

Rational Bias in Macroeconomic Forecasts

David Laster
Paul Bennett
In Sun Geoum

The authors thank Debbie Gruenstein, Lara Rhame, and Alka Srivastava for capable research assistance. Leonardo Bartolini, Ethan Harris, Valerie LaPorte, Nathaniel Leff, Elizabeth Miranda, Rick Mishkin, Frank Packer, Tony Rodrigues, two anonymous referees, and seminar participants at Columbia, CUNY, the Federal Reserve Bank of New York, the IMF, MIT, NYU, and Yale provided valuable comments. We also wish to thank Robert Eggert for help in categorizing the organizations represented in his newsletter, *Blue Chip Economic Indicators*.

Rational Bias in Macroeconomic Forecasts

Abstract

This paper develops a model of macroeconomic forecasting in which the wages firms pay their forecasters are a function of their accuracy as well as the publicity they generate for their employers by being correct. In the resulting Nash equilibrium, forecasters with identical models, information, and incentives nevertheless produce a variety of predictions in order to maximize their expected wages. In the case of heterogeneous incentives, the forecasters whose wages are most closely tied to publicity, as opposed to accuracy, produce the forecasts that deviate most from the consensus.

We find empirical support for our model using a twenty-year panel of real GNP/GDP forecast data from the survey *Blue Chip Economic Indicators*. Although the consensus outperforms virtually every forecaster, many forecasters choose to deviate from it substantially and regularly. Moreover, the extent of this deviation varies by industry in a manner consistent with our model.

(JEL D84,E17,L86)

David Laster
Paul Bennett
In Sun Geoum
Federal Reserve Bank of New York
Research and Market Analysis Group
33 Liberty Street
New York, NY 10045
Tel: (212) 720-6215 (Laster)
e-mail: david.laster@frbny.sprint.com

Rational Bias in Macroeconomic Forecasts

Empirical tests of the rational expectations hypothesis as it applies to professional macroeconomic forecasts generally examine whether predictions of a macroeconomic variable are unbiased and efficient.¹ These analyses assume that because forecasters have strong economic incentives to be accurate, the numbers they produce represent their best estimates. This assumption, however, invites further examination: Is the forecast that best advances the economic goals of an individual forecaster always the same as the “best” forecast in a statistical sense? Zarnowitz and Braun (1993) have documented that group mean (“consensus”) forecasts are more accurate than virtually all individual forecasts. Because consensus forecasts are available publicly on a timely basis, it is unclear why firms continue to produce forecasts that are likely to fall short of the consensus in accuracy. A related puzzle, noted by Lamont (1996), is that some experienced forecasters consistently produce projections that are outliers relative to other professional forecasts. These observations suggest that some forecasters, when making their projections, act according to goals unrelated to the pursuit of accuracy.

In this paper we develop a model in which forecasters’ wages are based on two criteria: their accuracy and their ability to generate publicity for their firms. Accuracy is defined in the usual sense of minimizing expected forecast error. Publicity comes from producing the forecast that outperforms all competing forecasts. In our model, these two goals can conflict -- attracting publicity often compromises accuracy.

¹This is the approach followed by numerous researchers in the empirical rational expectations literature, such as Figlewski and Wachtel (1981) and Keane and Runkle (1990). For a current discussion of the issues, see Croushore (1996).

The model demonstrates that even in the case where all forecasters have identical information and identical incentives forecasters' efforts to maximize their expected wage will lead many of them to deliberately produce biased projections to differentiate their views from the consensus. Thus, in contrast to the standard rational expectations approach, our model is one of "rational bias." The model has an additional implication for the case where the incentives forecasters face vary by industry: Forecasters working in industries that place the highest value on publicity will make predictions that deviate most from the consensus.

After presenting the model, we test some of its implications using a twenty-year panel of forecasts of annual real GDP and GNP growth from *Blue Chip Economic Indicators*, a widely used survey of professional forecasters. Sorting the forecasters into six industry categories, we find that those who work for banks and industrial corporations -- the types of firms that might be expected to make forecast accuracy a top priority -- tend to produce forecasts that are closest to the consensus. Independent forecasters, who stand to benefit most from favorable publicity, tend to produce outlying forecasts.

The plan of this paper is as follows. The first section briefly reviews earlier research and describes the forecasting industry. The next section develops a model of rational bias in forecasting and examines several of its implications. Empirical support for the model is provided in the third section. Section four considers some of the implications and limitations of the model. The final section offers a summary and conclusion.

1. Background

1.1 Previous Work

Research examining whether macroeconomic forecasts are consistent with models of rational expectation has produced mixed results. McNees (1978) finds only limited support for the hypothesis that the forecasts of unemployment, inflation, and real GNP growth by three major econometric forecasting firms are efficient and unbiased. Figlewski and Wachtel (1981) conclude that individual inflation forecasts from the Livingston survey are biased and have serially correlated errors, inconsistent with rational expectations. Questioning the methodology of previous empirical studies of the rational expectations hypothesis, Keane and Runkle (1990) set out to perform more carefully controlled tests. These provide evidence that macroeconomic forecasts from the ASA-NBER survey of professional forecasters incorporate full information and are unbiased, “salvaging the possibility that the rational expectations hypothesis is empirically valid” (p. 730). Bonham and Cohen (1995), however, dispute these results on the grounds that Keane and Runkle’s regressions fail to include key economic variables that were publicly available at the time the forecasts were made. Jeong and Maddala (1996) reject the rational expectations hypothesis for a set of interest rate forecasts from the ASA-NBER survey. In short, a preponderance of statistical evidence calls into question the notion that professional forecasts are rational in the sense of being efficient and unbiased.

Casting further doubt on how closely forecaster behavior conforms to rational expectations models are recent studies that suggest that macroeconomic forecasts are influenced by the incentives forecasters face. Ito (1990) finds that exchange rate forecasts are systematically biased toward scenarios that would benefit the forecaster’s employer. He terms this bias

“wishful expectations,” but leaves unresolved whether it stems from unconscious wish fulfillment or intentional behavior by individuals responding to corporate incentive structures. Lamont (1996) hypothesizes that a forecaster’s willingness to make predictions that deviate from the consensus may vary systematically with his or her level of experience or seniority. Analyzing forecasts of GNP and GDP from an annual *Business Week* survey, Lamont finds that forecasters who have been in the industry longer exhibit a greater willingness to deviate from the consensus. Ehrbeck and Waldmann (1996) consider several models in which forecasters, wishing to signal their competence, adjust their forecasts in response to new information in ways that compromise accuracy. Examining forecast data for U.S. Treasury bill rates, the authors are unable to find evidence consistent with their models of rational bias.

This paper makes two contributions to the literature on forecaster behavior. First, it develops an original model in which identical forecasters deliberately differentiate their predictions, creating the impression that there is a divergence of views when there is in fact none. The model therefore provides an answer to Ehrbeck and Waldmann’s challenge of reconciling forecaster behavior with strategic rationality. Second, the paper uses a new panel of data to demonstrate that the extent to which forecasters deviate from the consensus varies by their industry of employment. Our findings have implications, spelled out below, for the way that forecast data should be used to test the rational expectations hypothesis. Before presenting our model of forecaster behavior, we briefly discuss the job of the professional forecaster.

1.2 Roles of the Professional Forecaster²

Professional macroeconomic forecasters work for a variety of employers such as banks, securities firms, nonfinancial corporations, and consulting firms specializing in econometric modeling or other types of economic analysis. For many forecasters, a key part of the job is to provide analysis internally: to track economic variables, make projections, and share insights with the organization's decision makers. Thus, forecasters contribute to the firm through the accuracy of their forecasts.

Another fundamental role of economic forecasters is to provide marketing for their employers. This role applies to economic consultants, who are themselves the product that they must market, and to forecasters employed by major banks and securities firms, whose main duty is to provide economic advice to the firms' clients. The clients, many of whom receive advice from several sources, seek analysis that is relevant, informative, and entertaining. Through public speaking engagements, magazine articles, television interviews, and quotations in the press, economists keep their employers' names visible before current and prospective clients and cultivate an image of expertise and originality. In this arena, where the audience is of a broad range of backgrounds, the manner in which economists present their analysis may count as much as the content of the analysis. Forecasters who operate successfully in this environment provide publicity for their firms.

