WORKING PAPER

2002

Department of Economics
Tufts University
Medford, MA 02155
(617) 627 – 3560
http://ase.tufts.edu/econ

Who Herds?*

Dan Bernhardt
Department of Economics
University of Illinois
Champaign, IL 61820

Murillo Campello
Department of Finance
University of Illinois
Champaign, IL 61820

Edward Kutsoati
Department of Economics
Tufts University
Medford, MA 02155

March 21, 2002

Abstract

This paper develops a test for forecast bias that is robust to (a) correlated information amongst analysts; (b) common unforecasted industry-wide earnings shocks; (c) information arrival at different points of the forecast horizon; and (d) the possibility that the measure of earnings that analysts forecast differs from that which the econometrician observes. Empirically, we find that forecasts are biased, but that analysts do not herd. On the contrary, anti-herding obtains: Analysts systematically issue biased *contrarian* forecasts that overshoot the publicly–available consensus forecast in the direction of their private information. The forecast bias is economically large and declines with the amount of information at an analyst's disposal. The magnitude of the bias, its systematic variation with analyst following, and the pattern of bias in forecast revisions indicate that the bias is strategically chosen. In particular, the data *cannot* be explained by any form of analyst overconfidence.

*We are grateful to the Institutional Brokers Estimate System (I/B/E/S), a service of I/B/E/S International Inc. for providing data on analysts forecasts. The data was provided as part of a broad academic program to encourage earnings expectation research. We thank Long Chen and Selim Tepaloglu for their useful comments and suggestions. Comments of seminar participants at the Federal Reserve Bank of Chicago, University of Southern California, Claremont-McKenna College, University of Colorado, and the University of Illinois are also appreciated. The usual disclaimer applies.

1 Introduction

"Going against the consensus is the way an analyst makes his or her mark."
— Steven Garmaise, *Midland Walwyn Capital*

Both the financial press and academic research regularly suggest that security analysts herd toward the consensus forecast, issuing forecasts that under-weight their own information (Trueman (1990, 1994), Hong, Kubik and Solomon (2000)) These herding stories have arisen because forecasts seem "too clustered." For example, Gallo, Granger and Jeon (2001) find that forecasts of GDP converge as the date at which GDP is announced draws nearer, but that invariably final forecasts are either uniformly too low or too high. The authors conclude that forecasters herd.

But, clustered forecasts need not imply that analysts herd. First, earlier forecasts may contain information that subsequent analysts, as proper Bayesians, incorporate into their forecasts. Second, analysts rely on common information sources. In particular, all earnings analysts rely heavily on a company's chief financial officer for information. If a CFO tells each analyst the same thing, their forecasts will reflect this common (perhaps mis)information, so that their forecasts will tend all to be too high or too low relative to announced earnings. Indeed, if analysts have identical information, each unbiased analyst will make the same forecast, which will invariably differ from actual earnings. Third, common unanticipated market-wide shocks can cause most, if not all, forecasts to be too low or too high relative to earnings. For example, in 2000, most earnings forecasts for technology stocks greatly exceeded earnings. However, the fact that the forecasts systematically differed from earnings does *not* imply that the forecasts were biased; the huge decline in stock prices of technology firms suggests that this earnings decline was also unanticipated by the entire market. Fourth, the measure of earnings that analysts forecast may differ from the earnings that the econometrician sees. For example, analysts may not seek to forecast some exceptional items that appear in reported earnings (Abarbanell and Lehavy 2001). Finally, information arrival may cause forecasts to appear "too scattered," because analysts forecasting at later dates have access to more information, and analysts who report earlier often do not revise their forecasts (Trueman 1990).

This paper tests for herding and bias in the forecasts of earnings analysts, extending a test developed in Bernhardt and Kutsoati (2001). Our test for forecast bias is robust to <u>each</u> of the concerns highlighted above. The key idea underlying the test is that if an analyst is unbiased, then his forecast should be as likely to exceed realized earnings as fall short, both unconditionally, and conditional on anything in his information set. If, instead, an analyst herds, biasing his forecast toward the consensus *given* his information, then his forecast will be located between the consensus and his private estimate of what earnings will be. Consequently, if an analyst herds, then when his forecast exceeds the consensus, it should fall short of realized earnings more than half of the time. So too, when a herding analyst's forecast falls short of the consensus, it should exceed earnings more than half of the time. The opposite outcome is predicted if an analyst seeks to issue forecasts that distinguish him from other analysts. Then, an analyst will issue a forecast that will overshoot his private estimate of what earnings will be in the direction away from the consensus.

The key feature that these observations share is that no assumptions are made about how an analyst's information set is formed. This implies that our tests for unbiasedness or herding are unaffected by correlation in signals or information arrival. Essentially, we estimate two conditional probabilities: the conditional probability that a forecast exceeds realized earnings given some conditioning information, and the conditional probability that a forecast falls short of earnings (given some other conditioning event). To control for possible unforecasted earnings shocks, our test statistic averages these two conditional probabilities: under a null of unbiasedness, an unforecasted earnings shock has off-setting impacts on the frequency with which each over-shooting event occurs.

Bernhardt and Kutsoati (2001) explore theoretically how relative performance compensation affects analyst forecasts, so that analysts care about both absolute and relative forecast accuracy. If the forecast industry uses relative forecast accuracy to gauge analyst ability, then relative forecast compensation arises naturally.[1] If compensation is a convex function of relative forecast accuracy, then analysts try to separate their forecasts from those of other analysts. As a result, analysts issue forecasts that overshoot their best guess of earnings in the direction away from the consensus. In contrast, concave relative performance compensation induces analysts to issue forecasts that are closer to the consensus forecast than their information would suggest; *i.e.* analysts will herd so that forecasts will be overly clustered. This theoretical analysis suggests likely variables on which to condition in order to test for the unbiasedness of forecasts: (1) the amount by which a forecast exceeds the consensus, (2) whether the forecast was a revision, (3) analyst coverage, (4) the number of prior forecasts, and so on.

No matter where we search, we find overwhelming evidence *against* herding behavior. Rather, brokerage houses and money management funds appear to employ only the Thomas Kurlaks of the world to do their forecasting. Of Kurlak, Merrill Lynch's semiconductor analyst, it is said that "When Kurlak likes a company, his estimates tend to be much higher than everyone else's; when he is down on a company, they're much lower" (*"Who Really Moves The Market"*, FORTUNE, October, 1997). Analysts systematically issue biased *contrarian* forecasts that overshoot the con-

[1]The intense competition in the forecasting industry to win clients gives rise to intense competition to identify able analysts. There is ample evidence that analysts with better forecasting record are seen as more skilled, and, as the subjects of robust competition, receive large wage increases and bonuses.

sensus forecast in the direction of their private information. The conditional probability that an analyst's forecast overshoots actual earnings per share (EPS) in the direction away the consensus is 0.60. A forecast that falls short of the consensus falls short of EPS 62% of the time; and when it exceeds the consensus, it exceeds EPS 58% of the time. Further, the probability an analyst's forecast overshoots EPS *away* from the consensus rises rapidly with the magnitude by which an analyst's forecast differs from consensus. We find similar results for all analysts — the second, third, next-to-last, last, and so on. Our test statistic is remarkably stable, varying by less than five percentage points across all years and analyst-coverage, despite large variations in common earnings shocks. These findings are the opposite of those predicted by herding.

The other contribution of this paper is to quantify the magnitude of the forecast bias, and to distinguish between strategic contrarian and overconfidence explanations of the data. An overconfident analyst places excessive weight on his private signal, so that if his signal exceeds those of other analysts, then his forecast will tend to overshoot earnings in the direction of his signal (Ehrbeck and Waldmann 1996). What makes distinguishing between strategic contrarian and overconfidence explanations such a challenge is that there is a wide range of overconfidence alternatives reflecting different magnitudes of over-weighting of the private signal in the forecast: we have to show that no overconfidence alternative alone can be consistent with the data.

Distinguishing between these explanations has important policy and modeling implications. With regard to policy, there is considerable evidence that private investors may be overconfident and behave in predictable suboptimal ways (Odean (1999)). However, the actions of experts working for large institutions likely have far greater impacts on aggregate market outcomes, rendering the presence of myopic private investors less important. With regard to modeling, the economics profession has a stake in rejecting overconfidence explanations in the following sense: it is hard to imagine environments where the assumption that agents behave rationally in accord with their incentives is more appropriate than that here. Analysts are trained and experienced, and compensation is highly sensitive to performance. If trained analysts with so much on the line did not act rationally, then it would suggest that behavioral models must be almost universally considered.

To see why and how we reject overconfidence alternatives, consider the overshooting implications of the most extreme formulation of overconfidence possible. To be specific, suppose that signals are independently and identically drawn from a normal distribution, and analysts report unbiased forecasts given only their *own* signals, ignoring *all* information contained in earlier forecasts. Then, for every one percent of price by which the second analyst overshoots the consensus (the first forecast), he would overshoot earnings on average by one-half of a percent, as the best estimate of earnings given the two forecasts is their average. More generally, the j^{th} analyst would overshoot on average by $1 - \frac{1}{j}$ of a percent, so that overshooting magnitudes should rise for later analysts.

