the growth or inflation rates). However, in the presence of many forecasters the individual forecaster may perceive his own forecast as having negligible influence on the state variable. For simplicity, in this paper we assume that forecasts have no impact on the distribution of the state variable. The more general case is left for future research.

[5]This observation is consistent with a finding reported by Zarnowitz and Lambros (1987) on the ASA-NBER survey (later known as the Survey of Professional Forecasters). In addition to point forecasts, this survey initially asked forecasters to report probability distributions. Zarnowitz and Lambros documented that forecast dispersion is positively correlated with a measure of forecast uncertainty constructed with the subjectively reported probability forecasts. When the perceived prior uncertainty is higher, the posterior beliefs of the forecasters are less tight and the resulting forecasts more variable across forecasters.

[6]In its widely circulated *World Economic Outlook*, even the International Monetary Fund (2001) laments that forecasters typically fall short of predicting changes in macroeconomic performance (cf. pages 6-8). In particular, forecasts are often too optimistic when the economy enters a recession, as happened recently. This is interpreted as a possible indication that economic forecasters use their information inefficiently.

literature does not fall into this trap, and instead focuses on a possible correlation between forecasts and their errors: if the direction of the error can be predicted immediately after the forecast has been issued, why didn't the forecaster adjust the forecast accordingly? Our two theories offer some possible rationales.

Our first theory of strategic behavior posits that forecasters compete in a *forecasting contest* with pre-specified rules. A number of forecasting contests, such as the Wall Street Journal semi-annual forecasting survey, are run regularly among economists.[7] More generally, evaluators often rank forecasters by their relative accuracy (see e.g. Stekler (1987)). We find that reporting the best predictor on the state (the posterior mean conditional on the signal observed) is not necessarily an equilibrium in the contest. The equilibrium strikes a balance between two contrasting effects. Individual forecasters wish to report an accurate forecast, which is most likely to be on the mark, but at the same time they have an incentive to differentiate their prediction from those of the competitors. The second effect is due to the fact that the other forecasters might be even closer to the realized state. We solve for the equilibrium in a simple winner-take-all contest in which forecasters put too much weight on their private signals. Compared to honest forecasting, equilibrium forecasts are excessively differentiated.[8]

According to our second theory, forecasts are "cheap talk" made by forecasters who wish to foster their reputation for being well informed. This *reputational cheap talk* theory posits that forecasters are motivated by their professional prospects, as summarized by the reputation about their forecasting talent. The market evaluates the talent on the basis of the forecasts released and the realization of the state of the world. Somehow counter-intuitively, honest forecasting typically cannot occur in equilibrium. The model's predictions depend on the level of market rationality. When the market naively believes in honest forecasting, the forecast which guarantees the highest expected reputational payoff is not the best predictor of the state. When the market is naive, the optimal individual forecast turns out to be equal to the posterior expectation conditional on a signal equal to the true conditional expectation on the state. This is because forecasters want to pretend they have a signal equal to the posterior mean, which is the one most likely to be observed by a well-informed forecaster. Deviation forecasts are then biased toward the prior mean. At equilibrium, the market must have rational expectations regarding the forecasters' behavior, so the incentive to herd is self-defeating. Equilibrium forecasts are then unbiased but systematically less precise than if the forecaster were non strategic. This theory predicts that the economists' desire to be perceived as good forecasters results

[7] For a meteorological forecasting contest see e.g. the National Collegiate Weather Forecasting Contest.

[8] As discussed in more detail later, Laster, Bennett and Geoum (1999) consider a similar forecasting contest in a symmetric information environment. In the absence of heterogeneity of information, there is no reason to reward accurate forecasters. It turns out that the symmetric equilibrium in the forecasting contest is much more appealing once private information is introduced, being in pure rather than mixed strategies.

in them becoming poor forecasters.[9]

The paper is organized as follows. Section 2 sets up the statistical model. Section 3 develops the forecasting contest theory. Section 4 illustrates the reputational cheap talk theory. Section 5 compares the predictions of the different theories and extends the model to allow for a common unpredictable component. Section 6 discusses the role of anonymity in forecasting surveys. Section 7 concludes.

2. Model Setup

Information. There are n forecasters who issue forecasts on an uncertain state of the world x. It is common prior belief that the state is normally distributed with mean μ and precision ν, $x \sim N(\mu, 1/\nu)$, with p.d.f. $q(x) = \sqrt{\nu/2\pi}\exp\left(-\nu(x-\mu)^2/2\right)$. Each forecaster i has the private signal $s_i = x + \varepsilon_i$ on the state of the world. Conditionally on state x, signal s_i is assumed to be normally distributed with mean x and precision τ_i, $s_i|x \sim N(x, 1/\tau_i)$, with the p.d.f. $g_i(s_i|x) = \sqrt{\tau_i/2\pi}\exp\left(-\tau_i(s_i-x)^2/2\right)$. Forecaster i's observation of signal s_i leads to a normal posterior belief on the state with mean $E[x|s_i] = (\tau_i s_i + \nu\mu)/(\tau_i + v)$ and precision $\tau_i + \nu$ (cf. DeGroot (1970)). The p.d.f. of this posterior distribution is denoted by $q(x|s_i)$.

Honest Forecasting. Our benchmark forecast is the honest report of the posterior expectation, $h_i(s_i) = E[x|s_i]$, as assumed by most empirical investigations. In the normal model the posterior expectation minimizes the mean of any symmetric function of the forecast error, such as the mean squared error (cf. Bhattacharya and Pfleiderer (1985)). The honest forecast incorporates all available evidence s_i.

The honest forecast $h_i(s_i)$ has the important statistical property of being uncorrelated with its forecast error $h_i(s_i) - x$: $E[E[x|s_i](E[x|s_i] - x)] = E[E[E[x|s_i](E[x|s_i] - x)|s_i]] = E[E[x|s_i]E[E[x|s_i] - x|s_i]] = 0$. This implies that the honest forecast does not carry information about its forecast error. This orthogonality property may seem a necessary property of rational forecasts, but is not so. Asymmetric scoring rules generally result in forecasts violating the property, as exemplified by Zellner (1986). In the symmetric normal model it is still most natural to impose a symmetric scoring rule on an individual forecaster. However we will show below that some symmetric games can lead rational players to use non-orthogonal forecasts. The empirical literature on rational expectations (e.g. Keane and Runkle (1990, 1998)) assumes that rationality implies honest forecasting and tests for orthogonality to see if forecasters are rational.

Maximum Likelihood. Forecaster i's maximum likelihood estimator (MLE) would instead be s_i, regardless of the prior information. Since $E[s_i|x] = x$, the MLE forecast is

[9] See "Dustmen as Economic Gurus" (*The Economist*, 3 June 1995) on the good performance of a sample of London dustmen as forecasters of key economic variables.

unbiased and the forecast error ε_i is independent of x. However, the maximum likelihood forecast violates the orthogonality property: $E[s_i(s_i - x)] = E[(x + \varepsilon_i)\,\varepsilon_i] = 1/\tau_i > 0$. Note that MLE arises from Bayesian updating with the uniform distribution on the real line as the improper prior on the state and signals normally distributed conditionally on the state. We have chosen instead to posit a normal prior on x to reflect the fact that typically forecasters have some prior idea about the location of the variable to be predicted. The presence of prior public information drives the results obtained below.[10]

3. Forecasting Contest Theory

It is natural to investigate whether the competitive pressure from other forecasters provides incentives to incorporate all available information in the forecast. As discussed in the introduction, forecasters often participate in contests in which the best forecasters are rewarded. Consider the following simultaneous winner-take-all forecasting contest. First, the n forecasters simultaneously submit their forecasts $c_i \in \mathbb{R}$ after each of them has observed their private signal s_i. Second, the true state x is publicly observed and the forecaster whose forecast c_i turns out to be closest to x wins a prize $n > 0$, assumed to be proportional to the total number of forecasters participating in the contest. In the case of a tie, all winners share the prize evenly.

There is remarkably little previous work on forecasters' behavior in contests. In the applied probability literature, Steele and Zidek (1980) were the first to study a sequential forecasting contest among two privately informed forecasters. They assumed away game-theoretic considerations by supposing that the first forecaster reports truthfully. After observing the first forecast, the second guesser faces a simple decision problem and has a clear advantage. Steele and Zidek and some follow-up papers focused on the characterization of the second guesser's probability of winning. Similarly, Kutsoati and Bernhardt (2000) adopt a sequential approach to forecasting contests. They provide evidence of strategic bias in the earnings forecasts released by financial analysts. The last forecaster strategically selects a forecast that overshoots the consensus forecast in the direction of the private information.