1.3 Surveys of Professional Forecasters

One way for forecasters to gain publicity for their firms is to participate in surveys of professional forecasters. These surveys, which appear in the business press and specialized

²For a more extensive discussion, see Henry (1989).

publications, call attention to the firm whose forecasts for the latest period came closest to the actual outcome. *Business Week*, for example, regularly features a collection of macroeconomic forecasts by business economists in its year-end issue. Accompanying the forecasts is a write-up of the economist whose prior-year projection was closest to the mark. The write-up, complete with photograph, also prominently cites the firm for which the economist works. Similarly, the *Wall Street Journal* rewards the most accurate participant in its survey of professional forecasters with a separate article. The publisher of *Blue Chip Economic Indicators*, a monthly newsletter compiling dozens of professional economic forecasts, holds an annual dinner at which the most accurate forecaster for the previous year is honored. The winning forecaster is also identified in later issues of the newsletter.³

To succeed, a survey must offer benefits to all parties involved: the publication, its readers, and the participating firms. The publication gains useful data that will interest existing readers and attract new ones. Readers gain information about the outlook for the economy. Firms that go to the trouble of responding to the survey and offering their predictions free of charge also gain something -- media exposure. This includes the chance to receive the favorable coverage that comes from producing the best forecast in a given period. The high-profile reporting of a firm's forecasting expertise may be even more effective than a paid advertisement in attracting new clients. Because media citation also enhances the reputation of individual forecasters, economists will be particularly willing to help their firms by participating in these

³A recent example illustrates how significant these contests have become. When Lawrence Meyer was nominated to the Board of Governors of the Federal Reserve System in February 1996, newspaper accounts noted that in recent years his economic forecasting firm had twice produced the top forecast in *Blue Chip Economic Indicators*.

surveys.

Media coverage of the forecasters participating in these surveys tends to be favorable. Just as the business and financial press have strong incentives to identify the most successful forecaster, reporters also have strong incentives *not* to call attention to the poorest performers. This asymmetric media coverage appears to be an efficient mechanism for encouraging firms to participate, and to continue participating, in forecasting surveys.

2. The Model

We now develop a model that shows how forecasters' efforts to balance the twin objectives of accuracy and publicity can lead them to produce biased macroeconomic forecasts. Heuristically, if all forecasters have similar information, the pursuit of accuracy alone will lead to forecasts that cluster tightly around a consensus. Forecasters seeking publicity, however, will not want to be among the cluster because their forecasts would then have little or no chance of winning widespread attention. Instead, they will predict values with a reasonable likelihood of occurring but which are not already being forecast by others. As an extension of this reasoning, forecasters who are especially publicity-conscious will be more inclined to make unconventional forecasts; those who emphasize accuracy will make projections that cluster around the most likely outcome. The following sections spell out the details of our model.

2.1 Timing and Information Structure

Assume that there are N firms, each of which employs an economic forecaster. At date 0, the forecasters announce their predictions of the next period's level of a macroeconomic variable

x , whose probability distribution function (pdf) is discrete.⁴ Forecasters have access to two types of information on which to base their predictions of x : (1) the pdf of x , $f(x)$; and (2) the contemporaneous distribution of forecasts made by those in the profession, denoted $n(x)$.

By assuming the existence of a pdf on which forecasters agree, the model seeks to explain the dispersion of forecasts without appealing to differences among forecasters' information sets, methodologies, or abilities. Although some differences of opinion among forecasters are inevitable, one strength of the model is its ability to explain how a dispersion of forecasts can occur even in the absence of these differences.

Our other key assumption is that each individual, when forecasting, is aware of the distribution of contemporaneous forecasts, $n(x)$, a function that meets the condition

$$(1) \quad \sum n(x) = N.$$

The assumption that forecasters know $n(x)$ reflects an industry environment in which forecasters reveal and actively debate their views with one another. The forecasts that appear in the December issue of *Blue Chip Economic Indicators*, for example, are very similar to those appearing in the November issue. Even if important new information has arrived in the interim, forecasters can ascertain from published reports, statements in the press, or personal conversations how others in the profession have adjusted their views. Thus, in the model, the dispersion of published forecasts is due not to differences in information and methodologies, but to the strategic behavior of forecasters reacting to one another.

⁴The discreteness assumption reflects the way in which widely followed macroeconomic variables are reported. For example, the real GDP growth rate and the unemployment rate are rounded to the nearest tenth of a percent even though the level of GDP and the number of people working and looking for work are measured with greater precision.

At date 1, the realized value of the predicted variable, x_0 , is announced. A forecaster's wage is then set based on ex post performance.

2.2 Forecasters' Compensation

A forecaster is paid according to how well he or she fulfills the dual role of internal adviser and source of media attention. As internal adviser, a forecaster producing an accurate forecast helps the firm to plan wisely; a poor forecast will create inefficiencies and missed opportunities. More specifically, we assume that the opportunity cost L of an inaccurate forecast is a function of forecast error:

$$(2) \quad L(x) = L(x_0 - x).$$

L represents the difference between the profits a firm actually realizes from operations and the profit it would have earned had its forecaster been exactly on target.⁵ (This excludes any gains in profitability because of publicity, which we measure separately). The function is assumed to attain its minimum value of zero when its argument is zero.

A forecaster can also contribute to the firm by enhancing its reputation. A firm whose forecaster is the one who correctly predicts the value of x receives favorable publicity worth P . If several forecasters predict this value, however, the publicity must be shared among their firms. More formally, if the realized value of the forecast variable is x_0 , the value of publicity derived from a forecast of x equals

$$(3) \quad A(x) = P/n(x) \quad \text{if } x = x_0; \text{ and}$$

⁵We have chosen this specification to be parsimonious. If accuracy were more important in some states of the economy (e.g., recession) than in others, however, L could be generalized to assume the form $L(x_0-x, x_0)$.

$$A(x) = 0 \quad \text{otherwise.}^6$$

Forecasters are paid for their contribution to their employers, as measured by their accuracy and the advertising they generate. Assuming a linear pay structure, forecaster i earns

$$w_i = v_i - s_i L(x_i) + b_i A(x_i), \quad i = 1, 2, \dots, N$$

where w_i is forecaster i 's wage, v_i is a constant, $s_i \geq 0$ measures how closely forecaster i 's pay is tied to accuracy, and $b_i \geq 0$ reflects the size of the bonus paid to forecaster i for making a prediction that garners publicity by being among the most accurate.

Specifying L to be a quadratic loss function will simplify the analysis considerably.⁷

Thus, the wage of the i^{th} forecaster, who predicted x_i when the variable's realized value is x_0 , equals

$$(4) \quad w_i(x_i) = v_i - s_i(x_0 - x_i)^2 + b_i A(x_i), \quad i = 1, 2, \dots, N.$$

If forecasters are assumed to be risk neutral, their optimization problem reduces to choosing the value of x_i that maximizes their expected wage:

$$E w_i(x_i) = v_i - s_i E(x_0 - x_i)^2 + b_i P f(x_i) / n(x_i), \quad i = 1, 2, \dots, N,$$

where the final term is the expected value of the forecaster's bonus for correctly predicting x .

If we set $E(x_0)$ equal to μ , this expression simplifies to

$$(5) \quad E w_i(x_i) = w_i^* - s_i(\mu - x_i)^2 + b_i P f(x_i) / n(x_i), \quad i = 1, 2, \dots, N,$$

⁶An additional assumption that would make the model more realistic is that if no forecaster correctly predicts the value of x , all those who come closest will equally share the publicity:

$$A(x) = P / [n(x) + n(2x_0 - x)] \quad \text{if } n(x_0) = 0 \text{ and } |x - x_0| \leq |z - x_0| \\ \text{for all } z \text{ for which } n(z) > 0.$$

As a practical matter, this will not alter the equilibrium because, as we observe in footnote 13, every value in the range of possible outcomes will be forecast by someone.

⁷Alternative functional forms will produce very similar results. If, for example, L were a function of absolute error instead of squared error, forecasters' expected wage would be penalized for deviations from the *median* of the pdf as opposed to its mean. The analysis would be essentially the same.

where $w_i^* = v_i - s_i \text{Var}(x)$.⁸ The constant w_i^* , which is beyond the control of forecaster i , can be interpreted as his or her expected base wage, absent bonus, for forecasting the value μ .

The final two terms in equation (5) summarize the trade-off between accuracy and publicity that forecasters face. At one extreme, they can simply forecast μ , the expected value of x . This will minimize their expected squared error, but it would rule out the possibility of earning a large bonus if many others also forecast μ , that is, if $n(\mu)$ is large. Alternatively, to have a chance at winning a large bonus, they can choose a value of x for which $n(x)$ is small. Such a forecast would likely be biased, however, raising their expected squared error. The choice that a given forecaster makes will depend on two factors: differences between the pdf and the distribution of forecasts, as measured by $f(x)/n(x)$, and the relative emphasis his or her employer places on accuracy as opposed to advertising, measured by s_i/b_i . Next we consider the resulting equilibrium.

2.3 Homogeneous Incentives

In the simplest version of the model, every employer places the same emphasis on advertising and accuracy. Because the parameters b_i and s_i are the same for all forecasters, we can omit the subscripts from our discussion. There are three possible cases.

