Analyst compensation and forecasts: theory, tests and evidence*

Dan Bernhardt† Edward Kutsoati‡

May 17, 2001

Abstract

This paper derives and tests the implications of relative performance incentives for security analysts. If, in addition to compensation for absolute forecast accuracy, analysts are compensated on the basis of how their forecasts compare with those of other analysts, then earlier forecasts affect later announcements. The theoretical implications of the curvature of the relative performance component of compensation for the nature of the bias in the *last* analyst's forecast are derived. Importantly, we develop a frequency test for bias that is robust to (a) correlated information amongst analysts for a given firm; (b) unforecasted common industry-wide earnings shocks; and (c) information arrival so that later forecasts may be based on more information. Our frequency test provides overwhelming, incontrovertible evidence that the *last* analyst issues a biased forecast. We find this bias to be strategic and show how structural estimates of the bias can be obtained. The last analyst issues a contrarian forecast that tends to overshoot EPS *away* from the consensus (mean) forecast, in the direction of his private information. The economic magnitude of the bias is large: For every one percent that the last analyst's forecast overshoots the consensus, it overshoots earnings on average by between $\frac{3}{4}$ to $\frac{4}{5}$ of a percent, depending on the analyst following. Finally, we find evidence that investors do not unravel the bias in the last analyst's forecast.

*We are grateful to the Institutional Brokers Estimate System (I/B/E/S), a service of I/B/E/S International Inc. for providing analysts' forecasts of earnings per share. An earlier version of this paper was previously circulated as "Can Relative Performance Compensation Explain Analysts' Forecasts of Earnings?". We thank Lucy Ackert, Murillo Campello, George Deltas, Burton Hollifield, Ted Juhl, Owen Lamont, Gregor Smith, Brett Trueman, Mike Waldman, and Ján Zábojnik for helpful suggestions. The usual caveat applies.

†Department of Economics, 463 Commerce West, University of Illinois, Champaign, IL. 61820

‡Department of Economics, Tufts University, Medford, MA, 02155

1 Introduction

Security analysts' forecasts of firms' earnings are an integral part of the services provided by brokerage houses and money management funds. The advice brokerage houses give clients, and the consequent trading decisions made by these clients, rely heavily on the forecasts reported by the analysts. It is not surprising that the reliability of these forecasts have drawn attention both in popular financial press and among academic researchers. Among the fundamental questions addressed in our paper are the following: *Are* analysts' forecasts systematically biased? If they are biased, *how* biased are they, and what is the economic *impact* of the bias? Can we unravel the *reason* for the bias from the pattern of forecasts? Do investors *unravel* the bias in the forecasts?

This paper first develops a relative performance compensation based-theory for why forecasts might be biased, and derives the theoretical implications. In particular, we derive multiple testable restrictions on both the direction of the strategic bias, and for how the "amount" of uncertainty about earnings per share that the analyst faces affects the magnitude of the bias.

We then develop a frequency test for forecast bias that is *robust* to (a) correlated information amongst analysts for a given firm, (b) common unforecasted earnings shocks across firms, and (c) information arrival at different points in time. We find overwhelming evidence that forecasts are biased: evidence that is consistent with analysts strategically selecting contrarian forecasts that tend to overshoot EPS *away* from the consensus (mean) forecast, in the direction of their private information. These findings are sharply inconsistent with herding.

The bias that we find then leads us to determine both its economic impact, and whether investors unravel the bias. If analysts' forecasts are only marginally biased, then from an economic perspective, the forecast bias may be unimportant, especially if investors unravel the bias. We show how our model permits structural estimates of the strategic bias that analysts introduce into their forecasts. We find that the strategic bias is economically significant, and declines with measures of the "amount" of uncertainty about earnings per share, as predicted by our theory. Finally, we provide modest evidence that investors do not unravel the bias in the last analyst's forecast.

Many papers postulate plausible reasons as to why forecasts may be biased in one way or another. Laster, Bennett and Geoum (1999) and Hong, Kubik and Solomon (2000) explore how reputational or career concerns may cause analysts to rationally bias their forecasts away from, or closer to, the consensus.[1] Ehrbeck and Waldman (1996) suggest that analysts may be placing too much weight on their new (private) information hence their forecasts tend to be biased away from the consensus given all available information. Trueman (1990, 1994) argues that if investors attempt to infer analyst ability from forecasts, less able analysts will mimic the forecasts of others, thus

[1]Dugar and Nathan (1995) discuss other incentives that may impact on analysts' forecasts. Lin and McNichols (1998) and Michealy and Womack (1999) find that analysts working for investment banks that under-write equity offerings tend to issue more favorable forecasts and recommendations about the underlying firm's earnings. See also Lim's (2000) concern that analysts may issue optimistic forecasts in order to retain access to management.

herding toward the consensus. Also, analysts may be reluctant to revise forecasts upon receiving new information if revisions suggest that the initial forecast was less precise so that the analyst will be perceived as inept.