Next note that overconfident analysts should overshoot earnings by less if either (a) in addition to being overconfident, they herd; (b) signals are positively correlated; (c) later analysts have better information; or (d) analysts are overconfident, but not excessively so. To see this, first note that herding effectively raises the forecast weight on the consensus. Second, if signals are positively correlated or if later analysts have better information, then there is less incremental information in earlier forecasts. Finally, less overconfident analysts place greater weight on earlier forecasts. Each of these possibilities implies reduced overshooting magnitudes. Since correlated information and information arrival are both important, and extreme overconfidence is unreasonable, this $1 - \frac{1}{j}$ bound grossly overstates the overshooting that could be consistent with overconfidence. We then show that the forecast bias is *too large* to be explained by overconfidence. For example, for every one percent of price that the second analyst overshoots the consensus, he overshoots earnings by 0.73 of a percent, significantly exceeding the upper bound on overshooting consistent with overconfidence.

Under the assumption that analysts choose forecast biases that are linear functions of the expected error in the consensus forecast, we derive estimates of both the strategic bias and its expected dollar impact. The overshooting magnitudes imply a strategic bias of 1.7 times the difference between the analyst's (unobserved) true estimate of earnings and the consensus forecast if he were the 2^{nd} to report; and if he reports last (20^{th}), the strategic bias falls to 0.44 times the difference. Thus, the forecast bias falls rather than rises with the information at an analyst's disposal, but it always remains economically significant. Using reasonable bounds on correlation, information arrival and overconfidence, our findings indicate that analysts overshoot by too much to be explained by overconfidence alone. That is, strategic contrarian behavior is *necessary* to explain the data: we can reject unbiasedness and herding explanations in combination with all forms of overconfidence.

We also find bias in revised forecasts that is inconsistent with overconfidence. If analysts revise their forecasts toward the consensus, overshooting frequencies are quite high; but if a forecast revision moves away from the consensus, the revised forecast is essentially unbiased. Overconfidence predicts the opposite. That is, forecasts that move further away from the consensus should be *more* likely to overshoot. Strategic contrarian behavior can reconcile these findings: analysts who move toward the consensus do so reluctantly because they want to preserve separation from other forecasts; whereas analysts who move away from the consensus are confident that their forecasts are more accurate, so that they do not bias their forecasts. In sum, our paper provides evidence that strategic contrarian considerations and not overconfidence or herding underlie forecast choices.

Most earlier attempts at detecting herding did so by estimating the deviation of each forecast from the mean of all forecasts reported in that forecast period. With this approach, Hong, Kubik and Solomon (2000) find young and inexperience analysts are more likely to herd toward the consensus, and are more likely to lose their jobs for making forecasts that deviate away. Lamont (1995) offers similar findings for forecasters of GNP and other macroeconomic indicators.[2] A concern with this testing strategy, however, is that it does not account for either correlation in information or information arrival. Welch (2000) takes a different approach. Using a dataset of analysts' recommendations, he finds that the prevailing consensus and the two most recent revisions have a positive influence on the next analyst who makes a recommendation. Although Welch interprets his results as evidence of herding, they may also reflect the fact that recent revisions contain useful information that analysts incorporate into their own forecasts.

Our results contrast sharply with those of Keane and Runkle (1998), who claim that "professional stock market analysts make unbiased forecasts of earnings." As in our paper, Keane and Runkle try to control for correlated signals. However, they restrict their analysis to 21 very heavily-followed firms. We find that forecasts biases are smallest for such firms, and conjecture that sample selection may explain their failure to uncover the contrarian bias. Indeed, our findings provide strong support for the theory that convex relative performance compensation drives forecasts. In particular, Bernhardt and Kutsoati (2001) show that convex relative performance compensation causes analysts to issue contrarian forecasts, and that the strategic forecast bias should fall when analysts have more information. These results complement those of Zitzewitz (2001), who finds that analysts issue forecasts that exaggerate their information.

The rest of the paper is organized as follows. Section 2 develops our testable hypotheses, and presents our estimation approach. Section 3 explains our sample selection procedure, presents our empirical findings, and explores their economic significance. Section 4 concludes.

2 Do analysts herd?

An analyst is said to issue an unbiased forecast if and only if the forecast is equal to the analyst's posterior estimate of the median earnings per share *given* the information at the analyst's disposal.[3] An analyst's information includes both the information uncovered through his own analysis, as well as public information released by the firm, and forecasts by earlier analysts. Herding is a choice by an analyst to announce a forecast of earnings that is closer to the *publicly known consensus* than

[2]See also Laster, Bennett and Geoum (1999), and Ehrbeck and Waldmann (1996).

[3]The median corresponds to the mean if an analyst's posterior distribution over earnings is symmetric.

his information suggests, so that his forecast is located between the consensus and his posterior estimate of the median of earnings. This is the appropriate notion of herding in an environment in which earlier forecasts contain information about earnings and/or reflect new information arrival, and analysts update their beliefs about earnings following forecast observations.

Formally, for the i^{th} analyst to report an earnings forecast for firm j in quarter k, let $\hat{E}_{ijk}$ be the median of his posterior distribution $G_{ijk}(\cdot)$ over earnings per share given his information. Also, let F_{ijk} be his forecast, and let $\overline{F}_{ijk}$ be the outstanding consensus (average) forecast at the moment analyst i issues his forecast. Then i reports an unbiased forecast if and only if $F_{ijk} = \hat{E}_{ijk}$. While we do not observe $\hat{E}_{ijk}$, if i's forecast is unbiased, it should be as likely to exceed earnings as fall short, both unconditionally, and conditional on anything in i's information set. In particular, the likelihood with which i's forecast exceeds actual earnings should not vary with the amount by which i's forecast exceeded or fell short of the consensus.

If, instead, analyst i herds toward the consensus, then his reported forecast is between his private estimate of earnings and the consensus, so that either $\overline{F}_{ijk} < F_{ijk} < \hat{E}_{ijk}$ or $\overline{F}_{ijk} > F_{ijk} > \hat{E}_{ijk}$. As a result, given that a forecast exceeds the consensus, it should fall short of realized earnings more than half of the time; and given that a forecast is less than the consensus, it should exceed realized earnings more than half of the time.

Finally, if analyst i tries to distinguish himself from the others, then when his private estimate of earnings exceeds the consensus, his forecast will also exceed his private estimate. That is, if $\overline{F}_{ijk} < \hat{E}_{ijk}$ then $\overline{F}_{ijk} < \hat{E}_{ijk} < F_{ijk}$; and if $\overline{F}_{ijk} > \hat{E}_{ijk}$ then $\overline{F}_{ijk} > \hat{E}_{ijk} > F_{ijk}$. Such "anti-herding" behavior gives rise to overshooting: if i's forecast exceeds the consensus, then his forecast should exceed realized earnings more than half of the time; and if his forecast is less than the consensus, it should fall short of realized earnings more than half of the time.

Let τ be an index for a generic forecast by an analyst in some firm-quarter. We first select conditioning events z_τ^+ and z_τ^-, where z_τ^+ implies that the analyst's forecast exceeds the outstanding consensus; and z_τ^- implies that the forecast fell short of the consensus. For example, z_τ^+ might be the event that (a) forecast F_τ exceeded the consensus, $\overline{F}_\tau$, by at least x percent of the stock price at the end of the previous quarter, *i.e.* $\frac{F_\tau - \overline{F}_\tau}{P_{\tau-1}} \geq x$, (b) the analyst was the third analyst to report, and (c) the analyst had issued a prior forecast, so that F_τ represented a revised forecast. Similarly, z_τ^- might be the event that $\frac{\overline{F}_\tau - F_\tau}{P_{\tau-1}} \geq x$, and F_τ was a revised forecast issued by the third analyst.