Case I: Only Accuracy Matters ($b=0$ and $s>0$)

When only accuracy matters to employers, it follows from equation (5) that, because $b=0$, the expected wage is maximized when $(\mu-x)^2$ is at a minimum. The optimal forecast of x will

⁸The key step is to note that
$$E(x_0 - x_i)^2 = E[(x_0 - \mu) + (\mu - x_i)]^2 = E(x_0 - \mu)^2 + 2(\mu - x_i)E(x_0 - \mu) + (\mu - x_i)^2.$$

The middle term is zero by definition and the first term is the variance of x . Substituting,
$$Ew_i(x_i) = v_i - s_i \text{Var}(x) - s_i(\mu - x_i)^2 + b_i Pf(x_i)/n(x_i),$$
 which is equivalent to (5).

therefore be its expected value μ . If all employers compensate their forecasters based solely on accuracy, everyone will make the same forecast.

Case II: Only Publicity Matters ($b>0$ and $s=0$)

In the opposite case, forecasters are rewarded exclusively for the publicity they generate for their employers. Since $s=0$, expression (5) implies that each forecaster will choose the value of x for which $f(x)/n(x)$ is maximized. When all forecasters try to maximize $f(x)/n(x)$, the resulting distribution of forecasts $n(x)$ will be exactly proportional to $f(x)$, the pdf of x .

This result can be established through a proof by contradiction. Suppose that the ratio $f(x)/n(x)$ were not equal to a constant in equilibrium. Then there would exist values for x_1 and x_2 such that $f(x_1)/n(x_1) > f(x_2)/n(x_2)$. This outcome would not constitute an equilibrium. People forecasting x_2 , commanding a lower expected wage than those forecasting x_1 , would change their forecasts from x_2 to x_1 . The migration would raise $n(x_1)$ and lower $n(x_2)$, until the discrepancy is eliminated.⁹

Together, this general condition of proportionality and equation (1) imply that when publicity is all that matters,

$$(6) \quad n(x) = Nf(x) \quad \text{for all } x.$$

What is striking about this result is that even though everyone agrees which value of x is most likely to occur, forecasters will nonetheless be drawn to distribute themselves in a fashion that mimics the pdf of x .

Case III: Accuracy and Publicity Both Matter ($b>0$ and $s>0$)

⁹This argument implicitly assumes that $n(x)$ can take on non-integer values. For further elaboration, see the next footnote.

Having examined the two extreme cases, we next consider the intermediate case in which publicity and accuracy both matter. Expression (5) states expected wage Ew as a function of the forecasted value of x . For a distribution of forecasts $n(x)$ to constitute a Nash equilibrium, it must have the property that no forecaster will be able to increase his or her expected wage through a change of forecast. What this latter condition implies for equilibrium depends on whether or not $n(x)$ is assumed to take on only integer values. Provided that $n(x)$ can assume any real value, a necessary condition for equilibrium is that all forecasts must yield the same expected wage \hat{w} :

$$(7) \quad Ew(x) = \hat{w} \quad \forall x \ni n(x) > 0.^{10}$$

To simplify the analysis, we assume that the mean of $f(x)$ is one of the discrete values that x can assume. Reparameterizing so that $\mu_x = 0$ and substituting (7) into (5) yields

$$\hat{w} = w^* - sx^2 + bPf(x)/n(x),$$

which can be solved for n :

$$(8) \quad n(x) = \frac{Pf(x)}{k(\hat{w}) + rx^2}, \quad \text{where } k(\hat{w}) = (\hat{w} - w^*)/b \\ \text{and } r = s/b.$$

The expression $n(x)$ represents a distribution of forecasts that will cause all forecasters to have the same expected wage, \hat{w} . We can think of this function as the forecasters' iso-expected wage curve or, more simply, their iso-wage curve. For this expression to constitute an

¹⁰If, however, $n(x)$ is restricted to integer values, there will be a somewhat weaker condition that $n(x)$ must meet to constitute a Nash equilibrium, namely, that:

(7') under any distribution $n'(x)$ for which $n'(x_1) = n(x_1) + 1$ and $n'(x_2) = n(x_2) - 1$, where x_1 and x_2 are possible forecasts,

$$Ew(x_1) \leq Ew(x_2).$$

In this paper we assume that n can assume non-integer values, so that equation (7) applies. The larger the pool of forecasters, the better an approximation this is. Our analysis of the case in which n is restricted to integer values, not reported here, demonstrates that the main results continue to hold and that the resulting equilibrium is essentially the same.

equilibrium, it must also satisfy equation (1), which states that there are exactly N forecasters.

There exists a unique market-clearing wage \hat{w} for which equation (8) characterizes an equilibrium with N forecasters. To understand why, consider the terms on the right-hand side of equation (8). The pdf $f(x)$ is known and the parameters P (the value of publicity) and r (the relative emphasis that employers place on accuracy) have fixed values. The only variable in the expression is k , which is a positive linear function of the market expected wage \hat{w} . Thus, to each value of \hat{w} , there corresponds a distinct iso-wage curve, which we can label $n^{\hat{w}}(x)$ (Chart 1). Because P , $f(x)$, and r are all nonnegative, it follows from (8) that $n^{\hat{w}}(x)$ is monotonically decreasing in \hat{w} .¹¹ Moreover, by varying \hat{w} , we can make $\sum n^{\hat{w}}(x)$ arbitrarily large or small.¹² This implies that for any given number of forecasters N , there will exist a market-clearing expected wage \hat{w} and a unique equilibrium distribution $n(x)$ that satisfies equation (1).¹³

The equilibrium will be stable because forecasters predicting a value of x that is selected by more than $n(x)$ forecasters will earn a substandard expected wage, prompting them to change their forecasts. Conversely, a value of x chosen by fewer than $n(x)$ forecasters will command an expected wage above the industry standard. These forces will create incentives for forecasters to change their predictions until equation (8) holds.

2.3.1 Comparative Statics

How does the relative emphasis that employers place on accuracy as opposed to publicity

¹¹Intuitively, an increase in the reservation wage of forecasters will induce some to leave the business, raising the expected wage of those who remain (because there will be fewer forecasters with whom to share the available publicity).

¹²This is because $\forall x, n^{\hat{w}}(x) \rightarrow 0$ as $\hat{w} \rightarrow \infty$; and $n^{\hat{w}}(0) \rightarrow \infty$ as $\hat{w} \rightarrow w^*$.

¹³Expression (8) also implies that for all values of x for which $f(x)$ is positive, $n(x)$ will also be positive. In other words, there will always be someone forecasting every possible value of x .

affect the equilibrium distribution of forecasts? Cases I and II show that, in the extreme, the more employers reward accuracy, the greater the tendency for forecasters to herd. We would therefore expect that, for intermediate cases as well, the greater the value of $r=s/b$, the stronger the tendency to herd. Proposition 1 makes this relationship more precise.

PROPOSITION 1: *If $n_1(x)$ and $n_2(x)$ each represent an equilibrium distribution of N forecasters such that n_1 reflects a greater relative emphasis on accuracy than n_2 (that is, $r_1 > r_2$), then there exists a positive constant, a , so that $n_1(x) > n_2(x) \Leftrightarrow |x| < a$.*

This proposition states that an increased emphasis on accuracy will raise the number of forecasts of values within a symmetric interval around zero and will decrease the number of forecasts for values outside that interval (Chart 2).

To establish the proposition, we first note that (8) implies that

$$(8a) \quad n_i(x) = \frac{Pf(x)}{k_i(\hat{w}_i) + r_i x^2}, \quad i=1,2$$

where subscripts denote the two alternative distributions. It follows from the hypothesis $r_1 > r_2$ and (8a) that if $k_1 \geq k_2$, then $n_1(x) \leq n_2(x)$ for all x , with strict inequality holding for all nonzero values of x . This would imply that $\sum n_1(x) < \sum n_2(x)$, which violates the assumption that the two alternative distributions have the same number of forecasters. So k_1 must be less than k_2 . By substituting the values $i=1,2$ into (8a) and simplifying, we arrive at

$$n_1(x) > n_2(x) \Leftrightarrow x^2 < (k_2 - k_1)/(r_1 - r_2).$$

The desired result follows from setting $a = [(k_2 - k_1)/(r_1 - r_2)]^{.5}$.

Equation (8) and Proposition 1 demonstrate that the distribution of forecasts we observe for a given period reflects two underlying factors -- the pdf of the variable being forecast and the relative emphasis that employers place on accuracy. If the incentives forecasters face vary little

from year to year, then we can interpret a change in the extent to which forecasters cluster as an indication that the variable's (unobserved) pdf has changed. This implication accords with the empirical finding of Zarnowitz and Lambros (1987, p. 607) that “greater interpersonal differentiation of expectation is a symptom of greater uncertainty.” Conversely, secular trends in the incentives that forecasters face, as measured by r , will affect the equilibrium distribution of forecasts, even absent any changes in the forecast variable's pdf.