To determine which, if any, of these conflicting predictions is correct, one must first confront a fundamental problem: determining empirically *whether* and *how* analysts bias their forecasts is a formidable task, one that has largely defied solution. The problem is that one must develop a test for bias that is robust to

- Correlation in the information that analysts receive.
- Common, unforecasted earnings shocks that hit all firms in an industry.
- Information arrival over time; later forecasters will typically have better information on which to base their forecasts than earlier ones.

Correlated information amongst analysts creates problems for inference about bias because what may seem to be biased "herding" of forecasts toward the consensus may just be unbiased forecasts by analysts who make common errors due to the correlated information that they receive.[2] For example, if good news about earnings arrives after all forecasts are made, then existing forecasts will tend to be clustered too low, but the forecasts may have been unbiased estimates given the information analysts had when the forecasts were issued. Disentangling bias from common errors due to correlated signals is a particular concern for earnings forecasts because earnings analysts rely heavily on a single common source, the firm's Chief Financial Officer, for information. If the CFO misinforms analysts, then forecasts will reflect this common mis-information, and all will be too low or too high, even though each analyst's forecast may be unbiased given his information. This problem is most transparent if all analysts receive the *same* signal so that they make the *same* unbiased forecast: perfectly correlated signals give rise to perfectly clustered unbiased forecasts that are invariably too high or too low. The problem is to distinguish biased herding of forecasts from unbiased forecasts with common information.

To see how common, unforecasted earnings shocks create problems for inference, observe that since 1989, when reliable earnings forecasts were first available, earnings have tended to exceed forecasts. But this fact need not indicate that analysts' forecasts are biased downward (and/or biased together); rather it may reflect that firms did better in the past decade than everyone expected. The large increase in share prices over this period is consistent with the arrival of large common, positive earnings shocks that were unforecasted by both analysts and investors.

Finally, consider the impact of intertemporal information arrival. Forecasts may appear "too"

[2]For example, Gallo *et al.* (2001) find that forecasts of GDP converge as the date at which GDP is announced draws nearer, but that invariably the final forecasts are either uniformly too low or too high. They interpret this as herding, but an alternative explanation is that analysts' signals are correlated.

scattered if analysts make forecasts at different points in time, and earnings information arrives after some, but not all, forecasts have been made and earlier analysts do not issue revised forecasts.[3]

Our empirical work was motivated by our theoretical analysis of how relative performance compensation causes analysts to bias their forecasts. Relative performance compensation emerges naturally in environments where potential employers and clients are trying to infer analysts' abilities from their forecasts.[4] For example, if good equity researchers make more accurate forecasts, and hence better investment recommendations, they will win more clients, generate higher trading volumes and more brokerage commissions for their firms. Because the relative accuracy of an analyst's forecast helps firms and investors unravel the analyst's ability, competition among firms for better analysts leads to relative performance compensation: relatively more accurate analysts will be compensated both directly through bonuses, and indirectly through job placement, compensation for greater trading volume, and so on.[5]

When relative performance compensation is significant, then when reporting forecasts, analysts will seek to balance the aim of minimizing forecast error against looking good relative to other forecasters. We consider an environment in which forecasts are released sequentially and derive tight predictions about how the nature of relative performance compensation affects the forecast reported by the *last* analyst, relative to the consensus forecast.

We first show that if "better" relative performances are rewarded less than the losses incurred from "worse" performances, then the last analyst strategically biases his report toward the consensus. Such concave relative performance compensation may arise if sufficiently poor relative performance would cause the analyst to be fired.[6]

If, instead, relative performance compensation is a convex function of relative performance (as would be the case if, for example, being "right" when everyone else is wrong raises the number of clients desiring the analyst's services substantially (Laster *et al.* (1999), Henry (1989)), then the last analyst strategically biases his forecast in the direction of his private information, away from the consensus. We also show that this strategic bias is greater if the last analyst is more uncertain about earnings per share. Thus, the bias is predicted to be greater for smaller firms/firms that are followed by fewer analysts.

[3]Our focus is on whether forecasts are biased. Welch (2000) interprets herding as an increased probability of moving to the consensus (rather than just near). If the consensus reflects the views of many analysts, then an analyst's best posterior forecast heavily weights the consensus, so that unbiased analysts would generate the rational "herding" that Welch finds.

[4]Chevalier and Ellison (1997) document empirically how mutual fund manager's investments evolve over time, presumably reflecting learning by investors who direct investment funds according to performance.

[5]There is extensive evidence that analysts and employers care about the *relative* accuracy of analysts' forecasts. For example, the annual All-Star Analysts Survey (conducted by the Wall Street Journal and Zacks Research) rank security analysts on the basis of their earnings forecasting skills. Analysts' "skills" are measured by comparing the accuracy of their earnings estimates against those of other analysts following the same firm. Top-ranked analysts receive large bonuses (Stickel (1990), (1992)).

[6]Mikhail *et al.* (1999) finds that analysts who are relatively less accurate than their peers are more likely to lose their jobs, but that absolute forecast accuracy does not affect layoff probabilities.