Next, define conditioning indicator functions, γ_τ^+ and γ_τ^-, where

$$\gamma_\tau^+ = 1 \text{ if } z_\tau^+ \text{ occurred; } \gamma_\tau^+ = 0 \text{ if } z_\tau^+ \text{ did not occur,}$$
$$\gamma_\tau^- = 1 \text{ if } z_\tau^- \text{ occurred; } \gamma_\tau^- = 0 \text{ if } z_\tau^- \text{ did not occur.}$$

Similarly, define overshooting indicator functions, δ_τ^+ and δ_τ^-, where:

$$\delta_\tau^+ = 1 \text{ if } F_\tau > E_\tau \text{ and } \gamma_\tau^+ = 1; \text{ and } \delta_\tau^+ = 0 \text{ otherwise,}$$
$$\delta_\tau^- = 1 \text{ if } F_\tau < E_\tau \text{ and } \gamma_\tau^- = 1; \text{ and } \delta_\tau^- = 0 \text{ otherwise,}$$

and E_τ is the actual earnings in firm-quarter τ. Thus, $E[\delta_\tau^+|\gamma_\tau^+ = 1]$ is the probability that $F_\tau > E_\tau$ given that event z_τ^+ occurred; and $E[\delta_\tau^-|\gamma_\tau^- = 1]$ is the probability that $F_\tau < E_\tau$ given that event z_τ^- occurred. Under the null hypothesis that analysts issue unbiased forecasts

$$E[\delta_\tau^+|\gamma_\tau^+ = 1] = E[\delta_\tau^-|\gamma_\tau^- = 1] = 0.5.$$

If, instead, analysts herd, then because $z_\tau^+ = 1$ indicates that $F_\tau > \overline{F}_\tau$, herding implies that $\hat{E}_\tau > F_\tau > \overline{F}_\tau$. Consequently, with herding,

$$E[\delta_\tau^+|\gamma_\tau^+ = 1] < 0.5.$$

Similarly, with herding, $z_\tau^- = 1$ implies that $\hat{E}_\tau < F_\tau < \overline{F}_\tau$ so that

$$E[\delta_\tau^-|\gamma_\tau^- = 1] < 0.5.$$

Conversely, if analysts anti-herd,

$$E[\delta_\tau^+|\gamma_\tau^+ = 1] > 0.5 \qquad \text{and} \qquad E[\delta_\tau^-|\gamma_\tau^- = 1] > 0.5.$$

Figure 1 illustrates the events δ_τ^+ and δ_τ^-. If an analyst's forecast is median-unbiased, then *independently of anything in his information set, including the location of the consensus,* F_τ should correspond to his estimate of the median earnings per share, $\hat{E}_\tau$. If, instead, an analyst biases his forecast toward the consensus, then his forecast will be located between $\hat{E}_\tau$ and $\overline{F}_\tau$, making the conditional probability of overshooting (*i.e.*, the frequency of events δ_τ^+ and δ_τ^-) less likely; while if the analyst biases his forecast away from the consensus, this raises the likelihood of events δ_τ^+ and δ_τ^-. This central insight regarding how different forms of bias affect the probabilities of overshooting forms the crux of our test for forecast bias.

Under the null that forecasts are unbiased, both $\sum_\tau \delta_\tau^+$ and $\sum_\tau \delta_\tau^-$ are binomially distributed with mean 0.5. We use sample averages to estimate the population overshooting probabilities: $\frac{\sum_\tau \delta_\tau^+}{\sum_\tau \gamma_\tau^+}$ is our estimate of the conditional probability of overshooting earnings given z^+; while $\frac{\sum_\tau \delta_\tau^-}{\sum_\tau \gamma_\tau^-}$ is our estimate of the conditional probability of falling short of earnings given z^-. Our test statistic is the average of our estimates of the two conditional overshooting probabilities:

$$\mathbf{S}(z^+, z^-) = \frac{1}{2}\left[\frac{\sum_\tau \delta_\tau^+}{\sum_\tau \gamma_\tau^+} + \frac{\sum_\tau \delta_\tau^-}{\sum_\tau \gamma_\tau^-}\right]. \tag{1}$$

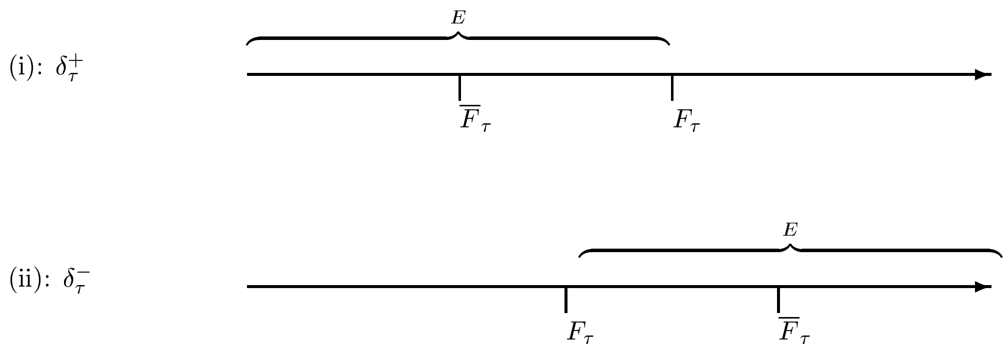

Figure 1: Realization of actual EPS relative to F_τ and $\overline{F}_\tau$

Crucially, the properties of **S** do not depend on how analysts form their posteriors over earnings, $G_\tau(\cdot)$. In particular, the posterior could reflect correlated signals amongst analysts for a given firm quarter and/or information arrival over time; neither affects our test of median unbiasedness.

Also, **S** has the property that its distribution is well-behaved asymptotically if the number of firm quarters is large, even if there are common unforecasted earnings shocks across firms in a given time period, and there are few time periods. Suppose earnings per share for firm i in quarter t contain an unforecasted market shock to earnings common to all firms in period t, ω_t:

$$E_{it} = \hat{E}_{it} + \omega_t + \epsilon_{it},$$

where ϵ_{it} is independently distributed across time and firms, and ω_t is unforecasted by the analyst. The unforecasted earnings shocks may be large. For instance, in 1999, the last analyst's forecast fell short of earnings 73% of the time, presumably due to large positive unforecasted earnings shocks.

Suppose first that $\omega_\tau = \omega, \forall\tau$, and that analysts' posteriors about earnings realizations for different firms are identically distributed about their medians, $G_\tau(\hat{E}_\tau - E_\tau|\Omega_\tau) = G_{\tau'}(\hat{E}_{\tau'} - E_{\tau'}|\Omega_{\tau'})$, $\forall\tau,\tau'$. Under the null hypothesis that forecasts are unbiased, $\sum_\tau \delta_\tau^+$ has a binomial distribution, $B(\sum_\tau \gamma_\tau^+, G(\hat{E}+\omega))$; and $\sum_\tau \delta_\tau^-$ has a binomial distribution, $B(\sum_\tau \gamma_\tau^-, 1 - G(\hat{E}+\omega))$. Under the null, **S** has mean,

$$\frac{1}{2}G(\hat{E}+\omega) + \frac{1}{2}(1 - G(\hat{E}+\omega)) = \frac{1}{2},$$

independently of ω and $g(\cdot)$, and variance

$$\frac{G(\hat{E}+\omega)(1-G(\hat{E}+\omega))}{4}\left[\frac{1}{\sum_\tau \gamma_\tau^+} + \frac{1}{\sum_\tau \gamma_\tau^-}\right] \leq \frac{1}{16}\left[\frac{1}{\sum_\tau \gamma_\tau^+} + \frac{1}{\sum_\tau \gamma_\tau^-}\right].$$

Thus, under the null of median unbiasedness, asymptotically,

$$\mathbf{S} \sim \mathcal{N}\left(\frac{1}{2}, \frac{G(\hat{E}+\omega)(1-G(\hat{E}+\omega))}{4}\left[\frac{1}{\sum_\tau \gamma_\tau^+} + \frac{1}{\sum_\tau \gamma_\tau^-}\right]\right).$$

The variance of $\mathbf{S}$ is bounded from above by $\frac{1}{16}\left[\frac{1}{\sum_\tau \gamma_\tau^+} + \frac{1}{\sum_\tau \gamma_\tau^-}\right]$, which does not depend on ω or $g(\cdot)$. Since neither the mean of $\mathbf{S}$ nor this variance bound are affected by variations in ω_τ and $g_\tau(\cdot)$ across firm quarters, if the null hypothesis of unbiasedness is rejected using the upper bound on the variance, then it can always be rejected. In sum, if forecasts are unbiased, then $\mathbf{S} = 0.5$. The alternative of herding implies that $\mathbf{S} < 0.5$; and anti-herding implies that $\mathbf{S} > 0.5$.

Importantly, a similar analysis reveals that our $\mathbf{S}$ test for median unbiasedness is robust to the possibility that analysts are not trying to forecast the earnings reported by I/B/E/S, but rather seek to forecast some other measure of earnings, $E_\tau + \varepsilon_\tau$, that is unobserved by the econometrician. For example, it may be that analysts do not try to forecast large negative earnings accruals that some firms take (Abarbanell and Lehavy 2001), but which are included in the earnings reported by I/B/E/S. As long as ε_τ is distributed independently from $F_\tau - \overline{F}_\tau$ (this would hold, for example, if ε_τ is the same for all analysts), then under the null hypothesis that analysts report the posterior median of the distribution over the earnings variable that they seek to forecast, the mean of $\mathbf{S}$ remains equal to 0.5 and the variance of $\mathbf{S}$ is reduced, so that the same upper bound on the variance can be used to reject the null of unbiasedness.

Finally, note that outliers do not have disproportionate effects on frequencies. Consequently, the tests are robust to failing to exclude data entry errors, or unusual "one-time" events such as large discretionary write-downs of assets that analysts do not seek to predict (see the concerns detailed in Keane and Runkle (1998) and Lim (2001)).