Proposition 1 also implies that the more employers emphasize accuracy over publicity, the lower the variance of forecasts.¹⁴ But recall from our discussion of Case II that when accuracy receives a zero weight, the distribution of forecasters will be identical to the pdf of x . These two statements together imply a

COROLLARY: *In equilibrium, the forecasts of a given variable will have a variance that is less than or equal to that of the variable itself.*

2.3.2 Consensus Forecast

In this model, the predictions of individual forecasters are often biased. What about the consensus forecast? Suppose that the pdf of x is symmetric about its mean, which is labeled zero, so that

$$f(-x) = f(x) \quad \text{for all } x.$$

Together with expression (8), this condition implies that $n(x)$, the distribution of forecasts, will also be symmetric about zero. From this symmetry it follows that the distribution of individual forecasts of the variable will have the same mean as the variable's pdf, a result worth emphasizing.

¹⁴This result follows from the property that, for all c , the sum $\sum (n_1(x) - n_2(x))$, when calculated for $-c \leq x \leq c$, is nonnegative. For further discussion see Mood, Graybill, and Boes (1974), pp. 74-5.

PROPOSITION 2: *If a variable has a probability distribution function symmetric about its mean, then the consensus (mean) forecast of the variable will be unbiased.*

Proposition 2 is premised on a very strong assumption -- perfect symmetry. More generally, it will hold approximately true to the extent that $f(x)$ approximates a symmetric function. The key insight is that when a large proportion of forecasters opt to make high forecasts, there will be a strong incentive for some to switch to low forecasts (and vice versa) to increase their expected bonuses. This incentive will cause the consensus forecast to be unbiased, or nearly so. The unbiasedness of the consensus has an important empirical implication, namely, that the consensus forecast should have the minimum expected mean squared error. We will return to this point in our empirical work below.

To summarize, we have shown that in the case where all employers pay their forecasters according to a common wage function, a stable, unique equilibrium distribution of forecasts exists. The variance of this distribution is less than that of the pdf, is positively correlated with the relative emphasis employers place on publicity, and reflects the objective uncertainty of the variable being forecast. If the pdf is symmetric, the consensus forecast will be unbiased.

2.4 Heterogeneous Incentives

Now consider the more general case in which the wage parameters w_i^* , s_i , and b_i can vary by industry. Firms that hire forecasters to function chiefly as internal advisors have a high value of s_i relative to b_i because accuracy is what matters most. Other employers place special emphasis on publicity and assign a much higher value to b_i relative to s_i .

The equilibrium we consider is one in which all forecasters receive the same expected wage. Expected wage is equalized *within* each industry because any wage inequality will

motivate forecasters to change their forecasts until the disparity is eliminated. Expected wage will also be equalized *across* industries because forecasters will migrate from low-wage industries to high-wage industries until interindustry discrepancies disappear.

Suppose that the pdf $f(x)$, the value of publicity P , the market-clearing expected wage \hat{w} , and the wage vector $\{w_i^*, s_i, b_i\}$ are known for each of m industries. Will an equilibrium distribution of forecasts exist, and, if so, how many forecasters will choose to work in each industry and what values will they forecast?

To solve for the equilibrium, we follow the same approach as in the one-industry homogeneous incentives case. The iso-wage curve $n(x)$ in equation (8) states how many forecasters employed in a given industry will be prepared to forecast each value of x . Appending subscripts, we can use an analogous expression in the multi-industry case:

$$(8b) \quad n_i(x) = \frac{Pf(x)}{k_i(\hat{w}) + r_i x^2}, \quad i=1, \dots, m,$$

where $k_i(\hat{w}) = (\hat{w} - w_i^*)/b_i$
and $r_i = s_i/b_i$.

For any given market-clearing expected wage \hat{w} , we can aggregate the industry-specific iso-wage curves $n_i(x)$, to determine the equilibrium distribution of all forecasters. For each possible forecast value x , the expression $n_i(x)$ states how many forecasters compensated according to the industry i pay scale can predict x and still earn \hat{w} . It follows that, in equilibrium the only forecasters who will predict the value x will be those from the industry or industries for which $n_i(x)$ is a maximum. Forecasters from other industries will be discouraged from predicting the value x because doing so would cause them to earn an expected wage below the market rate \hat{w} . The aggregate distribution of forecasters $n(x)$ will therefore be the upper envelope of the

individual industry iso-wage curves (Chart 3a):

$$(9) \quad n(x) = \underset{i}{\text{Max}} n_i(x) \quad i=1, \dots, m.$$

Equations (8b) and (9) together determine the equilibrium distribution of forecasters for the multi-industry case. The salient characteristic of this equilibrium is the close connection between the way an industry pays its forecasters and the forecasts that they generate. This is expressed in

PROPOSITION 3: *If forecaster 1 is rewarded relatively more for accuracy than forecaster 2 ($r_1 > r_2$), then their forecasts x_1 and x_2 will be such that $|x_1| \leq |x_2|$.*

The proof of this proposition follows directly from the optimizing behavior of forecasters.

The fact that forecaster 1 chooses value x_1 rather than x_2 implies that

$$Ew_1(x_1) \geq Ew_1(x_2).$$

Substituting into (5) and noting that $\mu=0$ yields

$$w_1^* - s_1 x_1^2 + b_1 P f(x_1)/n(x_1) \geq w_1^* - s_1 x_2^2 + b_1 P f(x_2)/n(x_2),$$

which simplifies to

$$(10) \quad P[f(x_1)/n(x_1) - f(x_2)/n(x_2)] \geq r_1(x_1^2 - x_2^2).$$

Similarly, forecaster 2's preference for x_2 over x_1 means that

$$Ew_2(x_2) \geq Ew_2(x_1), \quad \text{or}$$

$$(11) \quad r_2(x_1^2 - x_2^2) \geq P[f(x_1)/n(x_1) - f(x_2)/n(x_2)].$$

Combining (10) and (11) and subtracting gives

$$(r_2 - r_1)(x_1^2 - x_2^2) \geq 0.$$

The first of these two multiplicative terms is negative by hypothesis. Thus, $x_1^2 - x_2^2$ must be negative or zero. The proposition follows as a consequence.

Proposition 3 allows us to characterize the equilibrium distribution for forecasters. If the employers in each of m industries emphasize accuracy and publicity to varying degrees, we can number the industries so that

$$r_i > r_j \quad \text{for } i < j.$$

The proposition implies that forecasters employed in industry 1, which places the greatest weight on accuracy, will position themselves along a symmetric interval around zero, $[-c_1, +c_1]$.

Industry 2 forecasters will select values over the intervals $[-c_2, -c_1]$ and $[+c_1, -c_2]$, and so forth (Chart 3b). Finally, industry m forecasters, who receive the greatest relative rewards for attracting publicity, will position themselves along the tails of the probability distribution, at $(-\infty, -c_{m-1}]$ and $[+c_{m-1}, \infty)$.¹⁵

We can frame this result in a more realistic way. Instead of taking the industry wage parameters $\{w_i^*, s_i, b_i\}$ and the market-clearing wage \hat{w} to be exogenous, assume that each industry adjusts its pay scale to the level needed to attract a desired number of forecasters. Thus, for each industry, the parameter r_i is fixed while k_i can vary. If, in equilibrium, each industry i chooses to employ N_i forecasters, then Proposition 3 suggests an algorithm for determining the equilibrium distribution of forecasts as a function of $f(x)$, industry parameters r_i , and the desired number of forecasters in each industry, N_i (Appendix 1).

¹⁵There could, of course, be two or more industries for which r assumes the same value. In that case, it follows from (8b) that the one with the lowest value of k will bid away the forecasters from the other industries having the same value of r . If two industries have identical values for both k and r , they can be thought of as constituting a single industry for the purposes of the model.

This discussion implicitly assumes that the wage parameters for the m industries are such that some forecasters will be employed in each industry. If, however, the parameter k_j for a particular industry j is too high (that is, w_j^* is too low), forecasters might choose not to work in that industry. Stated in terms of expression (9), the uncompetitive pay package that industry j offers means that there does not exist a value of x for which $n_i(x)$ is at a maximum when $i=j$. This in no way affects our discussion. We can simply exclude uncompetitive industries such as j from the analysis.

To conclude, the multi-industry case of our model has a very strong empirical implication: the more an industry rewards its forecasters for generating publicity, the greater will be the tendency for forecasters in that industry to produce unconventional forecasts.

3. Empirical Results

In this section, we report statistical results that are generally consistent with the implications of the model. Our data consist of year-ahead forecasts of annual U.S. real GDP (before 1992, GNP) growth published in *Blue Chip Economic Indicators*. Participating forecasters were categorized by industry in consultation with Robert Eggert, the newsletter's editor. These industry categories are an objective way of grouping together forecasters who work for similar types of firms and can therefore be expected to face similar incentives. Before presenting our empirical results, we offer a few details about the data.

3.1 Data

Blue Chip Economic Indicators is a monthly newsletter that compiles several dozen professional forecasts of widely followed macroeconomic variables. The forecasts are produced by a variety of participating firms, some of which are designated as members of the Blue Chip panel. Our data set includes all forecasters, regardless of whether or not they are panel members. Once a firm is invited to participate in the survey, it remains a participant as long as it continues to submit forecasts. Although firms are not paid to participate, the newsletter offers them regular public exposure.