No-one seems to have noticed that forecasting in a contest is a problem of strategic location. A forecasting contest with no private information — or equivalently with perfectly correlated information, as in Laster et al. (1999) — is identical to Hotelling's (1929) pure location game. As is well known (cf. Osborne and Pitchik (1986)), equilibria in this classic location game crucially depend on the number of players and often involve mixed strategies. Our forecasting contest theory extends Hotelling's simultaneous location game to scenarios where the forecasters (firms) have private information on the distribution of

[10] This assumption is validated by a number of empirical studies. For example, Welch (2000) finds evidence that security analysts are strongly influenced by the prevailing consensus.

the state (location of consumers).[11] To simplify the problem, in this paper we consider the limit as n tends to infinity. In this case we obtain a simple equilibrium in pure strategies.

To play this game, forecasters form conjectures about the distribution of the opponents' forecasts and calculate their best response. Since the objective is to come closest to the true x, it is enough to conjecture about the distribution of the best forecast of the opponents, conditionally on x and s_i. In general, this distribution can be quite complicated. To simplify the analysis, we impose two further restrictions. First, we assume that the forecasters' signals s_i are statistically independent conditionally on x. Second, we assume that the signals of all the forecasters are of common precision τ.

Payoffs to forecaster i can then be re-defined as follows. Suppose forecaster i receives signal s_i and conjectures that the opponents' forecasts are distributed according to the conditional density $\gamma(c|x, s_i)$ on the real line. By conditional independence of the signals, x is sufficient for (x, s_i), so we can write $\gamma(c|x)$. Suppose forecaster i submits the forecast c_i. Since there is an infinity of opponents, their forecasts are dense on the support of γ, so forecaster i wins if and only if $x = c_i$. Conditional on winning, the number of players sharing the size-n prize is approximately $n\gamma(c_i|c_i)$. Thus, the expected payoff to forecaster i is

$$U_i(c_i|s_i) = \frac{q(c_i|s_i)}{\gamma(c_i|c_i)}.$$

3.1. Deviation

Since the logarithm is a strictly increasing function, it is equivalent to let the forecaster maximize $\log(q(c_i|s_i)) - \log(\gamma(c_i|c_i))$. Recall that $q(x|s_i)$ is the p.d.f. of the posterior normal distribution with mean $E[x|s_i] = (\tau s_i + \nu\mu) / (\tau + \nu)$ and precision $\tau + \nu$. Then $\log(q(c_i|s_i))$ is a concave quadratic function of c_i, and maximization of the first term alone would lead the forecaster to choose the honest forecast $h_i(s_i) = E[x|s_i]$. But honest forecasting fails to be an equilibrium because of the second term in the objective function $-\log(\gamma(c_i|c_i))$. In order to win the contest, it is useful to move into the territory of less common forecasts since the other contestants are then less likely to win the prize. There is an incentive to move further away from the prior mean μ because the opponents are more often close to μ:

Proposition 1 (Exaggeration in Contest Deviation) *If all other forecasters are forecasting honestly by applying the strategy $c(s) = (\tau s + \nu\mu) / (\tau + \nu)$, the contest drives forecaster i to exaggerate.*

[11] We refer to the companion paper Ottaviani and Sørensen (2001) for a more general analysis of forecasting contests.

Proof. Observe that $c|x$ is normally distributed with $E[c|x] = (\tau x + \nu\mu) / (\tau + \nu)$ and $V(c|x) = \tau / (\tau + \nu)^2$. Apart from an irrelevant constant term, we have

$$\log(\gamma(c|c)) = -\frac{(\tau+\nu)^2 (c - (\tau c + \nu\mu) / (\tau+\nu))^2}{2\tau} = -\frac{\nu^2 (c-\mu)^2}{2\tau},$$

a concave quadratic function of c with peak at μ.

The forecaster maximizes $\log(q(c_i|s_i)) - \log(\gamma(c_i|c_i))$, the difference of two concave quadratic functions. The objective function is concave when the first concave term prevails, i.e. for $(\tau + \nu) > \nu^2/\tau$ or equivalently $\nu/\tau < \left(1 + \sqrt{5}\right)/2$. If the precision ratio ν/τ is too high, the private signals are imprecise relative to the prior belief on x. This induces two effects. First, opponents with imprecise signals tend to put more weight on the prior, and therefore make $\gamma(c|c)$ more concentrated around μ. Second, forecaster i is less certain of the location of x, making $q(x|s_i)$ less concentrated around $E[x|s_i]$. The two effects go in the same direction, as $\log(q(c_i|s_i))$ becomes less concave and $\log(\gamma(c_i|c_i))$ becomes more convex.

The forecaster has a unique best reply when $\nu/\tau < \left(1 + \sqrt{5}\right)/2$. When $s_i \neq \mu$, this best reply cannot be the honest forecast $E[x|s_i]$, for at this value of c_i the first term has zero slope while the second term is sloping upwards away from μ. Honesty is not a robust strategy in this contest, since the extra term in the objective drives the best reply further away from the prior mean μ than $E[x|s_i]$. The first order condition for maximization is $((\tau+\nu) c_i - \tau s_i - \nu\mu) - \nu^2 (c_i - \mu)/\tau = 0$ solved by $c_i = (\tau^2 s_i + (\tau\nu - \nu^2)\mu) / (\tau^2 + \tau\nu - \nu^2)$.

When instead $\nu/\tau \geq \left(1 + \sqrt{5}\right)/2$, there is no best response for forecaster i to honest forecasting by the opponents. The incentive to move away from μ is now so strong that the forecaster wishes to go to the extremes $\pm\infty$. □

The forecaster is unsure about where to find x, but understands that the opponents are concentrated around μ. This 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.

3.2. Equilibrium

Having established that honest forecasting is not compatible with equilibrium, we are now ready to characterize the Bayes-Nash equilibrium. In equilibrium, individual forecasters apply for every s_i the best response to their conjecture about the opponents' distribution $\gamma(c|x)$, and this conjecture is correct given the strategies of the forecasters. Since all forecasters have the same precision, the game is symmetric and we find a symmetric equilibrium whereby all forecasters apply the same strategy:

Proposition 2 (Exaggeration in Contest Equilibrium) *For any values of ν and τ there exists a unique symmetric linear equilibrium $c(s) = Cs + (1 - C)\mu$ of the contest.*

Forecasters put more weight on their private information than according to the conditional expectation: $C = \left(\sqrt{\tau^2 + 4\nu\tau} - \tau\right)/2\nu > \tau/(\nu + \tau)$.

Proof. The proposition is proved by guessing that the forecasters use linear strategies of the form $c(s) = Cs + (1 - C)\mu$. Consider the best response problem for forecaster i. Given that the opponents use a linear strategy, conditionally on x their forecasts are normally distributed with mean $Cx + (1 - C)\mu$ and variance C^2/τ. Up to a constant term we have

$$\log(\gamma(c|c)) = -\frac{\tau\left(c - (Cc + (1 - C)\mu)\right)^2}{2C^2} = -\frac{\tau(1 - C)^2(c - \mu)^2}{2C^2}.$$

The necessary first order condition for maximization of $\log(q(c_i|s_i)) - \log(\gamma(c_i|c_i))$ is then $-((\tau + \nu)c_i - \tau s_i - \nu\mu) + \tau(1 - C)^2(c_i - \mu)/C^2 = 0$, solved by

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

The fixed-point condition for a symmetric Nash equilibrium is that this linear strategy be equal to the one posited, or $C^2\tau = \left(C^2(\tau + \nu) - (1 - C)^2\tau\right)C$.

Either C is zero, or this reduces to $(1 - C)\tau = C^2\nu$. For large values of C, and at $C = 1$ the right-hand side (RHS) of this quadratic equation exceeds the left-hand side (LHS). At $C = 0$ and $C = \tau/(\tau + \nu)$ the LHS instead exceeds the RHS. We conclude that the equation has two solutions, one negative, and the other in the interval $(\tau/(\tau + \nu), 1)$. Recall that $C = 0$ also solved the original equation, but this possibility can be ruled out directly, since in this case the opponents' forecasts are all concentrated on $c = \mu$. But then all replies other than μ yield forecaster i higher payoff, so that $c = \mu$ cannot be a symmetric equilibrium.