3 Empirical Analysis

3.1 Data and Sample Selection

Data on individual analysts' quarterly forecasts of earnings from 1989 to 1999 were taken from the Institutional Brokers Estimate System (I/B/E/S) Detail tapes. Each observation includes the company ticker, forecast horizon, codes that identify each analyst and brokerage house, the analyst's earnings estimate for the period-end and the date that the forecast was reported to I/B/E/S. We do not consider forecasts prior to 1989 because of the lag between the date of an analyst's forecast and the date the forecast was entered in the I/B/E/S database during this period.[4] After 1988, forecasts

[4]O'Brien (1988) reported an average publication lag of 34 trading days over the period 1975-1982.

were disseminated by I/B/E/S within 24 hours. Our sample selection ensures that publication dates are close to the actual dates that analysts released their forecasts and, more importantly, the analyst reporting *observes* earlier forecasts.

Actual earnings per share data, reported on the same basis (primary or diluted) as the corresponding forecasts, are also taken from the I/B/E/S database. Share prices and the number of shares outstanding at the end of each quarter are extracted from the Center for Research on Security Prices (`CRSP`) files. We filter the data for likely data entry errors, deleting any forecast with an absolute error value exceeding \$10.[5] Some forecasts are dated over 360 days before the actual earnings report, and could be data entry errors (entries for the wrong firm quarter). To deal with this, we include only the forecasts reported within 120 days of the actual earnings announcement. Then, for each firm-quarter, we count the number of analysts reporting forecasts (`cover`), and the order in which a forecast was reported in the sequence (`order`). We also track total forecasts, including revisions, reported each firm-quarter.

Analysts sometimes revise their old forecasts. Consequently, we only use the most recent forecast reported by each analyst when processing the consensus forecast. For every firm-quarter, we order forecasts by the date of release. We then define a forecast reported at time τ by F_τ and the mean of all other forecasts reported at least ℓ days before time τ as $\overline{F}_{\tau-\ell}$. $\overline{F}_{\tau-\ell}$ is then the *outstanding* consensus mean forecast of those reported at least ℓ days before date τ.

The procedure is illustrated with a hypothetical example in Table 1. The table presents a sequence of release of forecasts by five distinct analysts for a given firm-quarter. The total number of forecasts is 8. The data is sorted by the date of publication.

Table 1: Example of the sequence of forecasts and the processing of the consensus forecast

Obs. #	Analyst identity	Forecast	Date of publication (yy/mm/dd)	Order of forecast (for consensus estimation)
1.	1	0.75	92/04/17	1
2.	1	0.70	92/04/20	1
3.	2	0.72	92/04/27	2
4.	3	0.82	92/05/03	3
5.	4	0.82	92/05/03	3
6.	5	0.80	92/05/15	5
7.	3	0.84	92/05/15	4
8.	1	0.83	92/06/10	5

[5]O'Brien (1988) and Lim (2001) used a similar rule in deleting suspected data-entry errors.

If we require a lag of one day in calculating the outstanding consensus, then the last column will represent our notion of the *order* of release of the forecasts. For example, since the same analyst makes the first two forecasts, he is still the only analyst, so the order is 1. This means that the *consensus* forecast at the time analyst 2 reports his forecast on April 27 (see the 3rd observation) will be given by $0.70, the most recent forecast of analyst 1.

Also, since the 4th and 5th forecasts were made by different analysts on the same day, they are both 3rd in the *order* of release of forecasts. That is, both of them have knowledge of only two (distinct) analysts who have reported before them. Hence, they will both be facing the same consensus, which will be the average of the 2nd and 3rd forecasts. Notice that the 7th forecast (release by analyst 3) has an order of 4. This is because his forecast was made on the same day as the 6th observation (reported by analyst 5) and 3 *distinct* analysts had reported forecasts before May 15. Similarly, the 8th observation, reported by analyst 1 will be 5th in the *order* of release. The outstanding consensus at that point will then be the average of the most recent forecasts reported by all except analyst 1; that is, the mean of the 7th, 6th, 5th, and 3rd forecasts.

An important question is what analysts see or know about other forecasts when they release their own forecasts. The date of publication of an analyst's forecast in the I/B/E/S database could be later than the actual release date of his forecast to the market. Since it is unclear whether an analyst making a forecast will know of recent forecasts, we choose a lag of $\ell = 3$ days. Alternative lag choices of 1, 5 or 7 days yield qualitatively identical findings.

Table 2: Sample Statistics: # of obs $= 320,311$

			Percentiles				
	min	max	25	50	75	Mean	Std. dev
Number of analysts (`cover`)	2	38	6	9	14	10.12	5.75
Forecasts-Analysts ratio	1	4	1.14	1.33	1.57	1.39	0.34
Consensus *error* (as a percent of stock price)	-48.5	160.4	-0.07	0.001	0.18	0.23	2.0
Forecast *error* (as a percent of stock price)	-72.6	159.4	-0.09	0	0.11	0.14	1.87
Forecast minus *consensus* (as a percent of stock price)	-100.58	59.80	-0.10	-0.01	0.04	-0.08	0.94
Days between *any* two forecasts	0	116	1	6	13	9.58	12.46
Days between *own* forecasts (revised forecasts only)	0	117	21	36	56	38.59	22.08

Table 2 gives the distribution of analyst coverage and the average number of forecasts per analyst. An average of about 10 analysts follow a firm in a given quarter in our sample, but some firm-quarters have as many as 38 analysts reporting at least one forecast. The average number of forecasts per analyst is about 1.4, indicating that one-quarter ahead forecasts are not revised too often. On average, there are about 12 days between the release dates of successive forecasts.

We also report the distribution of *error* in individual forecasts as a percent of the stock's price as at the end of the previous quarter, and the difference between the last and consensus forecasts. To remove the effects of outliers in percentage errors and differences, we drop observations with stock price less than \$5.00 (see Lim 2001). The mean stock price in the resulting sample is about \$35.00, suggesting that analysts on average overshot earnings in the direction away from the consensus by about \$0.05. The summary statistics show that the distribution of the percentage forecasts error has extremely fat tails. Such outliers may have excessive impact on estimates using mean regressions. We address this problem in section 3.3.

3.2 Results of Frequency Tests

Let $\mathbf{S}(x)$ be our test statistic when z^+ is the event that $\frac{F_\tau - \overline{F}_{\tau-\ell}}{P_{\tau-1}} > x$, and z^- is the event that $\frac{\overline{F}_{\tau-\ell} - F_\tau}{P_{\tau-1}} > x$. The results, reported in Table 3, show that $\mathbf{S}(0)$ equals 0.60. That is, 60% of the time, analysts overshoot earnings in the direction away from the consensus. Figure 2 shows that $\mathbf{S}(x)$ rises rapidly with x: larger deviations from the consensus substantially raise the likelihood of overshooting. Thus, we overwhelmingly reject both the null of unbiasedness and the alternative of herding. Rather, analysts issue forecasts that overemphasize their private information.

Analysts' forecasts exhibit a contrarian bias *both* when they have positive information and when they have negative private information relative to the consensus (see the last 3 columns). Even though unconditionally, forecasts exceed earnings only 46% of the time, if a forecast exceeds the consensus, it exceeds earnings 58% of the time. So, too, if a forecast is less than the consensus, it falls short of earnings 62% of the time. The difference in overshooting rates (62% vs 58%) almost certainly reflects the fact that, unconditionally, forecasts tend to undershoot earnings (see also Abarbanell and Lehavy 2001). Note, however, that if analysts issued unbiased forecasts that corresponded to their posterior beliefs about the m-th percentile of earnings (m arbitrary), then this has offsetting impacts on our two estimates of the conditional overshooting frequencies, leaving the mean of our test statistic $\mathbf{S}$ under the null unchanged at 0.5, but with a reduced variance. Thus, pessimism *cannot* explain our results.

In short, our tests are robust to the features of the data highlighted by Abarbanell and Lehavy. Indeed, to the extent that (a) forecasts are more likely to be pessimistic, or (b) analysts are

Table 3: Frequency Test of Bias in Analysts's Forecasts

Average conditional probability, **S**, in different sub-samples. z_τ^+ is the event that the forecast exceeds the *consensus*, and z_τ^- is the event that the forecast is less than the *consensus*. 95% confidence intervals in square brackets.