Our data consist of the forecasts of year-ahead real GNP/GDP growth appearing in the December issues of the newsletter. During the sample period, 1976-95, the number of forecasters in any one year ranged from thirty-two to eighty-one. A total of 129 firms contributed

real GNP/GDP forecasts at one time or another. We divide these firms, listed in Appendix 2, into six industry categories -- banks, econometric modelers, independent forecasters, industrial corporations, securities firms, and other. When analyzing how forecaster behavior is related to industry of employment, we use the full panel of data, which consists of 1,197 individual forecasts. When analyzing the forecasts of individual firms, however, we restrict the sample to the forty-one firms with twelve or more annual real GDP forecasts. Of these, we use the thirty-eight firms that published twelve or more year-ahead December forecasts in the years before 1995 when comparing forecasts to actuals. The industry breakdown of our sample is as follows:

<u>obs)</u>	<u>Category</u>	<u># in survey</u>	<u># in subsample (>12 annual</u>
	Banks	30	12
	Econometric modelers	12	5
	Independent forecasters	38	9
	Industrial corporations	18	4
	Securities firms	14	4
	<u>Others</u>	<u>17</u>	<u>7</u>
	Total	129	41

Our analysis also uses the *Blue Chip's* consensus forecast of real GDP growth, which is reported in each issue. The consensus is calculated as the mean prediction of forecasters on the Blue Chip panel, rounded to the nearest tenth of a percent. The Blue Chip consensus never differs by more than 0.1 percent from the mean of our more inclusive sample of forecasters.

The "actual" figures with which we compare the forecasts are the official figures released in the January following the year in question. For example, the year-over-year actual for 1986 was the Commerce Department's January 1987 measure of 1986 constant-dollar GNP, expressed as a percentage change from the January 1987 measure of 1985 constant-dollar GNP, as reported in the *Survey of Current Business*. While these figures continue to be revised for several years

after their initial release, we use the first-released figure. McNees (1989) convincingly argues that this is an appropriate benchmark against which to measure forecast accuracy.

3.2 Consensus Forecasts

Researchers have found that consensus forecasts are substantially more accurate than individual forecasts. Zarnowitz and Braun (1993, Table 1.9), for example, compare the accuracy of individual and consensus forecasts of real GNP in the NBER-ASA quarterly surveys for the 1968-90 period. They find that, over horizons ranging from one to five quarters, the consensus forecast has a root mean squared error (RMSE) 23 to 27 percent below that of individual forecasts.

Table 1, which ranks the individual forecasters in our sample and the Blue Chip consensus by RMSE, shows a very similar result. The first two columns correspond to subperiods 1977-86 and 1987-95, and the third column corresponds to the entire sample period. Of thirty-eight firms, four beat the consensus in the first subperiod, as did ten in the second subperiod. Not one of the firms that outperformed the consensus in the first subperiod managed to do so again in the second. Over the entire period, the consensus outperformed all but one firm and had RMSE 22 percent below that of the average forecaster.

The superior performance of the consensus is consistent with our model. As noted in the discussion of Proposition 2, the model suggests that the consensus forecast will be unbiased, or very nearly so. Many individual forecasters, by contrast, will find it in their interest to make biased forecasts. Our model, therefore, predicts that the consensus will outperform individual forecasters. There are, of course, other, established explanations of why the consensus does so

well. For example, if forecasters have heterogeneous information sets or models, then averaging their projections will tend to cancel out their idiosyncratic errors. But explanations such as this pose a problem. If forecasters are behaving rationally and the consensus tends to perform best over time, why don't they simply forecast values close to the consensus?

Ehrbeck and Waldmann (1996, p. 37) propose an answer to this puzzle based on psychology: forecasters put too little weight on the consensus because they are genuinely overconfident. Our model offers an alternative explanation -- that forecasters behave strategically in the pursuit of publicity. The next section provides some evidence that, while necessarily indirect, points to the importance of strategic behavior.

3.3 Deviations from Consensus

Chart 4 shows the frequency distribution of the deviations of individual GDP (GNP) forecasts from the consensus. Are these deviations the result of strategic behavior or do they mainly reflect idiosyncratic differences in forecast methodologies and information?

One type of bias that is *not* borne out by our data relates to Ito's wishful expectations hypothesis, which posits that forecasters working in an industry tend to predict outcomes favorable to that industry. We do not find that forecasters' *mean* deviations from the consensus -- allowing positive and negative values to cancel -- varies systematically by industry. In Table 2 we regress the deviation of individual forecasts from the consensus on industry dummies. The OLS results, listed in the left-hand column of the table, show that five of the six industry categories have mean deviations from consensus that are less than a tenth of a percent and are statistically insignificant. The only industry category in which forecasters deviated significantly from the consensus was "other," a catch-all grouping consisting of financial publications,

government agencies, industry associations, insurance companies, and ratings agencies. When this broad grouping is broken down into its component subcategories, only the financial publication subcategory exhibited a significant mean deviation from the consensus (right column). All told, industry affiliation explains less than one percent of the variation in forecasters' deviations from consensus. Using an F-test, we cannot reject the null hypothesis that industry affiliation has no bearing on these deviations. In short, we find little support for the hypothesis that the GNP/GDP forecasts by the firms in our sample exhibit a consistent industry-specific bias.¹⁶

3.4 Absolute Deviations from Consensus

We next test the implications of our model more formally. The model suggests two empirical hypotheses: First, if the importance attached to accuracy relative to publicity varies across industries but is similar within industries, forecasters' mean *absolute* deviations (MAD) from the consensus should be related to their industry. Second, we can make informed a priori guesses about which industries will tend to emphasize publicity most and therefore produce forecasts that deviate most from the consensus.

Nonfinancial corporations, which use forecasts mainly for internal planning purposes, emphasize accuracy. At the other extreme, a consulting firm or advisory service seeking recognition would find the media attention invaluable and would place less emphasis on forecast accuracy in the traditional sense. Banks and brokerages occupy an intermediate position: they

¹⁶Our failure to detect significant differences across industries in mean forecasts should not be construed as a rejection of the wishful expectations hypothesis for our data set. The negative result we obtain could partially be due to the way in which we classify firms. Whereas the profitability of Ito's importers and exporters is clearly tied to the dollar/yen exchange rate, our data set has no analogous partition between firms that benefit more or less from strong real GDP growth.

welcome favorable publicity as a way of attracting clients, but must still maintain their credibility. Econometric forecasting firms need to emphasize accuracy, but are also under pressure to outshine their competitors. Thus, we hypothesize that the mean absolute deviations from the consensus will be relatively small for industrial corporations, large for independent forecasters, and somewhere in between for banks, securities firms, and econometric modelers.

Table 3 reports regressions that measure the MAD of forecasts produced by different categories of firms. The dependent variable is the absolute value of each forecast's deviation from the consensus. Because the equation includes a constant term, the coefficient for each industry dummy indicates whether its MAD is greater or smaller than that of the overall sample. The significance of four industry dummies supports the hypothesis that MADs differ across industry groups. In particular, industrial corporations and banks deviate least from consensus, followed by econometric modelers. The securities firms' MAD was appreciably greater. Independent forecasters had the largest MAD. Summing the constant and industry dummies, we see just how compelling these differences are. The MAD for independent forecasters is 0.82, more than double that for industrial corporations (0.37). This result is consistent with Lamont's (1996) finding that forecasters with firms bearing their own names tend to make unconventional forecasts. Overall, industry affiliation alone explains fully 10 percent of the variation in absolute deviations from the consensus. An F-test strongly rejects the null hypothesis that these affiliations are unrelated to absolute deviations. When year dummies are included to control for intertemporal changes in the distribution of GDP/GNP forecasts (second column), the results are essentially the same.

4. Implications and Extensions

4.1 Rational Expectations in the Aggregate

To place our model and empirical evidence in context, it is useful to note the distinction that Haltiwanger and Waldman (1989) draw between the idea of rational expectations as originally conceived and how it is now used. Muth's seminal article (1961, p. 316) describes a notion of rational expectations in the aggregate: "expectations of firms...tend to be distributed, for the same information set, about...the 'objective' probability distribution of outcomes." Later researchers, however, have taken rational expectations to mean that the expectations of each *individual* is consistent with the objective pdf of outcomes.

Our model has distinct empirical implications for these two versions of the rational expectations hypothesis as applied to macroeconomic forecasts. It provides a new rationale for why forecasts should exhibit rational expectations in the aggregate. The distribution of forecasts represents a competitive equilibrium. If it were substantially biased, some forecasters would deem it profitable to change their forecasts in response, mitigating this bias. The reason why the consensus should exhibit little or no bias therefore reflects more than just a tendency for a large number of idiosyncratic forecast errors to cancel one another out -- it reflects market forces.

The implication for individual forecasters is different. The model suggests that one should not rely exclusively on the forecast of any individual, no matter how capable, because such a forecast might be skewed in response to economic incentives. Thus, our model provides an argument for why macroeconomic forecasts should obey rational expectations in the aggregate but not at the individual level.