The second-order condition for the forecaster's optimization requires that the quadratic objective is concave. This is satisfied when $\tau + \nu \geq \tau(1 - C)^2/C^2$. By using $(1 - C)\tau = C^2\nu$, this condition can be reduced to $\tau \geq -\nu C$. The positive solution for C clearly satisfies this condition. Inserting $-\tau/\nu$ into the quadratic equation for C, its LHS becomes $(\nu + \tau)\tau/\nu$ while the RHS becomes τ^2/ν and we see that the LHS exceeds the RHS. The negative solution to the equation is then below $-\tau/\nu$, and it consequently violates the second order condition. In conclusion, $C = \left(\sqrt{\tau^2 + 4\nu\tau} - \tau\right)/2\nu$ in the unique linear symmetric equilibrium. □

As in the honest forecast, the weight on the signal is increasing in τ and decreasing in ν. For all values of ν and τ, this weight is larger than in the honest forecast, so the contest gives an incentive to move away from μ.[12] The equilibrium strikes a balance: opponents

[12] It turns out that this equilibrium in linear strategies exists for all parameter values. This is slightly surprising in light of the fact that a bounded best reply to honest forecasting by the opponents only existed for a subset of the parameter space. Intuitively, when the opponents increase their weight on the signal, they become less concentrated around μ, and this mitigates the inventive to move away from μ. See also Section 5.3 on this.

disperse themselves to such an extent that forecaster i is happy to reply precisely with the same dispersion. As with the deviation forecast, the equilibrium forecast is positively correlated with its error.

The contest equilibrium satisfies $C < 1$, so the forecast is not as extreme as the maximum likelihood estimate (MLE). However, the MLE is in fact robust to the contest if we take seriously the specification with the uniform distribution on the real line as improper prior on x. 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 this constant term does not distort the forecaster's problem. The contest distortion thus depends on the presence of prior information that anchors the forecasts of the opponents around μ. The tendency of opponents to be clustered around the prior mean drives forecasters away from it.

4. Reputational Cheap Talk Theory

In a forecasting contest, competition takes place according to rules which are clearly set out in advance. The way in which the market implicitly rewards ex-post successful performance is instead often subtler and less structured. For example, the labor and financial markets perform informal (or subjective) evaluations of the forecasters' track record and performance. This Section develops a positive theory of forecasting in which forecasters aim to impress the market with their talent. It is convenient to think of the market as performing the passive role of an evaluator. Instead of committing ex ante to a particular reward scheme, the market optimally evaluates ex post the forecasting ability based on all the information available. Forecasters with a better reputation can provide more valuable information and can therefore command higher compensation. To foster their careers, forecasters want to develop a good reputation for accuracy.

This theory is based on a model first proposed in the second part of Holmström (1982) to analyze the behavior of agents who aim at convincing the market that they are well informed. We follow Scharfstein and Stein (1990) by maintaining the assumption that the forecasts do not influence the realization or the observability of the state of the world.[13] Reputational forecasting is a game of cheap talk (Crawford and Sobel (1982)), since a forecaster's payoff does not depend directly on the forecast released, but only indirectly through the interest in the market's evaluation of ability.[14]

The structure of our basic statistical model needs to be extended to introduce a latent parameter representing the unknown talent t_i of forecaster i. We further assume that forecasters and the market share a common non-degenerate prior belief $p_i(t_i)$ on forecaster

[13]See Ottaviani and Sørensen (2000) for a discussion of the logical connection of the reputational herding model of Scharfstein and Stein (1990) and Trueman (1994) with the statistical herding model of Bikhchandani, Hirshleifer and Welch (1992) and Banerjee (1992).

[14]This is a crucial difference with the reputational signaling model of Prendergast and Stole (1996).

i's talent, with all the talents t_i and the state x statistically independent. Conditionally on state x and talent t_i, the signal s_i is distributed with p.d.f. $\tilde{g}(s_i|x,t_i)$. To obtain the normal distribution of signal s_i conditional on state x, $s_i \sim N(x, 1/\tau_i)$, we must have

$$\int_0^\infty \tilde{g}(s_i|x,t_i)\, p_i(t_i)\, dt_i = g_i(s_i|x) = \left(\frac{\tau_i}{2\pi}\right)^{1/2} \exp\left(-\frac{\tau_i (s_i - x)^2}{2}\right). \tag{4.1}$$

It is well known from statistics that it is impossible to satisfy equation (4.1) with $\tilde{g}(s_i|x,t_i)$ also normally distributed around x and a non-degenerate $p_i(t_i)$. Among the many possible specifications of $\tilde{g}$ and p_i satisfying equation (4.1), it is natural to posit that $s_i|x,t_i$ has an exponential power distribution (Box and Tiao (1973), page 517) with p.d.f.

$$\tilde{g}(s_i|x,t_i) = K_1 t_i^{1/4} \exp\left(-\frac{t_i (s_i - x)^4}{12}\right), \tag{4.2}$$

and let $1/t_i$ follow a Gamma distribution such that $t_i > 0$ has p.d.f.

$$p_i(t_i) = K_2 \tau_i^{3/2} t_i^{-7/4} \exp\left(-\frac{3\tau_i^2}{4t_i}\right). \tag{4.3}$$

The exponential power distribution is similar to the normal, except for the power 4 appearing where the normal has power 2. This information structure is still a location experiment with log-concave conditional density $\tilde{g}$. Lehmann's (1988) Theorem 5.3 then guarantees that the talent t_i parametrizes forecaster i's effectiveness. The precision parameter τ_i can be thought of as an average of the true underlying talent t_i, since $E[1/t_i] = 1/\tau_i^2$. Appendix A provides the constants K_1 and K_2 and verifies (4.1).

We further simplify the problem by removing the strategic interaction among forecasters.[15] For the remainder of this Section, we will then focus on a single forecaster and remove the subscript i. The game proceeds as follows. First, the forecaster issues a forecast (or message) m after observing the private signals s. Second, the true state x is observed by the market which uses (m,x) to update the belief $p(t)$ about the forecaster's precision t. The forecast m thus serves the role of a signal sent to the market about s, and the observation of additional information x allows for sorting. The forecaster's goal is to obtain a favorable updating on the precision, with the understanding that the market rewards a good reputation.

To update the beliefs about the forecaster's talent, the evaluator applies a *conjecture* on the forecaster's strategy mapping s into m and derives the distribution of m conditional on x and t, denoted by $\varphi(m|x,t)$. Bayes' rule is then used by the evaluator to

[15] This can be done by assuming that there is neither payoff nor statistical interaction among the forecasters. First, conditionally on x and $t_1, \ldots, t_n$, the n forecasters' private signals s_i are conditionally independent, so that forecaster i cannot signal anything to the market about t_j when $i \neq j$. Second, forecaster i's objective depends solely on the posterior beliefs about t_i, regardless of the posterior beliefs about t_j.

calculate the posterior reputation $p(t|m,x) = \varphi(m|t,x)\,p(t)/\varphi(m|x)$ where $\varphi(m|x) = \int_0^\infty \varphi(m|t,x)\,p(t)\,dt$. To model the forecaster's preferences over posterior reputations, we take a two-step von Neumann-Morgenstern formulation. We assume that the forecaster correctly knows the φ function used by the evaluator, so that the forecaster can predict how the posterior reputation is calculated. The utility of reputation $p(t|m,x)$ is then given by its expected value of the Bernoulli function $u(t)$,

$$W(m|x) \equiv \int_0^\infty u(t)p(t|m,x)\ dt. \tag{4.4}$$

Recalling that higher t means that the signal s_i is more valuable in a wide range of decision problems, the market is assumed to better reward those forecasters with a (first-order stochastically) better reputation. We secure this property by assuming 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 $q(x|s)$. As different forecasters have different beliefs about the state, there is some hope for separation in their strategies. The forecaster then chooses the message m which maximizes the expected W,

$$U(m|s) \equiv \int_{-\infty}^\infty W(m|x)q(x|s)\ dx. \tag{4.5}$$

4.1. Deviation

Consider first what happens when the evaluator conjectures that the forecaster applies the benchmark honest strategy $h(s) = (\tau s + \nu\mu)/(\tau+\nu)$. We prove that the forecaster who maximizes $U(m|s)$ has a simple deviation from this strategy:

Proposition 3 (Conservatism in Reputational Deviation) *If the evaluator conjectures honest forecasting $h(s)$, the forecaster biases the forecast towards the prior mean by reporting*

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

Proof. If the evaluator conjectures honest forecasting, observation of $m = h(s)$ and x allows the inference of the realized signal $\hat{s} = h^{-1}(m)$ and error $\hat{\varepsilon} = \hat{s} - x$. The updating of the reputation is then

$$p(t|m,x) = \frac{\tilde{g}(\hat{s}|x,t)\,p(t)}{g(\hat{s}|x)} \propto t^{-3/2}\exp\left(-\frac{t\hat{\varepsilon}^4}{12} - \frac{3\tau^2}{4t} + \frac{\tau\hat{\varepsilon}^2}{2}\right).$$

This posterior reputation satisfies two intuitive properties. First, due to the symmetry of the distributions, 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 $t < t'$, we see that the likelihood ratio $p(t|m,x)/p(t'|m,x) \propto \exp\left((t'-t)\hat{\varepsilon}^4/12\right)$ is increasing in $|\hat{\varepsilon}|$. A higher error shifts weight to lower values of t, making $p(t|m,x)$ worse in the first-order stochastic dominance order. These two properties imply that $W(m|x)$ is a strictly decreasing function of $|h^{-1}(m) - x|$ (see Milgrom (1981)).