	Freq. Test Statistic		Conditional Frequencies		
Sample	# obs.	$\mathbf{S}(0)$	$\Pr(F_\tau > E)$	$\Pr(F_\tau > E \mid z_\tau^+)$	$\Pr(F_\tau < E \mid z_\tau^-)$
Total sample	272,872	0.60	0.464	0.579	0.622
		[0.599, 0.602]			
1989 – 1992	88,890	0.607	0.551	0.673	0.542
		[0.604, 0.611]			
1993 – 1995	82,180	0.603	0.440	0.555	0.650
		[0.599, 0.606]			
1996 – 1999	101,802	0.590	0.406	0.514	0.667
		[0.587, 0.593]			
Segmented by Analysts following (`cover`)					
`cover` ≤ 5	60,672	0.605	0.450	0.571	0.639
		[0.601, 0.609]			
$6 \leq$ `cover` ≤ 14	150,248	0.605	0.470	0.590	0.620
		[0.602, 0.607]			
`cover` ≥ 15	61,952	0.586	0.459	0.561	0.611
		[0.581, 0.590]			
Segmented by timing of forecast (`order`)					
`order` $= 3$	34, 864	0.609	0.479	0.599	0.619
		[0.603, 0.614]			
`order` = 2nd-to-last	49,855	0.593	0.431	0.541	0.648
		[0.589, 0.598]			
`order` = last	58,339	0.592	0.421	0.528	0.657
		[0.588, 0.596]			
Segmented by timing `order` in small/large group `cover`					
`order` $= 2$ if `cover` ≤ 5	23,889	0.617	0.480	0.608	0.626
		[0.611, 0.623]			
`order` = *last* if `cover` ≤ 5	29,875	0.598	0.438	0.550	0.648
		[0.592, 0.604]			
`order` $= 2$ if $6 \leq$ `cover` ≤ 14	11,965	0.613	0.544	0.661	0.564
		[0.604, 0.622]			
`order` = *last* if $6 \leq$ `cover` ≤ 14	24,045	0.587	0.407	0.509	0.664
		[0.581, 0.593]			
`order` $= 2$ if `cover` ≥ 15	1,949	0.581	0.567	0.652	0.509
		[0.559, 0.603]			
`order` = *last* if `cover` ≥ 15	4,419	0.581	0.377	0.469	0.685
		[0.565, 0.596]			
Segmented by days between forecasts (`days`)					
$3 \leq$ `days` ≤ 7	80,525	0.611	0.467	0.591	0.631
		[0.607, 0.614]			
`days` ≥ 8	99,074	0.609	0.448	0.569	0.648
		[0.606, 0.612]			
Segmented by days prior to earnings announcement (`to-ann`)					
`to-ann` ≤ 15	41,981	0.583	0.385	0.484	0.683
		[0.579, 0.588]			
$15 <$ `to-ann` ≤ 30	46,582	0.570	0.412	0.497	0.642
		[0.565, 0.574]			
$30 <$ `to-ann` ≤ 60	86,254	0.595	0.477	0.588	0.601
		[0.591, 0.598]			
`to-ann` > 60	98,055	0.621	0.511	0.640	0.602
		[0.618, 0.624]			

forecasting a variable other than what we observe, or (c) there are unforecasted earnings shocks, the power of our test is reduced, making it harder to reject the null of unbiasedness. This renders the magnitude of overshooting bias that we uncover all the more remarkable.

Table 3 also reveals how the likelihood of overshooting varies with conditioning variables:

- The probability of overshooting is slightly higher for firms followed by fewer analysts. For example, the **S** statistic equals 0.61 for firms followed by 5 analysts or less, while it equals 0.59, for firms followed by at least 15 analysts.
- Earlier analysts are slightly more likely to overshoot. For example, the third analyst overshoots 61% of the time, while the last analyst overshoots 59% of the time.
- The second analyst is less likely to overshoot if the firm is covered by more analysts; but overshooting rates by the last analyst are not significantly affected by analyst coverage.
- There is almost no variation in overshooting rates across forecast-year, or as a function the number of days between forecasts. However, the overshooting bias is slightly higher for forecasts that are reported over 60 days prior to the earnings announcement.
- **S** varies little with the *order* of release for firms followed by many analysts, even though the unconditional probability that a forecast exceeds earnings falls for later forecasts.
- The conditional probability that a forecast falls short of earnings when it falls short of the consensus varies from a high of 66% in the late 1990s, when earnings were unexpectedly high, to a low of 54% in 1989 and early 1990s, when the economy was in an unexpected downturn. These results document both the importance of designing a test statistic that is robust to common unforecasted shocks, and our success in doing so.

Next we explore the bias in *revised* forecasts. Analysts are far more likely to revise forecasts down than up — 66.5% of the time, the revised forecast is lower than the original forecast, suggesting that revised forecasts are strategically chosen. Bernhardt and Campello (2001) provide evidence that these downward forecast revisions may be induced by firms seeking to reduce the consensus forecast, in order to generate a 'positive earnings surprise.'

We decompose revised forecasts in two ways. We first decompose revised forecasts according to whether the original forecast-consensus difference is *smaller* or *larger* than the forecast-consensus difference after the revision. We also decompose revised forecasts according whether the revised forecast, F^r_τ, is closer to the most recent consensus than the initial forecast, F^o_τ, was: *i.e.*, according to whether $|F^r_\tau - \overline{F}^r_\tau| < |F^o_\tau - \overline{F}^r_\tau|$, where $\overline{F}^r_\tau$ is the revised consensus.

On first pass, Table 4 might suggest that revised forecasts are less biased — revised forecasts overshoot earnings only 54% of the time, far less than the 60% rate at which all forecasts overshoot. Importantly, revised forecasts are *not* less biased due to the fact they are issued later in the forecast cycle: Table 3 indicates that the *last* analyst to report is only slightly less biased than earlier analysts, overshooting 59% of the time. Inspection of the last panel reveals that (a) if a forecast was revised, then the original forecast was awful, vastly overshooting earnings, and (b) revised forecasts are significantly better. This suggests that analysts are reluctant to revise forecasts.

Table 4: Test of Bias when Analyst Revises Forecasts

This Table reports average conditional probability, **S**, that a revised forecast overshoots earnings, in different sub-samples. 95% confidence intervals in square brackets.

			Conditional Frequencies		
Sample	# obs.	**S**	$\Pr(F_\tau > E)$	$\Pr(F_\tau > E \mid z^+_\tau)$	$\Pr(F_\tau < E \mid z^-_\tau)$
Revised forecasts	81,494	0.539 [0.535, 0.542]	0.446	0.496	0.582
Forecast-consensus difference after revisions					
Difference smaller (than before) $(\lvert F^r_\tau - \overline{F}^r_\tau\rvert < \lvert F^o_\tau - \overline{F}^o_\tau\rvert)$	44,871	0.549 [0.544, 0.553]	0.442	0.498	0.599
Difference larger (than before) $(\lvert F^r_\tau - \overline{F}^r_\tau\rvert > \lvert F^o_\tau - \overline{F}^o_\tau\rvert)$	36,623	0.528 [0.522, 0.535]	0.452	0.492	0.564
"Location" of old & revised forecast, relative to current consensus					
Revision *closer* to current consensus: $(\lvert F^r_\tau - \overline{F}^r_\tau\rvert < \lvert F^o_\tau - \overline{F}^r_\tau\rvert)$	39,979	0.568 [0.563, 0.573]	0.450	0.525	0.611
original forecast		0.845 [0.841, 0.850]	0.562	0.853	0.837
Revision *further* from current consensus $(\lvert F^r_\tau - \overline{F}^r_\tau\rvert > \lvert F^o_\tau - \overline{F}^r_\tau\rvert)$	41,515	0.505 [0.500, 0.511]	0.443	0.451	0.560
original forecast		0.571 [0.566, 0.575]	0.609	0.724	0.417

The results when we decompose forecast revisions further reveal that the bias characterization is more subtle than that suggested by the fact that, on average, revised forecasts are less likely to overshoot. The second panel shows that revisions that raise the difference between the consensus forecast and the analyst's forecast are *less* likely to overshoot than revisions that reduce the forecast–consensus difference. Analyst overconfidence predicts the opposite.

Most startlingly, the last panel of Table 4 reveals that forecast revisions that move *toward* the consensus do not move nearly far enough, so that they still overshoot frequently. In sharp contrast, forecast revisions that move away from the consensus are essentially unbiased. The result that, in aggregate, revised forecasts overshoot less reflects the averaging of these opposing results. Most importantly, the fact that revisions toward the consensus are *more* biased than those that move further away strongly suggests that analyst overconfidence does not underlie the forecast bias.

How might one interpret these results? One possibility is that analysts care about both absolute forecast accuracy, and forecast accuracy relative to the consensus. Consider the finding that revised forecasts which move away from the consensus are unbiased. If the consensus is sufficiently inaccurate given an analyst's updated information, then he may not want to bias his forecast further in order to separate away from the consensus. The revised forecast moves away from the consensus, but is unbiased. This is consistent with the possibility that after some point, there are decreasing returns to outperforming the consensus by more. Now consider the extreme overshooting bias in forecasts that move toward the consensus. An analyst who revises his forecast toward the consensus only does so when his original forecast is far 'too' extreme given his updated information. The fact that the analyst's revised forecast does not go far enough toward the consensus suggests a desire to distinguish his forecast from others.