4.2 Testing Forecasts for Rationality

This distinction could help explain why Figlewski and Wachtel (1981) found the inflation

forecasts of individuals to be biased and inefficient. Arguing that the common practice of testing the mean forecast for rationality fails to exploit relevant information, they analyze forecasts at the individual level. Yet our model suggests that each individual's forecast is a function of the forecasts made by others. Empirical tests of an individual's forecasts, which do not take this interaction into account, would therefore be biased toward rejecting rationality. Thus, our model provides an argument in favor of testing the consensus forecast, and not individual forecasts, for rationality.

Keane and Runkle (1990) have refined the approach of testing individual forecasts for rational expectations by making GLS adjustments to account for cross-sectional correlations in forecast error caused by unexpected macroeconomic shocks. To test whether these forecasts are efficient, they follow the standard approach of regressing forecast error on a vector of explanatory variables in the forecasters' information set. A significant regression coefficient for any of these variables would suggest that forecasters are not fully incorporating that variable in their projections, that is, that their forecasts are inefficient. One way our model might be tested using this procedure is to include among the explanatory variables a measure of each forecaster's deviation from the previous period consensus forecast. Our model, which predicts that forecasters deliberately choose to deviate from the consensus at the expense of accuracy, suggests that a forecast's deviation from the lagged consensus should be positively correlated with forecast error. Those who forecast above the consensus should, on average, be too high; those below, too low.

Ehrbeck and Waldmann (1996, Tables IV-V) perform just such a test for interest rate forecasts, with corroborating results. They find that a forecaster's deviation from previous month

consensus forecasts exhibits a large, significant positive correlation with forecast error. Their regressions imply that a 10 basis point increase in a forecaster's predicted T-bill rate relative to last month's consensus forecast will raise his or her forecast error by anywhere from 5.6 to 7.9 basis points.

Ehrbeck and Waldmann argue that this result is very difficult to reconcile with a model of strategic behavior. This is because forecasters who deviate from the consensus hurt themselves twice: first, by increasing their expected forecast error, and second, by creating a pattern of forecasts typical of less capable forecasters.¹⁷ Our model provides an answer to their dilemma. The puzzle disappears once we allow for the fact that a forecaster builds a reputation by making predictions that prove accurate, rather than just by adopting a forecasting style that mimics accurate forecasters.

4.3 Extensions

The model can be extended to capture several additional aspects of the forecasting industry, which we now briefly consider.

4.3.1 Heterogeneous Expectations

Our model demonstrates that the dispersion of professional forecasts that one observes can occur even in the absence of differences of information or views. Nonetheless, because such differences are inevitable, it is worth considering how they might change the results.

When forecasters hold differing views of variable x 's pdf, they have less need to behave

¹⁷Ehrbeck and Waldmann document that forecasters who tend to deviate most from the previous month's consensus forecast have significantly higher mean squared forecast error than their peers (1996, Table VI). In our study, a comparison of column 3 in Table 1 and Table 3 reveals a similar phenomenon: independent forecasters deviate most from the consensus and have the highest root mean squared errors.

strategically. If, for example, everyone agrees on a common pdf with a mean value of 2 percent, one forecaster might opt to predict 3 percent while another forecasts 1 percent. When, however, these same two forecasters have a slight difference of opinion -- one believes that the pdf peaks at 1.8 while the other holds that this occurs at 2.2 -- they will not have to shade their forecasts as much as before. Although strategic behavior would still contribute to the dispersion of forecasts, this dispersion would also reflect genuine differences of opinion. The greater these difference of opinion are, the less need there will be for forecasters to behave strategically, that is, to scatter. If we expand our model to allow both differences of opinion and strategic behavior, publicity-seeking forecasters will continue to scatter but accuracy-seeking forecasters might even herd (that is, strategically position themselves close to the mean).¹⁸

4.3.2 Specification of Industries

Our empirical analysis assumes that all firms in an industry assign similar weights to publicity and accuracy. These relative weights might also differ from firm to firm. Because nothing in our model limits the number of industries, however, we could in principle treat each firm as its own industry. Proposition 3 would therefore imply that each forecaster's tendency to herd exactly corresponds to the priorities of his or her employer.

4.3.3 Risk Aversion

Differences in attitudes toward risk might also help explain the observed dispersion of forecasters. Risk-averse forecasters who are willing to accept a lower expected wage in

¹⁸The net impact of herding behavior depends on its underlying cause. Thus, Scharfstein and Stein (1990) and Zwiebel (1995) are careful to distinguish between two motives for herding. If forecasters make appropriate Bayesian adjustments to their expectations based on the views expressed by others, their accuracy would improve. If, however, they mimic the consensus just to signal that they are well informed, their accuracy would be reduced.

exchange for reduced wage uncertainty would tend to make conventional forecasts. Forecasters who are less risk averse would make less conventional predictions, for which they would be compensated with higher expected wages. The relative wages of forecasters who are more or less risk averse would depend on the supply of such forecasters as well as the relative market demand for accuracy-oriented forecasters and advertising-oriented forecasters.

4.3.4 The Rolodex Hypothesis¹⁹

An alternative hypothesis, similar in spirit to our model, is that forecasters gain publicity by articulating extreme points of view whether or not they ultimately prove correct. Decision makers seeking a range of views might want to hear extreme forecasts and commentary to help them understand the potential upside and downside risks to the consensus forecast. Some forecasters may therefore choose to distinguish themselves by cultivating a reputation as a perennial bull or bear, causing journalists and business people to flip their Rolodexes to these forecasters' cards. This phenomenon -- the "Rolodex hypothesis" -- differs from our model in two basic ways. In our model, the payoff for being unconventional is endogenous and presumes a need to be correct, at least occasionally. The Rolodex hypothesis asserts that flouting convention is its own reward. Moreover, if this hypothesis is correct, we would expect forecasters who regularly deviate from the consensus to be either consistently high or consistently low. Our model, because of its one-period nature, makes no prediction about this.

5. Summary and Conclusion

This paper has developed a theory of rational bias in macroeconomic forecasting in which forecasters have an incentive to compromise the accuracy of their projections to gain publicity

¹⁹Our thanks to Ethan Harris for suggesting this idea.

for their firms. The theory rests on two key assumptions. First, that forecasters are fully knowledgeable about the true probability distribution of actual outcomes. Second, that individual firms value the publicity given to the forecaster who is most accurate in a given year. Our model predicts that rational forecasts are distributed in a way that reflects the true probability distribution of the variable being forecast, with the degree of clustering dependent on the variance of the true pdf as well as the relative emphasis employers place on accuracy. It shows how firms with identical information sets and models may nonetheless produce different forecasts. While we do not wish to deny that genuine differences of opinion exist, our analysis explains how a dispersion of forecasts can occur even without such differences. Moreover, our empirical finding that forecasters' deviations from the consensus are related to the types of firms for which they work provides evidence that professional forecasting has a strong strategic component.

Two refinements of this research -- one theoretical and one empirical -- seem promising. The first of these, outlined above, is to extend the model to allow for differences of opinion among forecasters. The second is to analyze forecaster behavior for a variety of financial and economic variables to determine whether the incentives for accuracy and publicity vary according to the measure being forecast. We hypothesize that forecasts of variables that have direct implications for investments, such as the long bond rate and exchange rates, will exhibit a greater emphasis on accuracy than will forecasts of other variables such as GDP.

This paper provides an approach that reconciles forecaster behavior with strategic rationality. The analysis also lends support to economists who caution against using published forecast data as a proxy for individual expectations. Yet the mechanism we identify also

explains why, despite this concern, the consensus forecast should accurately reflect the forecasting profession's collective insight.

References

- Bonham, Carl, and Richard Cohen. 1995. "Testing the Rationality of Price Forecasts: Comment." American Economic Review 85, no. 1 (March): 284-9.
- Croushore, Dean. 1996. "Inflation Forecasts: How Good Are They?" Federal Reserve Bank of Philadelphia Business Review, May-June: 15-25.
- Diebold, Francis X., and Jose A. Lopez. 1995. "Forecast Evaluation and Combination." Federal Reserve Bank of New York Research Paper no. 9525.
- Ehrbeck, Tilman, and Robert Waldmann. 1996. "Why Are Professional Forecasters Biased? Agency versus Behavioral Explanations." Quarterly Journal of Economics, February: 21-40.
- Figlewski, Stephen, and Paul Wachtel. 1981. "The Formation of Inflationary Expectations." Review of Economics and Statistics 63, no. 1 (February): 1-10.
- Haltiwanger, John C., and Michael Waldman. 1989. "Rational Expectations in the Aggregate." Economic Inquiry, October: 619-36.
- Henry, George B. 1989. "Wall Street Economists: Are They Worth Their Salt?" Business Economics, October: 44-8.
- Ito, Takatoshi. 1990. "Foreign Exchange Rate Expectations: Micro Survey Data." American Economic Review 80, no. 3 (June): 434-49.
- Jeong, Jinook, and G.S. Maddala. 1996. "Testing the Rationality of Survey Data Using the Weighted Double-Bootstrapped Method of Moments." Review of Economics and Statistics, May: 296-302.
- Keane, Michael P., and David E. Runkle. 1990. "Testing the Rationality of Price Forecasts: New Evidence from Panel Data." American Economic Review 80, no. 4 (September): 714-35.
- _____. 1995. "Testing the Rationality of Price Forecasts: Reply." American Economic Review 85, no. 1 (March): 289.
- Lamont, Owen. 1996. "Macroeconomic Forecasts and Microeconomic Forecasters." Paper presented at Federal Reserve Bank of Philadelphia/University of Pennsylvania Conference on Expectations in Economics: In Honor of the Fiftieth Anniversary of the Livingston Survey, October 3-4.
- McNees, Stephen K. 1978. "The 'Rationality' of Economic Forecasts." American Economic Review 68, no. 2 (May): 301-5.