Consider the best response of the forecaster with signal s to such updating by the evaluator. The forecaster's posterior distribution on the state is normal with mean $h(s)$ and variance $1/(\nu+\tau)$. If the forecaster submits message m, the evaluator infers that the signal is equal to $h^{-1}(m)$ and the forecast error is $h^{-1}(m) - x$. The inferred forecast error is then normally distributed with mean $h^{-1}(m) - h(s)$ and variance $1/(\nu+\tau)$. The forecaster's best reply maximizes the expected value of W, or equivalently minimizes the expected value of $-W$ which is a symmetric loss function of the error. The forecaster then chooses m such that the error has mean zero, by letting $h^{-1}(m) = h(s)$. □

This fundamentally new insight into herding is quite intuitive. The forecaster wants to be perceived as having a signal s equal to the posterior expectation $h(s)$ on the state, which is most likely to result in favorable reputation updating. But unless the posterior expectation puts zero weight on the prior belief (as with a perfectly informative signal or an improper prior), the signal is more extreme than the posterior expectation. So, if the market naively believes that the forecast reflects truthfully the forecaster's posterior expectation, the forecaster deviates by reporting $d(s) = h(h(s))$. 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.

4.2. Equilibrium

We have seen that if the market expects the honest forecasting strategy $h(s)$, the forecaster will deviate to the iterated $h(h(s))$. Soon the market should learn this fact, and it should start to base its updating on the conjecture that the forecaster's strategy is $h(h(s))$. However, by the same reasoning, the forecaster will then again fool the market by deviating to $h(h(h(s)))$ and so on. Avoiding this spiral of ever more sophisticated beliefs, a Nash equilibrium is a rest point of this process.

According to Proposition 3, honest forecasting is incompatible with equilibrium. As this is a cheap-talk game, a stronger implication follows: there exists no fully separating equilibrium. By definition, in a fully separating equilibrium, the strategy mapping signals into forecasts can be inverted. As before, the evaluator infers the signal through inversion of the strategy, but the forecaster with signal s then wishes to deviate to the forecast corresponding to signal $s' = E[x|s]$, which is different from s whenever $s \neq \mu$.

Non-existence of a fully separating equilibrium is not a particularly surprising finding in a cheap-talk setting. Another common property of cheap-talk games is that there exists an equilibrium with complete pooling. In such a babbling equilibrium, the forecaster issues the same message m regardless of the signal received, and any message received by the

evaluator is interpreted as carrying no information about the signal. More generally, not all information is conveyed to the evaluator. Equilibrium forecasting must necessarily involve some degree of pooling (or bunching) of signals into messages.

The reputational equilibrium forecast satisfies the orthogonality property. Due to the cheap talk nature of the game, the actual language used to send equilibrium messages is indeterminate. But the market can easily translate message m into the best estimate conditionally on m, namely $E[x|m]$. So, the forecaster is effectively communicating $E[x|m]$ to the evaluator. Being a conditional expectation of x, this forecast is uncorrelated with its error. We conclude:

Proposition 4 (Coarseness in Reputational Equilibrium) *There is no reputational cheap talk equilibrium with full revelation of information. There exist partially and fully pooling equilibria. Any equilibrium can be defined with a language such that the forecast has the orthogonality property.*

Because of the coarseness of the message, the evaluator loses information regarding not only the location of the state x, but also the forecaster's talent t. More precise information on t would be a useful input to assess the importance of future forecasts from the same person, so this will again imply a loss of information about future states of the world.

Next we show by example that there exists a partially separating equilibrium whereby some information is conveyed through the messages. In this example there are two messages, naturally defined from the model's symmetry around μ. The equilibrium strategy of the forecaster is to report a high message m_H whenever the signal s weakly exceeds the prior mean μ and a low message m_L when $s < \mu$. A forecaster is indifferent between these two messages when receiving signal μ. When observing a higher signal $s > \mu$, the forecaster expects high values of the state to be realized and prefers to send message m_H rather than m_L in order to indicate a positive signal, implying smaller average errors:

Result 1 (Binary Reputational Equilibrium) *In the reputational cheap talk model there exists a symmetric binary equilibrium.*

Proof. When observing message m_H and state x, the evaluator infers that the signal's error satisfies $\varepsilon = s - x \geq \mu - x$. Similarly, evaluator infers $\varepsilon < \mu - x$ upon observation of m_L and x. For any t, the error ε is equally likely as the error $-\varepsilon$ by (4.2). Sending message m_H in state x carries the exact same information about t as sending message m_L in state $2\mu - x$, which is equally far, from but on opposite sides of, μ. This implies the symmetry property $W(m_H|x) = W(m_L|2\mu - x)$.

As shown in the proof of Proposition 3, the smaller the error perceived by the evaluator, the more favorable the posterior reputation. When $x < \mu$, message m_H conveys that the error is at least as large as the positive number $\mu - x$. It can be verified that this message is worse news about the talent than the neutral message that $\varepsilon \geq 0$, which also contains some

of the more favorable errors close to zero. Thus we have $W(m_H|x) < W(m_H|\mu)$ for $x < \mu$. Symmetrically, it is good news about the talent to see that ε exceeds some negative number, since this observation includes more errors close to zero. Thus $W(m_H|x) > W(m_H|\mu)$ when $x > \mu$. Applying symmetry, we conclude from the two inequalities that $W(m_H|x) > W(m_L|x)$ when $x > \mu$.

Next, we show that the forecaster does not wish to deviate from the putative equilibrium strategy. By symmetry, it suffices to assume that $s \geq \mu$ and check that $U(m_H|s) \geq U(m_L|s)$. Using (4.5) and the symmetry of W, we have

$$\begin{aligned} U(m_H|s) - U(m_L|s) &= \int_{-\infty}^{\infty} (W(m_H|x) - W(m_L|x))\, q(x|s)\, dx \\ &= \int_{\mu}^{\infty} (W(m_H|x) - W(m_L|x))\, (q(x|s) - q(2\mu - x|s))\, dx. \end{aligned}$$

Since $q(x|s)$ is the p.d.f. of a unimodal 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 already had $W(m_H|x) > W(m_L|x)$ when $x > \mu$, so the integrand of the last integral is everywhere non-negative, implying that the integral is non-negative, i.e. that $U(m_H|s) \geq U(m_L|s)$ as desired.

Since the message space has not been restricted, we also need to specify the evaluator's beliefs following any third out-of-equilibrium message. We posit that the evaluator in this case assumes that the forecaster possessed a signal $s < \mu$, resulting in the same posterior reputation as message m_L. These beliefs satisfy the requirements of a perfect Bayesian equilibrium. □

Rather than performing direct tests of reputational cheap talk, most of the existing empirical literature provides indirect evidence of reputational concerns based mostly on heterogeneity across forecasters. Lamont (1995) finds that older forecasters tend to deviate more from the consensus. Chevalier and Ellison (1999) find that older mutual fund managers have bolder investment strategies. Hong, Kubik and Solomon (2000) conclude that the lower accuracy of older stock analysts is due the fact that they move earlier. Unfortunately, no one has so far attempted to model the endogenous timing of forecasts when the agents are concerned about their reputation or relative accuracy.[16]

The reputational cheap talk theory can be extended to allow for private information of forecasting ability, mixed objectives, and concern for relative reputation among forecasters. We refer to the companion paper Ottaviani and Sørensen (1999) for a more general theoretical analysis and discussion of the empirical literature. See also Zitzewitz (2001b) for an interesting model in which the market uses a simple econometric technique to evaluate the quality of the information contained in the forecasts.

[16] In Gul and Lundholm (1995) forecasters care about the absolute accuracy as well as delay. A forecaster with a more extreme signal acts earlier at equilibrium.

5. Discussion

5.1. Forecast Variability

Except for the reputational equilibrium forecasts, we have found 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 0 and 1. From the given distribution of s_i, we derive the conditional distribution of the linear forecast,

$$f_i|x \sim N\left(F_i x + (1 - F_i)\mu, F_i^2/\tau_i\right). \tag{5.1}$$

We now describe the variance of the forecasts under the different theories.