Were the last analyst unbiased, then his forecast should be as likely to exceed EPS as to fall short, independently of whether the last forecast exceeded or fell short of the consensus. But our frequency test provides overwhelming evidence that the *last* analyst's forecast is biased. The last analyst's forecast over-emphasizes his private information by announcing a contrarian forecast that tends to overshoot the EPS *away* from the consensus (mean) forecast. The frequency test shows that the last forecast overshoots EPS away from the consensus 60% of the time. It overshoots the EPS **both** when it falls short of the consensus (in which case, it falls short of EPS 66% of the time), and when it exceeds the consensus (also exceeding EPS 55% of the time). Our test statistic is remarkably stable, varying by less than five percentage points across all years, despite the large variation in common earnings shocks.

Our frequency findings incontrovertibly document that *last* forecasts are biased, a bias that is consistent with convex relative performance compensation, and inconsistent with herding theories. However, the frequency results do not provide information about the economic magnitude of the bias. We next show that if the last analyst's strategic bias is a linear function of the expected error in the consensus forecast given the *analyst*'s information, then we can obtain estimates of both the strategic bias chosen and the expected dollar impact of the bias, as a function of the observable difference between the last forecast and the consensus forecast. The estimated bias is economically significant: For every one percent that the last analyst's forecast overshoots the consensus, it overshoots earnings on average by between 0.75 to 0.82 of a percent, depending on the analyst following. This implies, for example, a strategic bias of about 3.6 times the difference between *last* analyst's (unobserved) true posterior estimate of earnings and the consensus forecast if two analysts follow the firm; while if 20 analysts follow the firm, the estimated strategic bias is still about 1.7 times the difference. As our theory predicts, the estimated bias rises with measures of the amount of uncertainty about earnings per share.

Our theory can also reconcile the empirical findings of other researchers. Ehrbeck and Waldmann (1996) find that forecasters of interest rates on 91-day U.S. Treasury bills revise their forecasts by too much and that large revisions tend to be correlated with large forecast errors. This finding is consistent with analysts issuing forecasts that over emphasize their private information because of convex relative performance compensation.

Our empirical results contrast sharply with those of Keane and Runkle (1998). Keane and Runkle attempt to control for the common signals that analysts receive, and argue that "The evidence strongly supports the view that professional stock market analysts make rational (unbiased) forecasts of earnings per share." The difference between our findings and Keane and Runkle's is startling. Importantly, our frequency test is robust: correlated information cannot explain the differences in our findings. A possible explanation for why Keane and Runkel fail to uncover the bias is that they do not distinguish among analysts. With convex compensation, robust predictions about the nature of bias in analyst's forecasts only obtain for the *last* forecaster, because earlier forecasters must take into account how their forecasts affect subsequent forecasts. However, while we obtain tight theoretical predictions only for the last analyst, in a related empirical work, Bernhardt, Campello and Kutsoati (2001) find qualitatively similar levels of biases for all analysts. A second possible reason is that Keane and Runkle consider forecasts for only 21 firms *very* heavily followed firms, where our theoretical analysis predicts the bias should be the smallest, and which our empirical work documents.

The rest of the paper is organized as follows: Section 2 presents a spare model of sequential forecasting, and derives the effect of relative performance compensation for announced forecasts. Section 3, describes the data, develop testable hypotheses and present our empirical findings, and discuss alternative explanations for our findings. Section 4 concludes. All proofs are in an appendix.

2 The Model

Our theoretical analysis focuses on the *last* of L security analysts who issue forecasts the earnings per share, E, of a firm. In practice, the last analyst will have several pieces of information at his disposal, including forecasts by earlier analysts since forecasts are produced and become publicly available in an almost continuous fashion. For example, subscribers (and contributors) to the Institutional Brokers Estimate System (I/B/E/S) EXPRESS receive the latest forecasts of earnings from a panel of security analysts. This product is delivered electronically and is accompanied by the mean (or consensus) and other summary statistics. Forecasts are also disseminated through newsletters, web-pages such as Yahoo, and so on.

We let $g(E|\Omega_L)$ denote the density summarizing analyst L's posterior beliefs about E given his information Ω_L, and let $G(E|\Omega_L)$ be the associated conditional cumulative distribution function. The only structure placed on L's information is: (i) the consensus forecast, F_m, is in L's information set; and (ii) $g(E|\Omega_L)$ is symmetric about the posterior median, $\hat{\theta}_L$, *i.e.* $g(E-\hat{\theta}_L|\Omega_L) = g(E+\hat{\theta}_L|\Omega_L)$, and is strictly positive on its connected support.