_____. 1989. "Forecasts and Actuals: The Trade-Off between Timeliness and Accuracy." International Journal of Forecasting 5: 409-16.

Mood, Alexander M., Franklin A. Graybill, and Duane C. Boes. 1974. Introduction to the Theory of Statistics. 3d ed. New York: McGraw-Hill.

Muth, John F. 1961. "Rational Expectations and the Theory of Price Movements." Econometrica 29, no. 3 (July): 315-35.

Scharfstein, David S, and Jeremy C. Stein. 1990. "Herd Behavior and Investment." American Economic Review 80, no. 3 (June): 465-79.

Zarnowitz, Victor, and Phillip Braun. 1993. "Twenty-Two Years of the NBER-ASA Quarterly Economic Outlook Surveys: Aspects and Comparisons of Forecasting Performance." In James H. Stock and Mark W. Watson, eds., Business Cycles, Indicators, and Forecasting. Chicago: University of Chicago Press.

Zarnowitz, Victor, and Louis A. Lambros. 1987. "Consensus and Uncertainty in Economic Prediction." Journal of Political Economy 95, no. 3 (June): 591-621.

Zwiebel, Jeffrey. 1995. "Corporate Conservatism and Relative Compensation." Journal of Political Economy 103, no. 1 (February): 1-25.

Appendix 1. An algorithm for determining the equilibrium distribution of forecasters given $f(x)$, $\{r_i\}$, and $\{N_i\}$, for $i=1,\dots,m$.

First, provisionally assume a value for $n(0)$, the number of forecasters predicting the value 0. This value will determine a value of k_1 and an associated iso-wage curve $n^1(x)$ for forecasters in industry 1. Working from 0 outward, assign $n^1(x)$ forecasters to values of x until all sector 1 forecasters are exhausted, at $x=\pm c_1$. Then assign enough sector 2 forecasters to $\pm c_1$ so that the total number of forecasters predicting these two values are, respectively, $n^1(-c_1)$ and $n^1(c_1)$. The values of $n^1(-c_1)$ and $n^1(c_1)$ will then determine an iso-wage curve, $n^2(x)$, for sector 2 forecasters. Continuing to move away from the origin, sector 2 forecasters will distribute themselves along this iso-wage curve until they, too, are exhausted. This recursive procedure is then repeated until either all N forecasters are assigned and some values of x are unaccounted for, or the entire range of x is blanketed but some forecasters remain unassigned. In either case, another value of $n(0)$ can be chosen and the procedure repeated until a value of $n(0)$ is found for which the entire range of values of x is covered and all forecasters are assigned.

Appendix 2

Forecasters Participating in the Blue Chip Economic Indicators

1976-95

Banks (30)

Bank of America
Bankers Trust Co.
Brown Brothers Harriman
Chase Manhattan Bank
Chemical Banking
Citibank
Comerica
Connecticut National Bank
CoreStates Financial Corp.
First Fidelity Bancorp
First Interstate Bank
First National Bank of Chicago
Fleet Financial Group
Harris Trust and Saving
Irving Trust Company
J P Morgan
LaSalle National Bank
Manufacturers Hanover
Manufacturers National Bank of Detroit
Marine Midland
Mellon Bank
National City Bank of Cleveland
Northern Trust Company
Philadelphia National Bank/ PNC Bank
Provident National Bank
Security Pacific Bank
Shawmut National Corp.
United California Bank
U.S. Trust Co.
Wells Fargo Bank

Econometric Modelers (12)

Chase Econometrics
Data Resources, Inc.
Fairmodel-Economics, Inc.
Georgia State University
Gil Heebner, Eastern College
Inforum - University of Maryland
Laurence H. Meyer & Associates
Merrill Lynch Economics
Michigan Quarterly U.S. Model
UCLA Business Forecasting
University of Illinois (B.T.)
Wharton Econometrics /WEFA Group

Industrial Corporations (18)

B.F. Goodrich
Caterpillar
Chrysler Corporation
Conrail
Eaton Corporation
DuPont
Ford Motor Company
General Electric Company
General Motors
Machinery & Allied Products
Monsanto Company
Motorola, Inc.
Pennzoil Company
Predex Corp.
Sears Roebuck
Union Carbide
Weyerhaeuser Co.
W.R. Grace

Independent Forecasters (38)

Albert T. Sommers
Argus Research
Arthur D. Little
Ben E. Laden Associates
Business Economics, Inc.
Center for Study of American Business
Computer Aided Production Planning Systems, Inc.
DePrince & Associates
DeWolf Associates
Econoclast
Econoviews International, Inc.
Evans Economics, Inc.
George Gols
Hagerbaumer Economics
Heinemann Economic Research
Helming Group
Herman I. Leibling & Associates
InfoMetrica, Inc.
Joel Popkin & Co.
Juodeika Allen & Co.
Leonard Silk, NY Times
MAPI
Morris Cohen & Associates
Moseley, Hallgarten, & Estabrook
Oxford Economics USA
Peter L. Bernstein, Inc.
Polyconomics
Reeder Associates (Charles)
Robert Genetski and Associates, Inc.
Rutledge & Co.
Schroder, Naess, and Thomas
Sindlinger Company, Inc.
SOM Economics, Inc.
Statistical Indicator Associates
Stotler Economics
The Bostian Group - HHG
Turning Points Micrometrics
Wayne Hummer & Company- Chicago

Securities Firms (14)

American Express / Shearson Lehman Co
Arnhold and S. Bleichroeder
A.G. Becker / Becker Associates
A.G. Edwards & Company
Chicago Capital, Inc.
CRT Government Securities
C.J. Lawrence, Inc.
Dean Witter Reynolds, Inc.
Goldman, Sachs Co.
Ladenburg, Thalmann, & Co.
Loeb Rhoades, Hornblower, & Co.
Morgan Stanley & Co, Inc.
NationsBanc Capital Markets Inc.
Prudential Securities, Inc.

Other (17)

Financial Publications

Cahners Economics
Financial Times Currency Forecaster
Eggert Economics Enterprises, Inc.
Fortune Magazine

Government Agencies

Bush Administration
Clinton Administration
Congressional Budget Office
Office of Management and Business

Industry Associations

Conference Board
Mortgage Bankers Association of America
National Association of Home Builders
U.S. Chamber of Commerce

Insurance Companies




Equitable Life
Metropolitan Life Insurance Co.
Prudential Insurance Co.

Ratings Agencies

Dun & Bradstreet
Standard and Poor's Corp.


Chart 1

A Family of Iso-wage Curves



Notes: The probability distribution function, $f(x)$, is a linear transformation of the binomial distribution with $n=400$, $p=q=.5$. It is scaled to have zero mean and a variance of one. The iso-wage curves are constructed for the parameter values [$P = \$5,000,000$; $w^* = \$50,000$; $s = 50,000$; and $b=1$].

Comparative Statistics: Equilibrium Distribution of Forecasters for Varying Degrees of Emphasis on Accuracy



Notes: The probability distribution function, $f(x)$, is a linear transformation of the binomial distribution with $n=400$, $p=q=.5$. It is scaled to have zero mean and a variance of one. The two alternative equilibrium distributions of forecasters were constructed using the parameter values:


$N = 100$
 $P = \$5,000,000$
 $r_1 = 50,000$
 $r_2 = 5,000$

Thus, n_1 reflects more of an emphasis on accuracy than does n_2 .

While the distributions of forecasters are sketched here as curves, they are actually discrete and take on values at intervals of 0.1.