Honest Forecast. The conditional variance of the honest forecast $V(h_i|x) = \tau_i/(\tau_i + \nu)^2$ is increasing in the signal precision if the forecaster is imprecise enough ($\tau_i < \nu$), but decreasing when instead the expert is precise ($\tau_i > \nu$). When the signal is poorly informative ($\tau_i \approx 0$), the honest forecast is concentrated on the prior mean and so the conditional variance is also near 0. When τ_i is very large, the conditional variance is again near 0 because a perfectly informative signal gives an honest forecast concentrated on the true state.

Contest Equilibrium Forecast. The conditional variance of the contest forecast is $C^2/\tau = \left(2\nu + \tau - \sqrt{\tau^2 + 4\nu\tau}\right)/2\nu^2$. The conditional variance of forecasts converges to $1/\nu$ in the limit as the private signal becomes uninformative $\tau \to 0$. This concords with Osborne and Pitchik's (1986) finding that with infinitely many symmetrically informed forecasters the distribution of equilibrium locations replicates the common prior distribution for x. Addition of private information has the desirable effect of inducing a symmetric location equilibrium in pure rather than mixed strategies. Note that the conditional variance of the distribution of the equilibrium contest forecasts decreases in τ and converges to 0 as $\tau \to \infty$. This is in sharp contrast to the non-monotonicity of the conditional variance of the honest forecast. In cases with imprecise private signals, one could check empirically whether forecasts are very widely dispersed as in the contest, or quite close together as in the case of honest forecasting.

Reputational Deviation Forecast. The conditional variance of the reputational deviation forecast is $\tau^3/(\tau + \nu)^4$ with variance first increasing in τ, maximal at $\tau = 3\nu$, and then decreasing in τ. Compared to the forecasts under truthtelling and in the forecasting contest, the reputational deviation forecast puts more weight on the prior mean, and is therefore less variable.

Reputational Equilibrium Forecast. The reputational equilibrium forecast of Result 1 is binomially distributed and therefore not directly comparable with the normally distributed forecasts which result in the cases described above. When forecasters mean what

they say, $m_L = E[x|s < \mu]$ and $m_H = E[x|s \geq \mu]$. By applying the well-known result that $E[y|y > 0] = \sigma\sqrt{2/\pi}$ for a normal variable $y \sim N(0, \sigma^2)$ (cf. Johnson and Kotz (1970)), we see that $E[x|s \geq \mu]$ is equal to

$$E[E[x|s]|s \geq \mu] = E[\frac{\tau s + \nu\mu}{\tau + \nu}|s \geq \mu] = \mu + \frac{\tau}{\tau + \nu} E[s - \mu|s > \mu] = \mu + \sqrt{\frac{2\tau}{\pi\nu(\tau + \nu)}}.$$

By symmetry, we have $m_L = \mu - \sqrt{2\tau/\pi\nu(\tau + \nu)}$. The more precise the signal of the forecaster (i.e., the greater τ), the greater the amount by which the message moves the prior beliefs on the state.

Given x, the chance of the forecast taking the high value $E[x|s \geq \mu]$ is $1-\Phi(\sqrt{\tau}(\mu - x))$ where Φ is the c.d.f. of the standard normal distribution. Thus

$$E[r|x] = \mu + \left(1 - 2\Phi\left(\sqrt{\tau}(\mu - x)\right)\right)\sqrt{\frac{2\tau}{\pi\nu(\tau + \nu)}}$$

and the variance of the binomial distribution is

$$V[r|x] = 4\left(1 - \Phi\left(\sqrt{\tau}(\mu - x)\right)\right)\Phi\left(\sqrt{\tau}(\mu - x)\right)\frac{2\tau}{\pi\nu(\tau + \nu)}$$

For any x, we have $\Phi(\sqrt{\tau}(\mu - x)) \in [0, 1]$ so $4(1 - \Phi(\sqrt{\tau}(\mu - x)))\Phi(\sqrt{\tau}(\mu - x)) \leq 1$, where the bound is tight being achieved for $x = \mu$. We conclude that $V[r|x] \leq 2\tau/(\pi\nu(\tau + \nu))$.

To compare this with the variance of the honest forecast, observe that

$$\frac{2\tau}{\pi\nu(\tau + \nu)} \leq \frac{\tau}{(\tau + \nu)^2}$$

if and only if $\tau/\nu \leq (\pi - 2)/2$. If the signal is not very precise, the conditional variance is uniformly higher under honesty than under the binary equilibrium. But the inequality is reversed when the signal has high precision and x is close to μ. In that case, the signals and thus the honest forecast is highly concentrated near μ, while the reputational forecast is highly variable with even chance of a positive and negative update of amount $\sqrt{2\tau/\pi\nu(\tau + \nu)}$.[17]

Comparison of Variances. Figure 2 plots the conditional variances as a function of τ when $\nu = 1$ for four forecasts: honest h, contest c, reputational deviation d, reputational binary equilibrium r. This graph makes clear that herding or exaggeration can be inferred

[17] Graham (1999) performs comparative statics on the second mover's incentives for deviation from a separating equilibrium with respect to changes in prior reputation, forecast ability and conditional signal correlation. Our analysis in this paper shows that equilibrium comparative statics is substantially different.

from forecast dispersion only after controlling for the quality of the forecaster's information. This point is also emphasized by Zitzewitz (2001a).

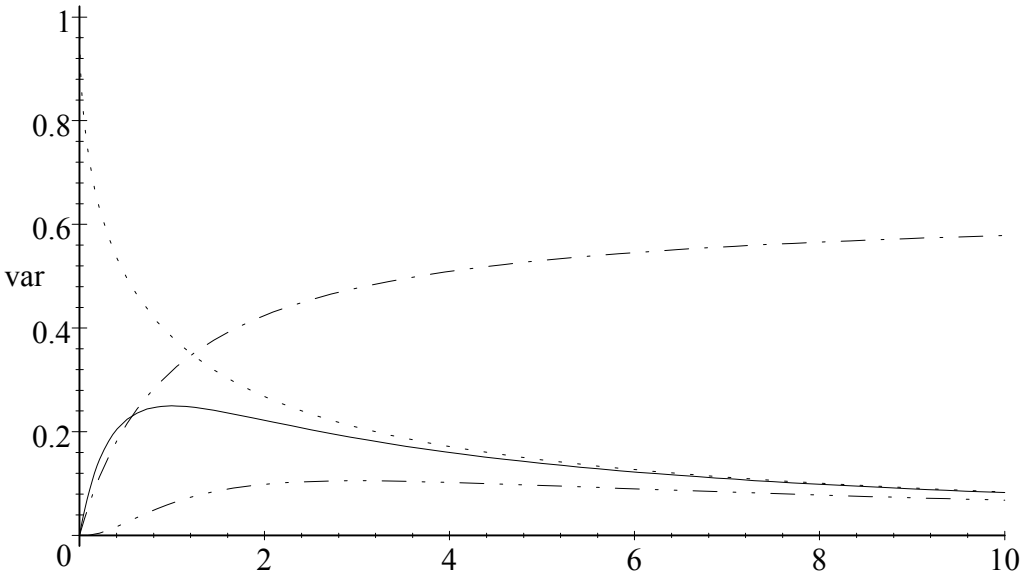

Figure 2. Conditional variances of forecasts as function of average precision τ. The solid line corresponds to the conditional variance of the honest forecast $V[h|x]$, long dashes to the reputational deviation forecast $V[d|x]$, medium dashes to the reputational equilibrium forecast $V[r|x=\mu]$, and short dashes to the contest forecast $V[c|x]$.

5.2. Forecast Error

The linear forecasts (5.1) with weight F_i less than 1 are not unbiased, since $E[f_i|x] > x$ when $x < \mu$ and $E[f_i|x] < x$ when $x > \mu$. The conditional mean of the linear forecast only goes part of the way from the prior μ to the true state x. We see that all these forecasts, including the honest one, possess the oft-lamented property that forecasters fail to predict extreme values. This property is therefore not evidence of inefficient conservativeness on the part of the forecasters.

In a similar vein, the forecast error $f_i - x$ is negatively correlated with x (since $E[(f_i - x)x] = E[(F_i\varepsilon_i + (1-F_i)(\mu - x))x] = -(1-F_i)/\nu < 0$). The forecast error can be predicted after the state x has been observed, even if forecasters are not conservative. When x is high (low) the error tends to be negative (positive). Furthermore, even if their signals are conditionally independent, there is correlation among the forecast errors of any pair i, j of forecasters. Thus, assume that ε_i is independent of ε_j. The covariance of $f_i - x = (F_i\varepsilon_i + (1-F_i)(\mu - x))$ with $f_j - x$ is $(1-F_i)(1-F_j)/\nu > 0$. This covariance is positive because forecasters tend to make equal-signed errors of opposite sign of x.