Our assumptions admit a *variety* of informational structures. We impose no structure on *how* the last analyst forms his posterior; he may well use the information in the forecasts of earlier analysts when updating. Earnings per share could be the sum of normal innovations: $E = \sum_{i=1}^{L+1} \theta_i$, where $\theta_1, ...\theta_{L+1}$ are drawn from some multi-variate normal distribution with mean zero and arbitrary variance-covariance matrix Ψ, and analyst L observes a noisy signal, $s_L = \sum_{i=1}^{L} \theta_i + \eta_L$, where η_L is an i.i.d., mean zero, normally distributed error term. Then, the last analyst's posterior is normally distributed, and hence symmetric. In particular, later forecasters can have better information about earnings than earlier forecasters, and analysts may receive correlated signals.

The other cornerstone for our analysis is that analyst L's total direct and indirect compensation depends on both his absolute forecast accuracy, and his relative forecast accuracy:

$$w_L = \overline{w} - \lambda|F_L - E| + R(|F_m - E| - |F_L - E|), \tag{1}$$

where $\overline{w}$ is a fixed wage, F_L is analyst L's forecast, and F_m is the mean of all other forecasts excluding F_L. The weight on absolute forecast accuracy, $|F_L - \text{EPS}|$, in compensation is given by $\lambda > 0$. Relative compensation, $R(|F_m - \text{EPS}| - |F_L - \text{EPS}|)$, rises if the last analyst's is relatively more accurate than the consensus, $R'(\cdot) > 0$. That is, as the consensus becomes relatively less accurate (*i.e.*, as $|F_m - \text{EPS}|$ rises relative to $|F_L - \text{EPS}|$), analyst L's compensation rises. Without loss of generality, we normalize $R(0)$ to zero. All of the theoretical findings that we test extend if we relax the linear structure on absolute performance compensation and assume only that absolute performance compensation rises monotonically with forecast accuracy.

The minimal structure imposed on the reduced form payoffs of analyst L ensures that they are consistent with many economic environments, including learning about analyst ability by employers and/or competing firms, tournaments to elicit effort among analysts, outside competition for able workers which induces compensation uncontrolled by the employer, etc.

We also impose no structure on *why* a particular analyst is last. In practice, there may be few reasons why an analyst chooses to report earlier. For example, empirically it turns out that earlier analysts tend to issue more optimistic forecasts than later analysts; the last analyst's forecast falls short of the consensus 58 percent of the time. This may be because analysts must make buy recommendations for clients to trade on for some subset of stocks; those with better initial signals place these stocks on buy lists, announcing earlier so as to give their clients more time to trade. In contrast, those with worse initial signals wait because they do not need to provide clients as much advanced warning, more information about the stock may arrive in the interim, and it may be easier to manipulate relative forecasts if one announces later, and hence can see earlier forecasts. *None* of our findings depend on the last analyst's identity, nor on why he is last; we just identify the consequences of different forms of relative performance compensation for the bias introduced by the last analyst given his information. Indeed, while we only obtain theoretical predictions for the last analyst (see remark 1 at the end of section 2.1), the *empirical* bias in earlier forecasts is similar to that for the last analyst (Bernhardt, Campello and Kutsoati (2001)).

We next derive and test the theoretical implications of different induced preferences, which, in turn, impose restrictions on the possible economic environments.

2.1 Relative Performance Incentives

To analyze how relative performance compensation affects the behavior of the last analyst, we first consider a linear relative performance compensation function as the benchmark. We show that if $R(\cdot)$ is linear, then the last analyst's forecast is unaffected by those made by others. If R is linear, then the analyst's compensation simplifies to

$$w_L = \overline{w} - (1+\lambda)|F_L - E| + |F_m - E|.$$

Even though the reports of other analysts still affect the last analyst's compensation — more accurate forecasts by others reduce his compensation — the last analyst maximizes expected compensation by reporting a forecast equal to the median of his posterior distribution, $F_L = \hat{\theta}_L$.

Proposition 1 *Suppose the relative compensation function of the last analyst is linear. Then the last analyst's equilibrium forecast equals the posterior median of his beliefs about earnings per share:* $F_L = \hat{\theta}_L$.

Simply put, if R is linear, introducing any bias in the forecast reduces both expected relative forecast accuracy and expected absolute forecast accuracy, lowering compensation. Note that if R is linear then the last analyst will report the median of his posterior over earnings per share **even if** $g(\cdot)$ is not symmetric.

We now show that independently of the form of $R(\cdot)$, some forecasts are strictly dominated given the last analyst's information and the consensus forecast:

Lemma 1 *If the consensus forecast is less than the last analyst's expectation of earnings, i.e.,* $F_m < \hat{\theta}_L$*, then his report exceeds the consensus,* $F_L > F_m$*. If, instead,* $\hat{\theta}_L < F_m$ *then* $F_L < F_m$.

Lemma 1 ensures that if the last analyst's expectation of EPS exceeds the consensus forecast, then so will his forecast; and if he expects earnings to be less than F_m, then his forecast will be less than F_m. For every F_L' such that $F_L' < F_m < \hat{\theta}_L$, since his posterior is symmetric about $\hat{\theta}_L$, the forecast $F_m + (F_m - F_L')$ dominates F_L' in terms of both expected absolute and relative forecast accuracy.