3.3 Economic Significance of Forecast Bias

Our frequency results demonstrate that analysts, rather than herd towards the consensus forecast, systematically attempt to distinguish themselves from other analysts by reporting biased contrarian forecasts. However, these findings reveal limited information about the magnitude of the forecast bias. We now estimate the size of the forecast bias.

We first estimate the relationship between forecast error and the difference between the forecast and outstanding consensus. Since earnings are reported on a per share basis, we express the error and difference as a percent of share price at the end of the previous quarter. Using the previous quarter share price removes the contemporaneous effect of recent forecasts on the stock's price. The forecast error is then

$$\text{Error}_\tau = \frac{F_\tau - E_\tau}{P_{\tau-1}}.$$

The difference between the forecast and the consensus is

$$\text{SFD}_\tau = \frac{F_\tau - \overline{F}_{t-1,\tau}}{P_{\tau-1}}.$$

Our analysis uses the following empirical model:

$$\begin{aligned}\text{Error}_\tau &= \alpha_0 + \sum_i \text{firm}_i + \alpha_1\,\text{SFD}_\tau + \alpha_2(\texttt{lord}) + \alpha_3(\texttt{lcov}) + \alpha_4(\texttt{lord})\,\text{x}\,(\texttt{lcov}) \\ &+ \alpha_5(\texttt{lord})\,\text{x}\,\text{SFD}_\tau + \alpha_6(\texttt{lcov})\,\text{x}\,\text{SFD}_\tau + \alpha_7(\texttt{lord})\,\text{x}\,(\texttt{lcov})\,\text{x}\,\text{SFD}_\tau + \epsilon_\tau \qquad (2)\end{aligned}$$

Under the null hypothesis that an analyst's forecast is unbiased, one can interpret the dependent variable $\text{Error}_\tau = \frac{F_\tau - E_\tau}{P_{\tau-1}}$ as a price-normalized forecast error if the analyst ran a regression using all available information to obtain his best forecast. Were forecasts unbiased, then all coefficient estimates should be zero, since the forecast error would be orthogonal to everything in the analyst's information set, including the independent variables included in this regression. To see whether the forecast error, and hence forecast bias, varies with analyst-coverage and/or information at an analyst's disposal, we include proxies for the timing of forecasts, using the log of `order`, *lord*, and the log of the number of analysts following the firm, *lcov*. To understand why we do this, notice that if the consensus contains more information, but an analyst ignores this when making his forecast, then his error will be larger for every percentage deviation from the consensus. Then, provided that a greater analyst following implies a more precise consensus, the average forecast error should rise with `lord`, `lcov`, and `lord x lcov`. In contrast, strategic agents who care about relative forecast performance will introduce a smaller bias to their forecast if there is less uncertainty about earnings, so that the average forecast error should fall with `lord`, `lcov` and `lord x lcov`.

A *significant* concern is the massively fat tails and left-skewness in the distribution of forecast errors in the initial sample (Table 5): outliers in the tail, especially those in the left tail, will have an excessive impact on estimates, giving rise to misleading results. It is also the case that most of the extreme left tail observations reflect large writeoffs of asset values that analysts may not be trying to predict, and hence should not be included in the sample (Abarbanell and Lehavy 2001).

We address these concerns in two ways. First, we drop observations in the extreme tails of the distribution of the *consensus* error and check the degree of skewness after each trim. Table 5 provides sample moments after successive trims of 0.5%, 1%, 2.5% and 5% of the tails in the distribution of the *consensus* error. We then estimate OLS firm fixed-effects regressions using the trimmed samples.[6] For these regressions we present both Huber-White's robust standard errors, which account for possible correlation of forecast errors within firms (Rogers 1993), and the bootstrapped 95% confidence intervals for all coefficient estimates. The bootstrapping procedure consists of repeating each of the regressions 100 times, with the sampling restricted to one observation per individual firm each time. These confidence bounds are, therefore, not influenced by

[6]Inspection reveals that a small 0.5% trim of the tails leaves the distribution of forecast errors (the dependent variable) relatively symmetric. We report results for all of the sub-samples, but only interpret results for the 0.5% trim in the text. See Lim (2001) for a similar approach.

standard arguments that within-firm error autocorrelation in panel regressions may 'too often' lead to the rejection of the null hypothesis (Type I error). Finally, we estimate equation (2) with quantile (*i.e.*, median) regression using the entire sample. An advantage of median regression is that, in contrast to ordinary least squares regression, its estimation procedure minimizes the mean absolute value of deviations, thereby producing estimates that are less sensitive to outliers. A disadvantage is that the large number of firms in our sample precludes us from allowing for fixed-firm effects in our median regression (private correspondence with Roger Koenker).

Structural interpretations of the parameter estimates can be obtained under the assumption that analysts choose to introduce a forecast bias that is a linear function of the difference between his (unobserved) true posterior estimate of earnings and the consensus forecast, so that

$$F_\tau = \hat{E}_\tau + a_{n,\tau}(\hat{E}_\tau - \overline{F}_\tau),$$

where $\hat{E}_\tau$ is the analyst's true posterior estimate of earnings in a firm-quarter τ, given all available information at time t, the time that analyst issued his forecast. Here, $\hat{E}_\tau$ corresponds to the median of the analyst's posterior distribution over earnings for our median regressions, and to the mean of his posterior for our mean regressions. We index the bias ($a_{n,\tau}$) because it should rise with the amount of uncertainty that the analyst faces about earnings (and hence fall with the number of other forecasts/information at his disposal).

The difference between the forecast and consensus forecast, as a function of this bias, is

$$\begin{aligned} F_\tau - \overline{F}_\tau &= \hat{E}_\tau + a_{n,\tau}(\hat{E}_\tau - \overline{F}_\tau) - \overline{F}_\tau) \\ &= (1 + a_{n,\tau})(\hat{E}_\tau - \overline{F}_\tau), \end{aligned}$$

so that

$$(\hat{E}_\tau - \overline{F}_\tau) = \frac{F_\tau - \overline{F}_\tau}{1 + a_{n,\tau}}. \tag{3}$$

We use this relationship to express the dependent variable in our regressions as the sum of a true (unobserved) forecasting error for an analyst, $\frac{\hat{E}_\tau - E_\tau}{P_{\tau-1}}$, which should be orthogonal to everything in the last analyst's information set, plus a bias term that can be written in terms of the difference between the forecast and the consensus:

$$\text{Error}_\tau = \frac{\hat{E}_\tau - E_\tau}{P_{\tau-1}} + \left(\frac{a_{n,\tau}}{1 + a_{n,\tau}}\right) SFD_\tau. \tag{4}$$

Subtracting $\left(\frac{a_{n,\tau}}{1+a_{n,\tau}}\right) SFD_\tau$ from both sides of our regression, the left hand side becomes a normalized regression error from the analyst's forecasting regression which is orthogonal to everything

Table 5: Distribution of Forecast Errors

This table reports the moments of distribution of forecast errors, after trimming tails to eliminate outliers and skewness in initial sample

	Percentiles			Sample moments			
	25	50	75	Mean	Std. dev.	Skewness	Kurtosis
No trim: $N = 320,311$							
Consensus *error* (as a percent of stock price)	-0.07	0.001	0.18	0.23	2.0	24.48	1212
Forecast *error* (as a percent of stock price)	-0.09	0	0.11	0.14	1.87	24.59	1357
Forecast minus *consensus* (as a percent of stock price)	-0.10	-0.01	0.04	-0.08	0.94	-19.71	1320
0.5% trim: $N = 317,108$							
Consensus *error* (as a percent of stock price)	-0.07	0	0.18	0.16	0.63	4.37	33
Forecast *error* (as a percent of stock price)	-0.09	0	0.11	0.09	0.80	1.06	190
Forecast minus *consensus* (as a percent of stock price)	-0.10	-0.007	0.037	-0.07	0.59	-10.6	656.5
1% trim: $N = 313,904$							
Consensus *error* (as a percent of stock price)	-0.07	0	0.18	0.14	0.63	3.27	19.97
Forecast *error* (as a percent of stock price)	-0.09	0	0.11	0.07	0.66	-1.36	310.3
Forecast minus *consensus* (as a percent of stock price)	-0.10	-0.007	0.036	-0.07	0.52	-11.4	857.1
2.5% trim: $N = 304,296$							
Consensus *error* (as a percent of stock price)	-0.065	0	0.16	0.11	0.43	2.19	10.21
Forecast *error* (as a percent of stock price)	-0.085	0	0.10	0.05	0.50	-5.34	647.8
Forecast minus *consensus* (as a percent of stock price)	-0.09	-0.007	0.034	-0.06	0.42	-12.9	1310.6
5% trim: $N = 288,280$							
Consensus *error* (as a percent of stock price)	-0.058	0	0.15	0.08	0.30	1.67	6.63
Forecast *error* (as a percent of stock price)	-0.078	0	0.09	0.037	0.40	-11.1	1487.4
Forecast minus *consensus* (as a percent of stock price)	-0.08	-0.006	0.032	-0.04	0.35	-18.4	2396.3

in the analyst's information set. Since the sum of the coefficients on $SFD_{t,\tau}$,

$$\beta_{n,\tau} = \alpha_1 + \alpha_5(\texttt{lord}) + \alpha_6(\texttt{lcov}) + \alpha_7(\texttt{lord}) \cdot (\texttt{lcov})$$

provides an estimate of the bias $\frac{a_{n,\tau}}{1+a_{n,\tau}}$, an unbiased estimate of the strategic bias coefficient when there are n analysts is

$$\hat{a}_{n,\tau} = \frac{\beta_{n,\tau}}{1-\beta_{n,\tau}}.$$

So, too, we can back out the expected dollar bias in the forecast as $\beta_{n,t}(F_\tau - \overline{F}_\tau)$.