To remove this correlation, one could alternatively study forecaster i's shock as $f_i - E[f_i|x] = F_i\varepsilon_i$. This error follows a normal distribution with mean 0 and variance F_i^2/τ_i,

and is uncorrelated with x and with the errors of other forecasters. This makes them useful observations for regression analysis. Empirically, $E[f_i|x]$ may not be known by the data analyst even when x has been realized, and it is a common approach to estimate it using the consensus forecast, to which we now turn.

When $n \geq 2$ forecasters have issued their forecasts, it is simple to calculate the unweighted average forecast $\bar{f} = \sum_{i=1}^{n} f_i/n$, often referred to as the *consensus forecast*. In general, this is not the optimal forecast given the n signals.[18] For example, when the individual errors ε_i are statistically independent, the correlation of the honest 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.[19] When all forecasters have equal precision and independent errors, we see that $\bar{f}$ converges almost surely to $E[f|x]$ as $n \to \infty$ by the strong law of large numbers. Thus, asymptotically the shocks relative to the consensus $f_i - \bar{f}$ have the desirable uncorrelation properties.

5.3. Orthogonality

Consider linear forecasting rules of the form $f_i(s_i) = F_i s_i + (1 - F_i)\mu$. We have already noted that under honesty the forecast is uncorrelated with its error $f_i - x$ when $F_i = \tau_i/(\nu + \tau_i)$. More generally, the correlation is

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

Clearly, this correlation has the same sign as $F_i - \tau_i/(\tau_i + \nu)$. There is positive correlation when F_i is larger as in the contest, and negative correlation when F_i is smaller as in the reputational deviation. The reputational equilibrium forecast satisfies orthogonality but is not efficient.

As noted above, all types of forecast had errors negatively correlated with the state x. Thus, after knowing x the sign of the errors could be predicted. However, this is an unreasonably strong test of the forecasters' abilities since x is still known when the forecasts

[18] An extensive literature in operations research studies the problem of how to optimally combine forecasts obtained with different methods or incorporating different information sets. See e.g. the early contributions by Bates and Granger (1969) and Bunn (1975). That literature does not, however, consider the possibility of strategic behavior by the forecasters.

[19] 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 $\bar{h} = (n\nu\mu + \tau\sum_{i=1}^{n} s_i)/(n\nu + n\tau) = (n\nu\mu + n\tau x + \tau\sum_{i=1}^{n}\varepsilon_i)/(n\nu + n\tau)$ and the error is $\bar{h} - x = (n\nu(\mu - x) + \tau\sum_{i=1}^{n}\varepsilon_i)/(n\nu + n\tau)$, so that the covariance is always negative: $E[\bar{h}(\bar{h} - x)] = -(n-1)\tau/n(\nu + \tau)^2 < 0$ for $n > 1$. Kim, Lim and Shaw (2001) suggest methods to correct for the loss of information in the consensus forecast.

are released. If they report honestly, their error cannot be predicted from the forecast. The contest forecast and the reputational deviation forecast fail instead to inherit this property, so that once a forecast has been released the sign of its error can be predicted.

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 forecast error on the forecast. Most studies report a positive correlation of the forecast and its error, consistent with the 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) cannot support any bias, as they note that the tests are not as powerful as is usually assumed. There is significant positive correlation among the residuals in this regression. This is evidence of a common error in the forecasts. In the case of GDP growth forecasts, over a year elapses from the submission of the forecasts before the data are realized. Within this year, unpredictable changes to the actual growth rate often occur. This causes the correlation found in the data.

5.4. Common Error

The model can be easily extended to account for such an ex-post innovation in the state. The n forecasters still receive signals $s_i = x + \varepsilon_i$, with the same distribution as before. We maintain the assumption that $x, \varepsilon_1, \ldots, \varepsilon_n$ are stochastically independent. However, the state later realized is not x but actually $y = x + \varepsilon_0$. We naturally assume that $\varepsilon_0 \sim N(0, 1/\tau_0)$ is independent of x and the other errors, so that the change is unpredictable. The error ε_0 plays the role of a common error in the signals about the observed state y, for forecaster i observes $s_i = y - \varepsilon_0 + \varepsilon_i$.

The honest forecast of y is the same as the honest forecast of x, since the subsequent error is unpredictable. Indeed, $E[y|s_i] = E[x + \varepsilon_0|s_i] = E[x|s_i]$ by the independence assumption. However, the posterior beliefs about y are less precise than the posterior beliefs about x due to the added error term. The variance of the normally distributed $y|s_i$ is $1/(\tau_i + \nu) + 1/\tau_0$. Let $\tilde{q}_i(y|s_i)$ denote the p.d.f. of this posterior belief.

5.4.1. Contest Theory

The contest winner is now the forecaster who gets closest to the realized state y. We maintain the assumption that all forecasters have equal precision $\tau_i = \tau$ for $i \neq 0$ and extend the definition of forecaster i's payoff as follows. Suppose forecaster i receives signal s_i and conjectures that opponents' forecasts are distributed according to the conditional density $\gamma(c|y, s_i)$ on the real line. If forecaster i submits the forecast c_i, the logarithm of the expected payoff is $\log(U_i(c_i|s_i)) = \log(\tilde{q}(c_i|s_i)) - \log(\gamma(c_i|c_i, s_i))$.

As before, the action comes from the second term. Observation of y and of s_i gives two independent sources of information on the location of x. Clearly, $x|y, s_i$ is normally

distributed with mean $(\nu\mu + \tau_0 y + \tau s_i)/(\nu + \tau_0 + \tau)$ and precision $\nu + \tau_0 + \tau$. Assuming that forecaster i conjectures that the opponents apply a linear strategy $c(s_j) = Cs_j + (1-C)\mu = Cx + C\varepsilon_j + (1-C)\mu$, the opponents' forecasts are normally distributed with $E[c|y, s_i] = (1-C)\mu + C(\nu\mu + \tau_0 y + \tau s_i)/(\nu + \tau_0 + \tau)$ and variance $C^2/(\nu + \tau_0 + \tau) + C^2/\tau$ with the last term coming from ε_j. Except for a constant term, $\log(\gamma(c_i|c_i, s_i))$ is equal to the following quadratic function of c_i:

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

Once the analysis of the forecasting contest is repeated for this extension, one finds that there is no symmetric linear Nash equilibrium when τ/ν is large relative to τ_0/ν. In this case, the forecasters have poor information on the location of y, and $\log(\tilde{q}(c_i|s_i))$ is not very concave. Still, they have good information about the state and about the other forecasters' signals, resulting in the convexity of $\log(\gamma(c_i|c_i, s_i))$ being too large when the other forecasters use best replies. When instead τ/ν is small relative to τ_0/ν, there is a linear equilibrium with similar features to the one studied in the benchmark model.

5.4.2. Reputational Theory

To extend the reputational theory, we must now take the interaction of the n forecasters more seriously, since the errors $\varepsilon_1, \ldots, \varepsilon_n$ are correlated conditionally on y. In order to update beliefs about the precision of forecaster i, the evaluator uses the information on the location of x contained in the realized y as well as in the n forecasts.

First, we verify that it is not a Nash equilibrium for every forecaster to honestly report $h(s_i)$ (or use any other fully separable strategy). We prove that forecaster 1 would wish to deviate as before — the same argument applies to all other forecasters. Under the assumption that all forecasters $2, \ldots, n$ report honestly, the market can infer their signals $s_2, \ldots, s_n$. Based on the n conditionally independent normal signals $s_2, \ldots, s_n, y$ about x, the market first updates its beliefs about x to arrive at a normal posterior as usual. From the point of view of forecaster 1, this updating is equivalent to the forecaster receiving a more precise signal z about x. The problem faced by forecaster 1 is then a slightly complicated version of the problem analyzed earlier. The market does not observe x, but only receives a signal z about it. The market assumes that 1 is honest, and inverts the strategy to infer the signal $s_1 = h_1^{-1}(m_1)$ from the message m_1. By the law of iterated expectations, the posterior reputation is $p(t|m_1, z) = E[p(t|m_1, x)|z]$ — the evaluator's beliefs on t is the average of the beliefs the market would have held when observing x. Using this fact, we arrive at $U(m_1|s_1) = \int_{-\infty}^{\infty} W(m_i|x) E[q(x|z)|s_i]\, dx$ where $E[q(x|z)|s_i]$ is the forecaster's expectation of the evaluator's beliefs on x, and where $W(m_i|x)$ is as before. It is simple to verify that $E[q(x|z)|s_i]$ is a normal density for x with a mean between μ and s_i. As before, forecaster 1 deviates in order to give the impression of having a signal equal to that posterior mean.