Lemma 1 details which forecasts are dominated given the analyst's information and hence would never be issued by the last analyst. We now characterize how the pattern of forecasts by the last analyst will vary with the nature of the relative compensation function. In particular, we determine when the last analyst will choose either to "under-emphasize" or to "exaggerate" his information.

We first show that if analyst L's compensation is a concave function of relative performance, then his forecast will under-emphasize his private information, causing him to locate more closely to the consensus than his information would suggest:

Proposition 2 *Suppose* $R(\cdot)$ *is concave. Then*
(i) If $F_m < \hat{\theta}_L$ *then* $F_m < F_L < \hat{\theta}_L$*. Conversely, if* $F_m > \hat{\theta}_L$*, then* $F_m > F_L > \hat{\theta}_L$.
(ii) $0 < \frac{d(\hat{\theta}_L - F_L)}{d(\hat{\theta}_L - F_m)} < 1$*: that is, the bias in* F_L *increases in* $(\hat{\theta}_L - F_m)$.

Concave relative performance compensation causes the last analyst to be averse to locating too far from the consensus, lest his forecast be relatively inaccurate. He willingly foregoes some of the compensation that derives from smaller absolute forecast errors in order to be closer to the consensus: If he is wrong, then there is comfort in 'numbers', in having others be wrong, too. Proposition 2 details that the consequence is that the last analyst underweights his private signals, biasing his forecast toward the consensus forecast. Hence, if analysts' payoffs are a concave function of relative performance, forecasts can become "clustered" even if individual private information suggests otherwise. Indeed, as the last analyst's expectation, $\hat{\theta}_L$, diverges further from F_m, his forecast diverges further from $\hat{\theta}_L$, but at a rate less than one. Equivalently, if the consensus forecast were revised, this would cause the last analyst to revise his forecast in the direction of consensus' revision, even though his beliefs about earnings are unaltered.

Proposition 2 forms the basis of our predictions about a concave $R(\cdot)$. The conservative bias introduced in the forecast implies that if F_L exceeds F_m, it is more likely that the last forecast will be less than the actual earnings per share. Conversely, if F_L is less than F_m, then F_L is more likely to exceed actual earnings. This generates a relationship between the pattern of forecasts and the direction of error in the last forecast.

Alternatively, relative compensation may be convex, $R''(|F_m - E| - |F_L - E|) > 0$, perhaps because one or two accurate forecasts that depart from the consensus may draw a substantial clientele. The next proposition details that convex relative performance compensation causes the last analyst to "exaggerate" his information by reporting an F_L that overshoots $\hat{\theta}_L$ away from F_m.

Proposition 3 *Suppose that $R(\cdot)$ is convex. Then if $F_m < \hat{\theta}_L$, then $F_m < \hat{\theta}_L < F_L$. Conversely, if $F_m > \hat{\theta}_L$, then $F_m > \hat{\theta}_L > F_L$.*

If the relative compensation function is convex, the last analyst has an incentive to bias his forecast to "locate" away from F_m: a relatively accurate forecast yields a much higher payoff than the losses incurred if his forecast is less accurate. This induces the last analyst to take forecasting risks by reporting contrarian forecasts. In particular, the last analyst will report an overly-optimistic forecast of the firm's EPS if $\hat{\theta}_L > F_m$, and an overly-pessimistic forecast if $\hat{\theta}_L < F_m$. Consequently, F_L is more likely to exceed actual earnings than fall short if F_L exceeds F_m. So, too, F_L is more likely to fall short of earnings actual earnings than exceed earnings if F_L exceeds F_m.

We next show that the bias that the last analyst introduces in his forecast rises with the 'amount' of uncertainty he faces about earnings.

Proposition 4 *Suppose that $R(\cdot)$ is convex, and consider two posterior belief distributions for the last analyst, $g_L(E|\Omega_L)$ and $\tilde{g}_L(E|\Omega_L)$, where $\tilde{g}_L(\cdot)$ is a mean preserving spread of $g_L(\cdot)$. Then for any given difference between the consensus and the last analyst's expectation, $(F_m - \hat{\theta}_L)$, the last analyst chooses a more severely biased report when his posterior beliefs are less precise:*

$$|F_L(\tilde{g}_L) - \hat{\theta}_L| > |F_L(g) - \hat{\theta}_L|.$$

Proposition 4 has no frequency analog. Although a less-dispersed posterior reduces the bias in the last analyst's forecast, there is more probability mass close to the median, so that the probability of overshooting may be higher or smaller if more analysts follow a firm. We implement Proposition 4's content empirically as follows:

Corollary 1 *Suppose that $R(\cdot)$ is convex and that the last analyst's posterior about E is tighter if he observes more forecasts about a firm's EPS. Then for any given difference between the consensus and the last analyst's forecast, $(F_L - F_m)$, the expected bias in the last analyst's forecast is less when there are more analysts:*

$$E[(F_L - \hat{\theta}_L)|F_L - F_m]$$

is a decreasing function of L.