Table 6 reports the findings. Note the high R^2 associated with our regressions, especially when we increase the trim slightly to eliminate gross outliers. The coefficient estimates that we care about are those on `SFD`, and its interactions with `lord` and `lcov`. Note that `lord` is bounded from above by `lcov`. Forecast bias clearly rises with `SFD`. However, the interactions make it difficult to read off both the magnitude of the forecast bias, and how the estimated forecast error varies with analyst order. Consequently, we consider the estimated impacts in the context of a firm followed by 20 analysts. Using the estimates from the OLS fixed-effects regression where tails are trimmed at 0.50%, we find that if 20 analysts follow a firm in a given quarter and the second analyst reports a forecast that exceeds the first analyst's by one standard deviation of the mean forecast–consensus difference, then, on average, his forecast overshoots actual earnings by about 1.02% of the stock price, *ex post*.[7] However, for the same difference between the forecast and the outstanding consensus, the 20^{th} analyst to report a forecast, on average, overshoots earnings by only 0.58% of the stock's price.

Using our structural estimates, these imply an approximate bias of 1.7 times the difference between the analyst's (unobserved) true estimate of earnings and the consensus forecast if he were the 2^{nd} to report (out of 20 analysts following the firm); while if he reports last (the 20^{th}), then the implied bias is approximately 0.44 times. If we trim the tails of the *consensus–forecast error* distribution by more, the predicted forecast bias rises, as does the forecast bias by analysts who forecast earlier, relative to those who go later. The implied biases are somewhat smaller using our median regression estimates, especially for analysts who report later: The implied bias from the median regression is 0.62 when an analyst reports 2^{nd}, and 0.01 for the 20^{th} analyst to report.

We now forcefully document that our results cannot be explained simply by analyst overconfidence. Recall that, in its most extreme form, overconfidence by the j^{th} analyst would imply average overshooting magnitudes of the order of $1 - \frac{1}{j}$ of a percent; and that this extreme bound requires

[7] The mean forecast–consensus difference, SFD, in our sample is -0.08% (of stock price) and the standard deviation is 0.94 (See Table 2).

Table 6: Economic Significance of Bias in Analysts' Forecasts

OLS firm fixed-effects regressions are estimated for total sample, and after triming tails in distribution of *consensus* error. For the OLS regressions, robust standard errors (in parentheses) were computed using the Huber-White method; and we also report (in square brackets) the 95% confidence interval of parameter estimates obtained by bootstrapping: (*), (**) and (***) indicate significance at 90%, 95% and 99% levels.

Dependent variable: $\text{Error}_\tau = \frac{F_\tau - E_\tau}{P_{\tau-1}}$

Regression	SFD	$lord$	$lcov$	$lord * lcov$	$lord * SFD$	$lcov * SFD$	$lord * lcov * SFD$
OLS-FE (0.5% trim)	0.277	−0.085	0.188	0.006	0.403	0.150	−0.181
Obs.: 317,111	(0.2124)**	(0.026)***	(0.016)***	(0.010)	(0.117)***	(0.064)**	(0.048)***
R^2: 0.163	[0.05, 0.54]	[−0.13, −0.04]	[0.15, 0.22]	[−0.01, 0.02]	[0.15, 0.66]	[0.01, 0.27]	[−0.27, −0.07]
OLS-FE (1% trim)	0.481	−0.078	0.171	0.006	0.248	0.097	−0.123
Obs.: 313,904	(0.121)***	(0.020)***	(0.013)***	(0.001)	(0.121)**	(0.059)*	(0.045)***
R^2: 0.277	[0.24, 0.66]	[−0.11, −0.04]	[0.14, 0.21]	[−0.007, 0.02]	[0.02, 0.47]	[−0.02, 0.22]	[−0.23, −0.04]
OLS-FE (2.5% trim)	0.542	−0.074	0.140	0.009	0.285	0.083	−0.128
Obs.: 304,296	(0.099)***	(0.016)***	(0.010)***	(0.006)	(0.111)***	(0.046)*	(0.041)***
R^2: 0.390	[0.36, 0.75]	[−0.11, −0.04]	[0.12, 0.16]	[−0.002, 0.02]	[0.04, 0.48]	[−0.03, 0.17]	[−0.20, −0.03]
OLS-FE (5% trim)	0.671	−0.071	0.106	0.014	0.261	0.039	−0.106
Obs.: 288,280	(0.086)***	(0.012)***	(0.007)***	(0.004)***	(0.098)***	(0.037)	(0.036)***
R^2: 0.532	[0.51, 0.84]	[−0.09, −0.05]	[0.09, 0.12]	[0.005, 0.02]	[0.02, 0.43]	[−0.03, 0.11]	[−0.16, −0.02]
Quantile (median)	0.186	−0.010	0.015	−0.001	−0.061	0.103	−0.034
Obs.: 320,311	(0.001)***	(0.0008)***	(0.0005)***	(0.0003)***	(0.001)***	(0.0005)***	(0.0003)***
pseudo-R^2: 0.086							

Table 7: Overshooting Magnitude and the Over–confidence Hypothesis

This table reports the coefficient on SFD returned from on OLS fixed effects (2.5regression of Error on SFD conditional on the order of the analyst report. Standard errors are Huber-White adjusted. We also report 95% confidence interval of parameter estimates obtained via bootstrapping.

j^{th}Analyst	Coefficient on SFD_τ	Std Error	95% conf. int.	R^2
2^{nd}	0.728	0.022	$[0.685, 0.771]$	0.402
3^{rd}	0.757	0.040	$[0.686, 0.827]$	0.399
4^{th}	0.780	0.047	$[0.684, 0.876]$	0.421
5^{th}	0.810	0.081	$[0.645, 0.974]$	0.497
6^{th}	0.678	0.029	$[0.624, 0.733]$	0.390
7^{th}	0.779	0.044	$[0.692, 0.868]$	0.470
8^{th}	0.640	0.042	$[0.566, 0.714]$	0.283
9^{th}	0.708	0.052	$[0.610, 0.806]$	0.336
10^{th} and higher	0.576	0.036	$[0.510, 0.648]$	0.254

that analysts ignore all information in other forecasts, that signals not be positively correlated, and that analysts who forecast later in the forecasting cycle not have better information. For the second analyst to report, this implies an upper bound on overshooting of 0.50%. We show that this bound is exceeded by regressing Error_τ on SFD_τ for the second analyst to report. Because the bound on overconfidences *does not* depend on analyst following, we can test it directly. As a result, we do not have to worry about whether our findings are driven by functional form. This regression reveals that for every one percent of price that the second forecast overshoots the first, it overshoots earnings on average by 0.73%. This difference exceeds 0.50% at a better than 0.001% test level, implying that we can confidently reject all overconfidence explanations. Table 7 reveals overshooting rates for analogous regressions for later analysts in the forecasting order. Overshooting rates are so high that we can reject all overconfidence explanations for the third, fourth and fifth analyst to report. For subsequent analysts, given any reasonable assumptions on signal correlation and the timing of information arrival, overshooting rates are so high that it is implausible for them to be consistent with any reasonable overconfidence hypothesis.

Summarizing our findings, analysts introduce substantial overshooting biases into their forecasts, the bias is less for later analysts, and it is less if analyst coverage is greater. The coefficient estimates support the argument that analysts strategically issue contrarian forecasts, to try to distinguish themselves from others, especially in an uncertain forecasting environment. In particular, the more by which a forecast overshoots the consensus, the more by which it is predicted to overshoot earnings, but the estimated amount by which the analyst overshoots falls with the amount of information at his disposal. The large magnitudes of the estimated overshooting forecast bias also reveal that they are not due to overconfident analysts underweighting the information in the forecasts of others.

4 Conclusion

This paper develops a robust frequency test for forecast bias. Detecting forecast bias in the securities industry is difficult because analysts rely on similar sources of information and are surprised by the same events. Our test is designed to be robust to correlated signals among analysts, common unforecasted shocks to earnings, and information arrival. This test can also be usefully employed to detect bias in the forecasts of macroeconomic variables. There, common data sources, and similar econometric models again imply that correlated signals are a significant concern to overcome when attempting to unravel forecast bias.