Chart 3a

Equilibrium Distribution of Forecasters: Multi-Industry Case




Notes: The probability distribution function, $f(x)$, is a linear transformation of the binomial distribution with $n=400$, $p=q=.5$. It is scaled to have zero mean and a variance of one. While the distribution of forecasters is sketched here as a curve, it is actually discrete and takes on values at intervals of 0.1. P is set equal to \$5,000,000. The values of r_i and N_i (the number of firms per industry) are assumed to be:

<u>Industry</u>	<u>r_i</u>	<u>N_i</u>
$i = 1$	250,000	20
2	50,000	20
3	25,000	20
4	5,000	20
5	0	20

Chart 3b

Equilibrium Distribution of Forecasters: Multi-Industry Case

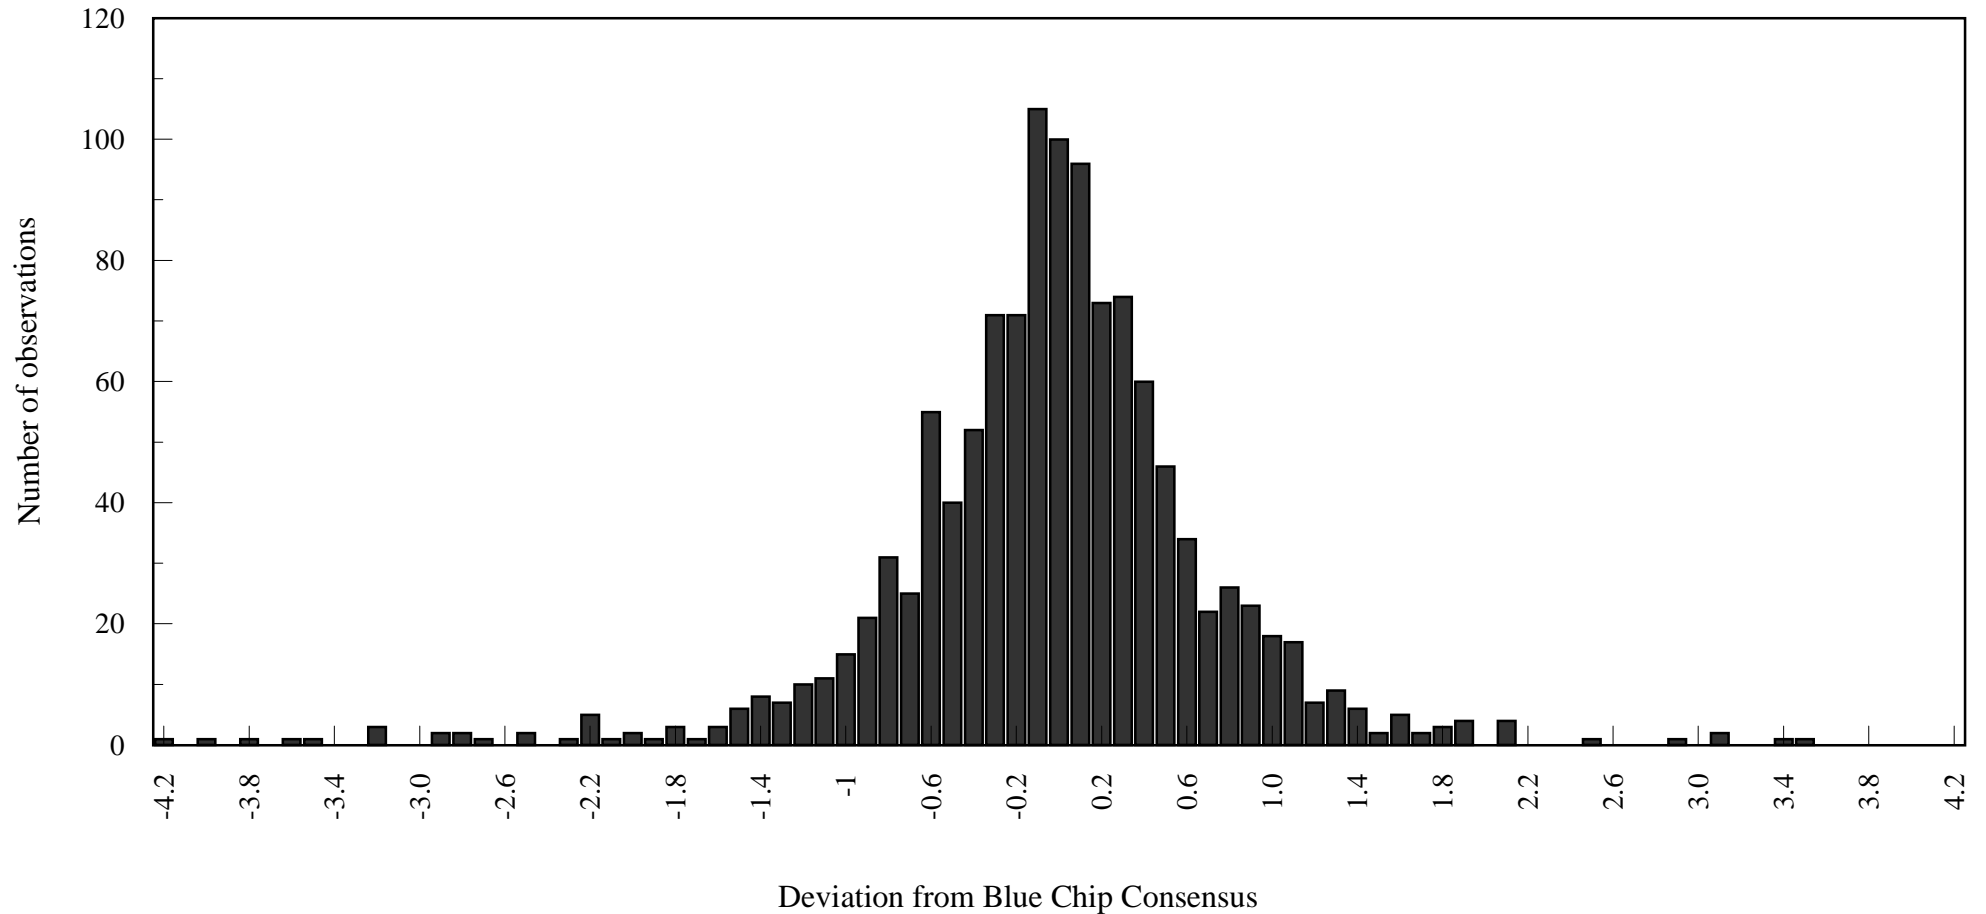


Notes: The probability distribution function, $f(x)$, is a linear transformation of the binomial distribution with $n=400$, $p=q=.5$. It is scaled to have zero mean and a variance of one. P is set equal to \$5,000,000. The values of r_i and N_i (the number of firms per industry) are assumed to be:

Industry	r_i	N_i
$i = 1$	250,000	20
2	50,000	20
3	25,000	20
4	5,000	20
5	0	20

Chart 4

Distribution of Individual Year-End Real GDP/GNP Forecasts Relative to Consensus Blue Chip Indicators, 1977-96



Source: *Blue Chip Economic Indicators.*

Table 3

Regressions of Absolute Deviations from Consensus as a Function of Industry Secto

Dependent variable: |GDP(i)-Consensus|

Sample Period: 1977-96

Explanatory Variables:				
	Coefficient	t-statistic	Coefficient	t-statistic
Constant	0.52**	33.89	0.52**	35.71
BAN	-0.11**	-4.44	-0.12**	-5.14
ECO	-0.09*	-2.16	-0.10**	-2.59
IND	0.30**	11.01	0.31**	11.80
COR	-0.15**	-3.93	-0.13**	-3.50
SEC	-0.00	-0.08	-0.02	-0.36
OTH	-0.07	-1.83	-0.06	-1.76
Y1977	--	--	-0.16	-1.87
Y1978	--	--	-0.09	-1.17
Y1979	--	--	0.08	1.02
Y1980	--	--	0.19*	2.56
Y1981	--	--	0.23**	3.09
Y1982	--	--	0.39**	5.31
Y1983	--	--	0.16*	2.14
Y1984	--	--	-0.13	-1.74
Y1985	--	--	-0.03	-0.44
Y1986	--	--	0.15*	2.26
Y1987	--	--	0.18**	2.71
Y1988	--	--	0.22**	3.86
Y1989	--	--	-0.02	-0.34
Y1990	--	--	0.03	0.58
Y1991	--	--	0.11*	1.98
Y1992	--	--	-0.02	-0.44
Y1993	--	--	-0.22**	-4.06
Y1994	--	--	-0.25**	-4.62
Y1995	--	--	-0.23**	-4.22
Y1996	--	--	-0.21**	-4.46
R-squared:	0.10		0.20	
F-statistic:	25.35**		12.17**	

Source: *Blue Chip Economic Indicators*.Notes: The regressions were run using the panel of 1,197 individual forecasts of real GDP appearing in the 1977-96 De *Blue Chip Economic Indicators*. Explanatory variables are industry sector dummies, defined as follows:

BAN = banks

ECO = econometric modelers

IND = independent forecasters

COR = industrial corporations

SEC = securities firms

OTH = other miscellaneous forecasters (financial publications, government agencies, industry associations, insurance companies, and rating agencies).

Industry dummies and year dummies are constrained to have an observation-weighted mean value of zero.

* significant at the 95 percent level.

** significant at the 99 percent level.