That is, the more information sources the last analyst has, the tighter is his posterior, and hence the smaller is the strategic bias in the last analyst's forecast. Consequently, for any given deviation in the last forecast away from the consensus, the expected error in his forecast falls with the number of analysts or related measures of the "amount" of uncertainty, such as firm size.

Remark 1 The reader may wonder why we do not characterize the forecasts of all analysts. In fact, some characterizations are possible. If $R(\cdot)$ is linear, then each analyst reports an unbiased forecast equal to the median of his posterior given his information. If $R(\cdot)$ is concave, and signals are i.i.d., then the first analyst issues an unbiased forecast, and subsequent analysts issue forecasts that are biased toward the extant consensus. If, instead, $R(\cdot)$ is convex, characterizations of earlier forecasts are not possible, because forecasts by later analysts vary with forecasts of earlier analysts, who, in turn, hope to separate from a consensus that will incorporate these later forecasts.[7] Indeed, when $R(\cdot)$ is convex, even if there are only two analysts,

- The *first* analyst may bias his forecast if doing so tends to cause the second analyst to increase his bias by enough: more dispersed forecasts benefit both analysts.
- The first analyst may expect higher payoffs, even though the second analyst has superior information. This is because the last analyst may issue a more biased forecast. For example, if their signals are perfectly correlated, and the first analyst issues an unbiased signal, then if the second analyst issues a biased signal, his expected payoff will be lower. *Thus, if relative performance compensation is convex, then even if the functional form of compensation does not vary with the forecast order, analysts in our model may not prefer to go last.*[8]

[7]Consider 3 analysts. Suppose that $F_1 < \hat{\theta}_2 = \hat{\theta}_3$, and relative compensation is convex. Analyst 2 wants to bias his forecast away from F_1 only if that bias raises $|F_2 - 0.5(F_1 + F_3)|$. But analyst 3 will bias his forecast away from $0.5(F_1 + F_2)$, so that if analyst 2 biases his forecast further away from F_1, it may cause analyst 3 to further bias his forecast away from F_1. But this may reduce $|F_2 - 0.5(F_1 + F_3)|$.

[8]In practice, another benefit of earlier forecasts is that clients have more time to trade on the information, even though the forecast may be less accurate.

Remark 2 Our qualitative theoretical predictions extend if the last analyst is penalized for bad relative performance just as much as he is rewarded for a good relative performance. More formally, relative performance compensation is said to be *anti-symmetric* about zero if:

$$R(|F_m - E| - |F_L - E|) = \begin{cases} R(|F_m - E| - |F_L - E|), & \text{if} \quad |F_m - E| > |F_L - E| \\ -R(|F_L - E| - |F_m - E|), & \text{if} \quad |F_m - E| < |F_L - E|, \end{cases}$$

with $R(0) = 0$ and $R'(\cdot) > 0$. The relative performance compensation function is *anti-symmetric* concave if $R(\cdot)$ is a concave function of positive relative performances, and *anti-symmetric* convex if $R(\cdot)$ is a convex function of positive relative performances. For example, anti-symmetric convex compensation may arise if an analyst received large bonuses for positive relative performances, but would be fired for substantially negative relative performances. Kutsoati (1998, ch. 3) shows that if the relative performance function is "anti-symmetric convex", then the last analyst will again bias his forecast toward his private signal and away from the consensus; if, instead, the relative performance function is "anti-symmetric concave," he will bias his forecast toward the consensus. The proofs mirror those here. The key is to exploit the symmetry in the analyst's beliefs about earnings and the fact that $R'(z) = R'(-z)$ for all z. The intuition is that relative to the consensus, the last analyst anticipates being right more often than being wrong, so that his forecast reflects more heavily the curvature assumptions on the payoffs for good relative performances.

3 Empirical Analysis

3.1 Data and Sample Selection

We extract firms' earnings per share and data on individual analysts' quarterly forecasts from 1989 to 1999 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.[9] After 1988, forecasts were disseminated by I/B/E/S within 24 hours (see I/B/E/S Research Bibliography (1996)). Our sample selection ensures that the publication dates are close to the actual dates at which analysts released their forecasts and, more importantly, the last analyst *observes* earlier forecasts. Data on the quarterly earnings announcement dates, daily returns, share prices and the number of shares outstanding at the end of each quarter are taken from the Center for Research on Security Prices (CRSP) files.

We first filter the data for likely data entry errors, deleting any forecast with an absolute error

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

value of more than \$10.[10] Second, for some firm–quarters, the number of days between the last forecast and the earnings announcement is quite high; we discard an observation if the number of days from the last forecast to the date of actual earnings announcement exceeds 60 days.[11] Third, since our analysis involves isolating the last forecast in each firm–quarter (according to the publication dates) and relating it to the outstanding *consensus forecast*, we drop firm–quarters when only one analyst issues forecasts of that period–end's EPS.