We find overwhelming evidence that earnings forecasts are biased. Analysts do not herd, but rather seek to separate from other analysts by issuing forecasts that over-emphasize their private information relative to the consensus. Analysts appear to issue strategically strongly contrarian forecasts, biasing forecasts *away* from the consensus forecast — 60% of the time, a forecast overshoots actual EPS in the direction away from the outstanding mean forecast.

Our structural estimates reveal that the economic magnitude of the bias in forecasts is significant. For example, on average, the second analyst (out of 20) biases his forecast by about 1.7 times the difference between the analyst's (unobserved) true estimate of earnings and the consensus forecast, while the last analyst to report has an implied bias of about 0.44 times the difference. Thus, the bias falls as uncertainty around earnings is resolved, but its magnitude remains large.

Both our frequency results and our structural estimates are consistent with the hypothesis that analysts strategically bias their forecasts to try to separate from the forecasts of others; and are inconsistent with the premise that forecast biases are driven by analysts' failure to incorporate useful information into their forecasts.

References

[1] Abarbanell, Jeffery, and Reuven Lehavy (2001), "Biased Forecasts or Biased Earnings? The Role of Reported Earnings in Explaining Apparent Bias and Over/Underrrecation in Analysts' Earnings Forecasts," *mimeo*, Haas School of Business, University of California, Berkeley.

[2] Bernhardt, Dan, and Edward Kutsoati (2001), "Analyst's Compensation and Forecasts: Theory, Tests and Evidence," *mimeo*, University of Illinois.

[3] Bernhardt, Dan, and Murillo Campello (2001), "The Dynamics of Earnings Management and Analyst Forecast Errors," *mimeo*, University of Illinois.

[4] Ehrbeck, Tilman, and Robert Waldmann (1996), "Why Are Professional Forecasters Biased? Agency versus Behavioral Explanations," *Quarterly Journal of Economics* 111, 41-64.

[5] Gallo, Giampiero M., Clive W. Granger, and Yongil Jeon (2001), "The Impact of the Use of Forecasts in Information Sets," *mimeo*, University of California, San Diego.

[6] Hong, Harrison, Jeffrey Kubik and Amit Solomon (2000), "Security Analysts' Career Concerns and Herding of Earnings Forecasts," *Rand Journal of Economics* 31, 121-44.

[7] I/B/E/S Research Bibliography (1996), "Annotated Bibliography of Earnings Expectation Research," Institutional Brokers Estimate System, Inc., 5th Edition.

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

[9] Lamont, Owen (1995), "Macroeconomic Forecasts and Microeconomic Forecasters," NBER Working paper # 5284.

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

[11] Lim, Terence (2001), "Rationality and Analysts' Forecasts Bias," *Journal of Finance* 56, 369-85.

[12] O'Brien, Patricia (1988), "Analysts' Forecasts as Earnings Expectations," *Journal of Accounting and Economics* 10, 53-83.

[13] Odean, Terrance (1999), "Do Investors Trade Too Much?", *American Economic Review,* 89, 5, 1279–1298

[14] Rogers, W. H. (1993), "Regression Standadrd Errors in Clustered Samples," STATA Technical Bulletin 13, 19-23.

[15] Trueman, Brett (1990), "On the Incentives for Security Analysts to Revise their Earnings Forecasts," *Contemporary Accounting Research* 7, 203-222.

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

[17] Welch, Ivo (2000), "Herding Among Security Analysts," *Journal of Financial Economics* 58, 369-96.

[18] Zitzewitz, Eric (2001), "Measuring Exaggeration by Analysts and Other Opinion Producers," *mimeo,* Stanford University, Graduate School of Business.

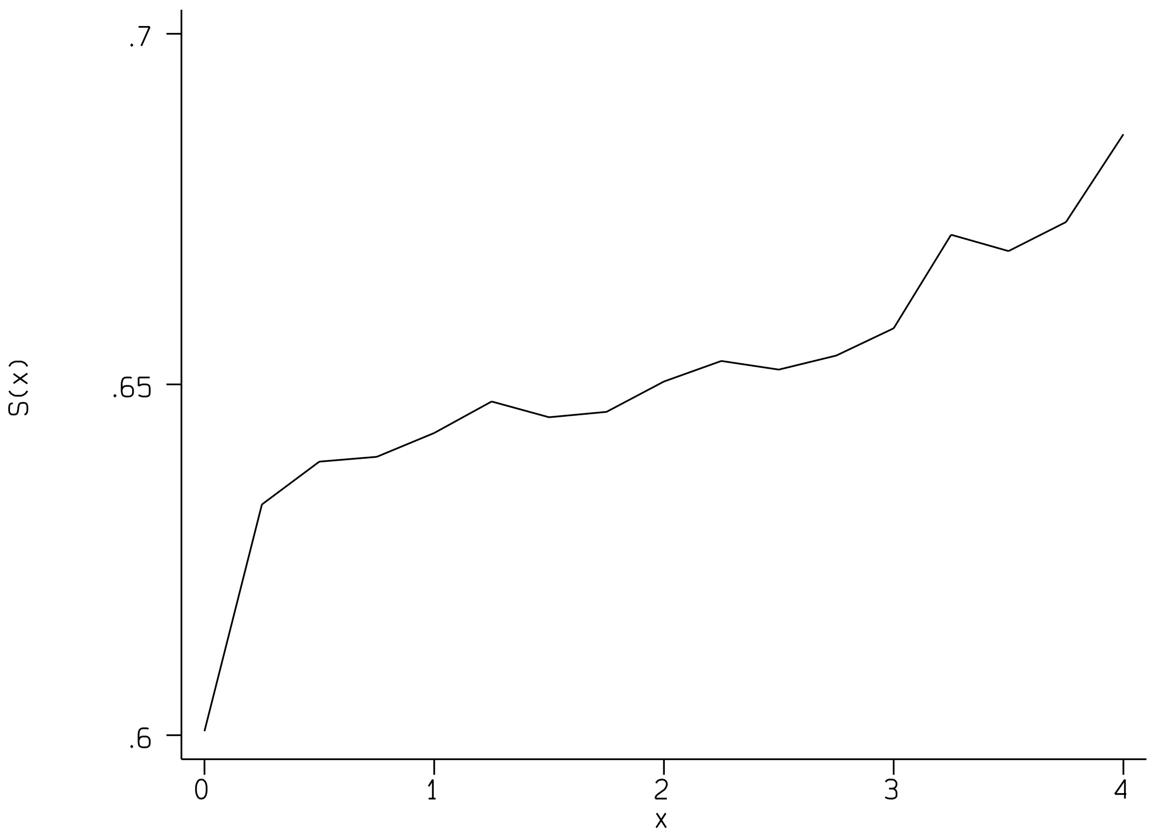

Figure 2: Average conditional probability that forecast overshoots earnings away from consensus, $\mathbf{S}(x)$, given that it exceeds or falls short of the consensus by at least $x\%$ of the stock's price. To avoid the impact of possible data entry errors, we drop observations for which a forecast exceeds or falls short of the consensus by more than $x = 20\%$ of the stock's price.

WORKING PAPER SERIES 2002

ase.tufts.edu/econ/papers

2002-01 FULLERTON, Don and Gilbert METCALF. The Distribution of Tax Burdens.

2002-02 EGLLESTON, Karen and Richard ZECKHAUSER. Government Contracting for Health Care.

2002-03 EGLLESTON, Karen, Robert JENSEN, and Richard ZECKHAUSER. Information and Communication Technologies, Markets and Economic Development.

2002-04 EGGLESTON, Karen, George NORMAN, and Lynne PEPALL. Managed Health Care and Provider Integration: a Theory of Bilateral Market Power.

2002-05 CHISHOLM, Darlene C. and George NORMAN. Heterogeneous Preferences and Location Choice with Multi-Product Firms.

2002-06 BROWN, Drusilla K., Alan V. DEARDORFF, and Robert M. STERN. Pros and Cons of Linking Trade and Labor Standards.

2002-07 BROWN, Drusilla K., Alan V. DEARDORFF, and Robert M. STERN. Impacts on NAFTA Members of Multilateral and Regional Trading Arrangements and Initiatives and Harmonization of NAFTA's External Tariffs.

2002-08 IOANNIDES, Yannis M. and Jeffrey E. ZABEL. Interactions, Neighborhood Selection and Housing Demand.

2002-09 METCALF, Gilbert E. and George NORMAN. Oligopoly Deregulation and the Taxation of Commodities.

2002-10 METCALF, Gilbert E. and George NORMAN. Oligopoly Deregulation in General Equilibrium: A Tax Neutralization Result.

2002-11 BABIKER, Mustafa H., Gilbert E. METCALF, and John REILLY. Tax Distortions and Global Climate Policy.

2002-12 BERNHARDT, Dan, Eric HUGHSON, and Edward KUTSOATI. Survival of the Unfittest: How Dodos Become Managers.

2002-13 BERNHARDT, Dan, Murillo CAMPBELLO, and Edward KUTSOATI. Who Herds?