Second, it is again a Nash equilibrium for every forecaster to use the binary strategy of Result 1. A forecaster with signal $s_i = \mu$ regards the two possible messages as symmetric, and is thus indifferent. A forecaster with signal $s_i > \mu$ thinks it more likely that other forecasters report that they saw $s_j > \mu$, and thinks it more likely that the market observes $y > \mu$ — alas forecaster i considers it more likely that the market's posterior beliefs on x are shifted upwards. Forecaster i then prefers to send the message which signals $s_i > \mu$. We conclude that our theoretical findings are robust to the addition of ex-post noise.

Since the evaluator does not have access to x, even honest revelation of s_i does not allow for the calculation of the individual ε_i. We have seen how the evaluator modifies the Bayesian procedure in this context, averaging over the possible values of x given the available information, including the messages from the other forecasters. Recall that under a linear strategy $f_i - E[f_i|x] = F_i\varepsilon_i$ is proportional to the error and that in the case of many forecasters, the consensus forecast provides a good approximation for $E[f_i|x]$. A forecast close to the consensus then indicates a small absolute value of ε_i and is therefore good news about the forecaster's ability.

6. The Role of Anonymity

In order to test the different theories, it might be useful to compare non-anonymous with anonymous forecasting surveys. Figure 3 plots the individual forecasts for the period 1972-1993 in the *Survey of Professional Forecasters* of the Federal Reserve Bank of Philadelphia, the most prominent anonymous survey of professional forecasters (see also Stark (1997). 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 forecasts 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 confirm 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, 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 (non-anonymous surveys).

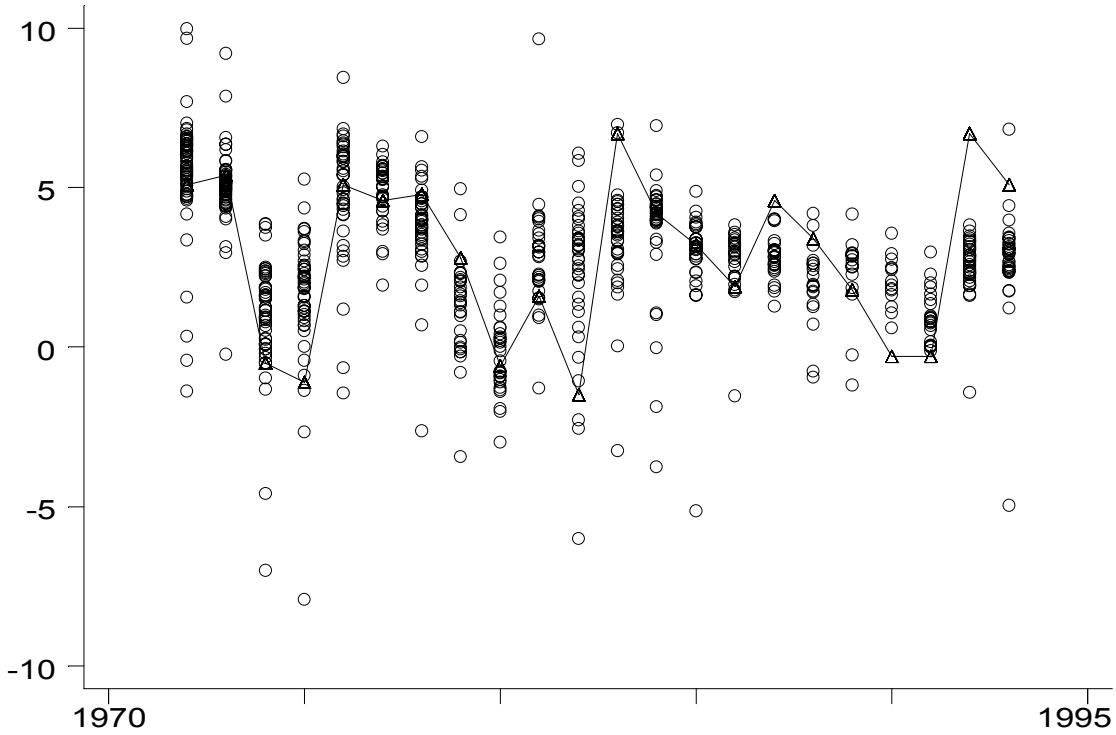

Figure 3. Comparison of individual forecasts of GNP growth from the Survey of Professional Forecasters with the realizations for the period 1972-1993. Axes and symbols as in figure 1. See Stark (1997).

A problem with the hypothesis of honest forecasting in anonymous surveys is that our theory does not predict behavior in this situation. As confirmed to us by Croushore and Lamont, 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 them to change their report, unless they have a strict incentive to do so; second, they might be concerned that their strategic behavior be uncovered to the editor of the anonymous survey.

As seen in Figure 3, the forecasts in the anonymous Survey of Professional Forecasters are widely dispersed throughout the entire period considered. In comparison with Figure 1, this suggests that less herding might be present in the anonymous survey. However, this possibility needs more careful investigation. The composition of the forecasters' panel of the Survey of Professional Forecasters has recently be made available to researchers. The joint hypothesis of honest reporting in anonymous surveys and strategic forecasting in non-anonymous surveys could nevertheless be tested by pooling in a single regression all the forecasters belonging to both data sets. It would then be possible to evaluate the usefulness of keeping anonymous surveys.

7. Conclusion

This paper has adopted a positive approach, but it suggests a number of natural normative questions. Motivated mainly by the needs of accurate probabilistic weather forecasts, a large literature in meteorology and statistics studies how to motivate forecasters to form and state their correct subjective probabilities (cf. Dawid's (1986) overview). Statisticians have developed scoring rules that elicit truthful information and avoid misrepresentation of the forecaster's beliefs (see e.g. de Finetti (1965) and Savage (1971)). Combining the insights of this literature with the economic theory of regulation, Osband (1989) studies the optimal provision of forecasting incentives in the presence of costly information acquisition.

The normative literature on scoring rules has focused on the case of single forecasters. Clearly, there is no gain from conditioning the state-contingent reward on the messages of other forecasters when the realized state is sufficient for such messages. When instead the state is observed with noise, or equivalently when forecasters possess conditionally dependent signals, conditioning the reward also on the competitors' forecasts might improve the incentives for forecast accuracy.

Forecasting is proving to be an apt laboratory for improving our understanding of strategic communication and positioning by non-partisan informed agents. The availability of data sets and the richness of institutional details can inspire and give discipline to our theorizing. The insights gained can be helpful in shedding light on a number of other social and economic problems, such as the choice of research topic by scientists. As also stressed by Banerjee (1989), rewarding originality counteracts the natural tendency to herd, but it is an imperfect tool.

Interesting phenomena emerge in simple and plausible models, without the need to depart from rationality. While this can be considered as a strength of our approach, it is also a limitation, in view of the experimental evidence on deviations from the Bayesian paradigm. For example, according to the representativeness bias (Kahneman and Tversky (1973)), forecasters often disregard prior information when making intuitive predictions. Experimental subjects tend to put excessive weight on the signal they receive and be overconfident in their predictions. The representativeness bias has similar implications to our contest model. The reputational theory can easily be extended to the case where the underlying talents of forecasters are *perceived* to be heterogenous.[20] More work needs to be done on building and testing behaviorally plausible models of forecasting.[21]

[20] In an early behavioral model of financial advice, Denton (1985) assumes that investors listen to financial advisers who have no real information. In some cases, this might well be the case. For example, Hartzmark (1991) found that futures forecasters depend on luck rather than forecasting ability. Zitzewitz (2001a) instead finds that security analysts differ greatly in performance, justifying the assumption that they are fundamentally heterogeneous.

[21] Ehrbeck and Waldmann (1996) run an early horse race between economic and phsycological explanations for biases in forecasting.

References

[1] Banerjee, Abhijit V., "Herd Behavior and the Reward for Originality," 1989, Princeton University mimeo.

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

[3] Batchelor, Roy and Pami Dua, "Conservatism and Consensus-seeking among Economic Forecasters," *Journal of Forecasting*, 1992, 11, 169–181.

[4] Bates, John M. and Clive Granger, "The Combination of Forecasts," *Operations Research Quarterly*, 1969, 20(4), 451-468.

[5] 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.

[6] Bhattacharya, Sudipto and Paul Pfleiderer, "Delegated Portfolio Management," *Journal of Economic Theory*, 1985, 36(1), 1-25.

[7] Bunn, Derek W., "A Bayesian Approach to the Linear Combinations of Forecasts," *Operations Research Quarterly*, 1975, 26(2), 325-329.