In practice, analysts occasionally revise their forecasts before quarterly EPS are made public by the firm, so that the *total* number of forecasts include both old and revised forecasts. Because older forecasts tend to be less accurate *ex-post* (O'Brien (1988)), we use only an analyst's latest forecast, discarding his earlier forecasts. Our consensus forecast is then equal to the mean of the latest forecast of each analyst, *except* the last analyst. Using the mean of the latest forecasts seems to be the appropriate measure of consensus, as it is highlighted in the I/B/E/S database, and by other researchers. Importantly, choosing the wrong measure for the "consensus" only reduces our chances of finding that the last forecast is systematically affected by the outstanding consensus.

Table 1: Sample Statistics

			Percentiles			Sample moments	
	min	max	25	50	75	Mean	Std. dev.
Number of analysts (cover)	2	38	3	6	10	7.14	4.90
Total forecasts[a]	2	85	4	8	14	10.38	8.49
Forecasts–Analysts ratio	1	4	1.11	1.33	1.60	1.40	0.37
Number of Days from *last* forecast till earnings report (days)	0	60	7	14	27	17.60	14.21
Error in *last* forecast (as a percent of stock price)[b]	-20.38	22.86	-0.101	-0.006	0.076	0.025	0.515
Difference in *last* and *consensus* forecasts (as a percent of stock price)	-21.78	22.93	-0.106	-0.009	0.035	-0.065	0.460

NOTES:
The sample has 63,711 firm-quarters. All firm-quarters have at least two forecasts reported by distinct analysts on different dates.
a. Total forecasts is the total number of forecasts (including both old forecasts and their revisions) reported by all analysts in a firm-quarter.
b. To reduce the sensitivity of our results to noise in prices and to mitigate the effects of outliers, we drop firm–quarters with share price under \$5 and observations in the 2.5% tails of the percentage error in the mean forecast.

Table 1 gives the distribution of analyst coverage, total forecasts reported, the average number of

[10] O'Brien (1988) and Lim (2001) used a similar rule in deleting suspected data-entry errors.
[11] None of these filtering rules have a significant impact on our findings.

forecasts per analyst, and the number of days from the last report till the actual EPS announcement date across all firms-quarters. We also report the distribution of *error* in the last forecast (as a percent of the stock's price as at the end of the previous quarter) and the *difference* in the last and consensus forecasts. To remove the effects of outliers in percentage errors and differences, we discard observations with stock price less than \$5.00, and the 2.5% tails of the distribution of the percentage error in the consensus (see Lim (2001)). Our final sample consists of 63,711 firm quarters.

Table 1 shows that an average of about 7 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 analysts infrequently revise their one-quarter ahead forecasts. Finally, in about half of the sample, the last forecast was reported about two weeks before the actual earnings report was released by the firm.

3.2 Frequency tests

In this section we develop a frequency test for biased forecasts that is robust to (a) correlated information amongst analysts, (b) common unforecasted earnings shocks, and (c) information arrival. If the last analyst reports an unbiased forecast equal to the posterior median of his beliefs about earnings per share, then his forecast should be as likely to exceed earnings per share as fall short, both unconditionally, and when we condition on anything in his information set. In particular, the probability with which the last forecast exceeds actual earnings should not vary with whether or not the last forecast exceeded the extant consensus forecast. We test this unbiasedness hypothesis, against the alternatives that relative performance incentives may cause the last analyst to bias his forecast toward or away from the consensus. Essentially, we use the frequencies with which the last forecast exceeds or falls short of true EPS, given the consensus forecast, to estimate the probability of overshooting.

We first group the data into mutually exclusive events according to the realization of earnings per share, E_τ, relative to the last forecast $F_{L\tau}$, and the outstanding consensus, $F_{m\tau}$, in each firm-quarter, τ. In practice, both earnings per share and forecasts by analysts are reported in cents: some unobserved rounding occurs. This creates a problem with inference whenever the last forecast equals the consensus or actual EPS. In particular, we cannot discern whether the last analyst's forecast exceeds the consensus or not; and whether it exceeds true EPS or not. Consequently, for the frequency tests, we drop observations where either $F_{L\tau} = E_\tau$ or $F_{L\tau} = F_{m\tau}$. Accordingly, let N be the number of observations in our sample, and define indicator functions that capture whether the last forecast exceeded the consensus:

$$\gamma_\tau^+ = 1 \text{ if } F_{L\tau} > F_{m\tau}, \text{ and zero otherwise,}$$

$$\gamma_\tau^- = 1 \text{ if } F_{L\tau} < F_{m\tau} \text{ and zero otherwise,}$$

and define indicator functions, δ_τ^+ and δ_τ^-, to take on the value one if the last forecast overshoots earnings per share in the same direction as it overshoots the consensus (see Figure 1):

$$\delta_\tau^+ = 1 \text{ if } F_{L\tau} > E_\tau \text{ and } F_{L\tau} > F_{m\tau}; \text{ and } \delta_\tau^+ = 0 \text{ otherwise.}$$

$$\delta_\tau^- = 1 \text{ if } F_{L\tau} < E_\tau \text{ and } F_{L\tau} < F_{m\tau}; \text{ and } \delta_\tau^- = 0 \text{ otherwise.}$$