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

[9] Crawford, Vincent and Joel Sobel, "Strategic Information Transmission," *Econometrica*, 1982, 50(6), 1431–1452.

[10] Croushore, Dean, "Introducing: The Survey of Professional Forecasters," Federal Reserve Bank of Philadelphia Business Review, 1993, November/December, 3-13.

[11] Croushore, Dean, "The Livingston Survey: Still Useful After All These Years," Federal Reserve Bank of Philadelphia Business Review, 1997, March/April, 15-26.

[12] Dawid, A. Philip, "Probability Forecasting", *Encyclopedia of Statistical Sciences: Volume 7*, 1986, edited by S. Kotz, N.L. Johnson and C.B. Read, New York: Wiley-Interscience, 210-218.

[13] de Finetti, Bruno, "Methods for Discriminating Levels of Partial Knowledge Concerning a Test Item", *British Journal of Mathematical and Statistical Psychology*, 1965, 18-I, 87-123. Reprinted in *Probability, Induction, and Statistics: The Art of Guessing*, Chapter 4, London: John Wiley & Sons, 1972, 25-63.

[14] DeGroot, Morris, *Optimal Statistical Decisions*, 1970, New York: McGraw Hill.

[15] Denton, Frank T., "The Effect of Professional Advice on the Stability of a Speculative Market", *Journal of Political Economy*, October 1985, 93(5), 977-93.

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

[17] Graham, John, "Herding Among Investment Newsletters: Theory and Evidence," *Journal of Finance*, 1999, 54(1), 231–268.

[18] Gul, Faruk and Russell Lundholm, "Endogenous Timing and the Clustering of Agents' Decisions," *Journal of Political Economy*, 1995, 103(5), 1039-1066.

[19] Hartzmark, Michael L., "Luck versus Forecast Ability: Determinants of Trader Performance in Futures Markets, *Journal of Business*, 1991, 64(1), 49-74.

[20] Holmström, Bengt, "Managerial Incentive Problems: A Dynamic Perspective," in *Essays in Economics and Management in Honor of Lars Wahlbeck*, Helsinki: Swedish School of Economics, 1982. Reprinted in *Review of Economic Studies*, 1999, 66(1), 169–182.

[21] Hong, Harrison, Jeffrey D. Kubik and Amit Solomon, "Security Analysts' Career Concerns and Herding of Earning Forecasts," *RAND Journal of Economics*, 2000, 31(1), 121-144.

[22] Hotelling, Harold, "Stability in Competition," *Economic Journal*, 1929, 39(153), 41-57.

[23] International Monetary Fund, *World Economic Outlook: Fiscal Policy and Macroeconomic Stability*, May 2001.

[24] Johnson, Norman L. and Samuel Kotz, *Distributions in Statistics: Continuous Univariate Distributions - 1*, New York: Wiley, 1970.

[25] Kahneman, Daniel and Amos Tversky, "On the Psychology of Prediction," *Psychological Review*, 1973, 80(4), 237–251.

[26] Kandel, Eugene and Neil D. Pearson, "Differential Interpretation of Public Signals and Trade in Speculative Markets," *Journal of Political Economy*, 1995, 103(4), 831-872.

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

[28] Keane, Michael P. and David E. Runkle, "Testing the Rationality of Price Forecasts: New Evidence from Panel Data," *American Economic Review*, 1990, 80(4), 714–735.

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

[30] Kim, Oliver, Steve C. Lim and Kenneth W. Shaw, "The Use of Forecast Revision in Reducing Built-in Biases in Mean Analyst Forecasts," *Journal of Accounting Research*, forthcoming.

[31] Kutsoati, Edward and Dan Bernhardt, "Can Relative Performance Compensation Explain Analysts' Forecasts of Earnings?" August 2000, Tuft University and University of Illinois mimeo.

[32] Lamont, Owen, "Macroeconomic Forecasts and Microeconomic Forecasters," 1995, NBER Working Paper 5284.

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

[34] Lehmann, Eric L., "Comparing Location Experiments," *Annals of Statistics*, 1988, 16(2), 521–533.

[35] Milgrom, Paul, "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 1981, 12, 380–391.

[36] Osband, Kent, "Optimal Forecasting Incentives," *Journal of Political Economy*, 1989, 97(5), 1091–1112.

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

[38] Ottaviani, Marco and Peter Sørensen, "Professional Advice," 1999, UCL Economics Discussion Paper 99–04.

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

[40] Ottaviani, Marco and Peter Sørensen, "The Positive Theory of Forecasting Contests," 2001, University of Copenhagen and UCL mimeo.

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

[42] Savage, Leonard, "Elicitation of Personal Probabilities and Expectation," *Journal of the American Statistical Association*, 1971, 66(336), 783-801.

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

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

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

[46] Stekler, Herman O., "Who Forecasts Better?" *Journal of Business and Economics Statistics*, 1987, 5(1), 155-158.

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

[48] Welch, Ivo, "Herding among Security Analysts," *Journal of Financial Economics*, 2000, 58(3), 369-396.

[49] Zarnowitz, Victor, "An Appraisal of Short-Term Economic Forecasts," NBER occasional paper 104, 1967.

[50] Zarnowitz, Victor and Phillip A. Braun, "Twenty-two Years of the NBER-ASA Quarterly Economic Outlook Surveys: Aspects and Comparisons of Forecasting Performance" in *Business Cycles, Indicators, and Forecasting*, NBER Studies in Business Cycles, Volume 28, edited by James H. Stock and Mark W. Watson, Chicago and London: University of Chicago Press, 1993, 11-84.

[51] Zarnowitz, Victor and Luis A. Lambros, "Consensus and Uncertainty in Economic Predictions," *Journal of Political Economy*, 1987, 95(3), 591–621.

[52] Zellner, Arnold, "Biased Predictors, Rationality and the Evaluation of Forecasts," *Economics Letters*, 1986, 21, 45-48.

[53] Zitzewitz, Eric, "Measuring Herding and Exaggeration by Equity Analysts," 2001a mimeo, Stanford GSB.

[54] Zitzewitz, Eric, "Opinion-Producing Agents: Career Concerns and Exaggeration," 2001b mimeo, Stanford GSB.

Appendix: Information

The distribution of t is described by the p.d.f.

$$p(t) = \frac{3^{3/4}\Gamma(1/4)\tau^{3/2}}{4\pi} t^{-7/4} \exp\left(-\frac{3\tau^2}{4t}\right)$$

on $(0,\infty)$, where $\Gamma(v) = \int_0^\infty u^{v-1}e^{-u}du$, the well-known Gamma function. We wish to prove that this integrates to 1. We show that transformation to $y = 1/t$ gives the well-known p.d.f. of the Gamma distribution with parameters $3/4$ and $3\tau^2/4$. From $dy = -y^2 dt$, we see that the p.d.f. of y on $(0,\infty)$ is

$$y^{-2}\frac{3^{3/4}\Gamma(1/4)\tau^{3/2}}{4\pi} y^{7/4} \exp\left(-\frac{3\tau^2}{4}y\right) = \frac{(3\tau^2/4)^{3/4}}{\Gamma(3/4)} y^{-1/4} \exp\left(-\frac{3\tau^2}{4}y\right),$$

as desired. The constants are correct since $(\tau^2)^{3/4} = \tau^{3/2}$, $4^{3/4} = 2\sqrt{2}$, and the Gamma function satisfies $\Gamma(y)\Gamma(1-y) = \pi/\sin(\pi y)$ so that $\Gamma(1/4)\Gamma(3/4) = \pi/\sin(\pi/4) = \sqrt{2}\pi$.

Next we verify that equation (4.1) is satisfied. The p.d.f. of the signal conditional on the state is then

$$\int_0^\infty \tilde{g}(s|x,t)\,p(t)\,dt = \int_0^\infty \frac{\sqrt{3}\tau^{3/2}}{2\pi\sqrt{2}} t^{-3/2} \exp\left(-\frac{3\tau^2}{4t} - \frac{t(s-x)^4}{12}\right) dt.$$

Note that the density of t appearing under the integral is that of the Inverse Gaussian distribution, and its well-known specification (use e.g. (4.2) of Johnson and Kotz (1970) page 139 with their parameters $\varphi = \tau(s-x)^2/2$ and $\mu = 3\tau(s-x)^{-2}$) allows us to continue with

$$= \frac{\sqrt{3}\tau^{3/2}}{2\pi\sqrt{2}} \frac{\sqrt{4\pi/3}}{\tau} \exp\left(-\frac{\tau(s-x)^2}{2}\right) = \sqrt{\frac{\tau}{2\pi}} \exp\left(-\frac{\tau(s-x)^2}{2}\right) = g(s|x),$$

the Normal distribution with mean x and precision τ.