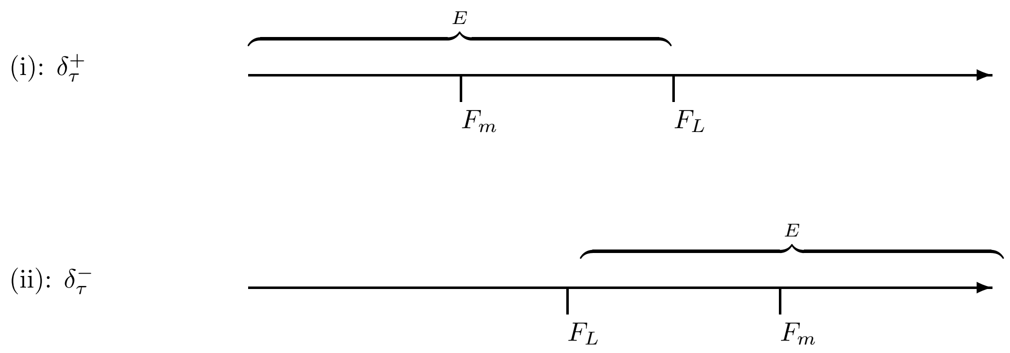

Figure 1: Realization of actual EPS relative to F_L and F_m

Thus, $E[\delta^+ + \delta^-]$ is the probability the the last forecast overshoots earnings per share in the same direction as it overshoots the consensus. Figure 1 illustrates the events δ_τ^+ and δ_τ^-. If the last analyst's forecast is median– unbiased, then *independently of anything in his information set, including the location of the consensus,* $F_{L\tau}$ should correspond to his estimate of the median earnings per share, $\hat{\theta}_{L\tau}$. If, instead, the last analyst biases his forecast toward the consensus, then his forecast will be located between $\hat{\theta}_{L\tau}$ and $F_{m\tau}$, making the conditional probability of overshooting (*i.e.* the frequency of events δ_τ^+ and δ_τ^-) less likely; while if the last 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 frequency test for bias.

Under the null that his forecast is unbiased, both $\sum_\tau \delta_\tau^+$ and $\sum_\tau \delta_\tau^-$ are binomially distributed with mean 0.5. Hence, using the normal approximation to the binomial, under the null, the sample average

$$\mathbf{S}' = \frac{\sum_\tau \delta_\tau^+ + \sum_\tau \delta_\tau^-}{N},$$

is (asymptotically) normally distributed with mean 0.5, and variance $\frac{0.25}{N}$. This sample average can be used to estimate the population probability of overshooting. It is also important to note that the properties of S' do not depend on how the last analyst forms his posterior, $g(\cdot)$. 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.

However, we also want our test to be robust to correlated unforecasted earnings shocks *across* firms. It is likely, for example, that in a given quarter, firms in an industry receive a common un-anticipated earnings shock, so that, *ex post*, forecasts are systematically too low or too high relative to realized earnings. Accordingly, we use

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

to estimate the population probability of overshooting, $E[\delta^+ + \delta^-]$. Note that $\frac{\sum_\tau \delta_\tau^+}{\sum_\tau \gamma_\tau^+}$ is our estimate of the conditional probability of overshooting actual earnings given that our forecast exceeds the consensus; while $\frac{\sum_\tau \delta_\tau^-}{\sum_\tau \gamma_\tau^-}$ is our estimate of the conditional probability of falling short of true earnings given that our forecast falls short of the consensus. $\mathbf{S}$ is therefore the average of our estimates of the two conditional overshooting probabilities. If the last forecast exceeds the consensus by more often than it falls short, $\mathbf{S}$ weights $\sum_\tau \delta_\tau^+$ by relatively less than it weights $\sum_\tau \delta_\tau^-$.

$\mathbf{S}$ 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:

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

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

The keys for the robustness of S are (i) ω_t does not affect whether the last forecast exceeds the consensus, *i.e.* whether $F_{Lt} > F_{mt}$, and (ii) $\omega_t \neq 0$ has exactly offsetting impacts on the frequencies of $F_{Lit} > E_{it}$, and $F_{Lit} < E_{it}$. For example, $\omega_t < 0$ raises the frequency with which event δ_{it}^+ occurs, but has an equally offsetting impact on the frequency with which δ_{it}^- occurs.

Suppose first that $\omega_\tau = \omega, \forall \tau$, and that analysts' posteriors about earnings realizations for different firms are identically distributed about their medians; $g_\tau(\hat{\theta}_\tau - E_\tau|\Omega_{L\tau}) = g_{\tau'}(\hat{\theta}_{\tau'} - E_{\tau'}|\Omega_{L\tau'})$ for firm quarter $\tau \neq \tau'$. Under the null hypothesis that forecasts are unbiased, $\sum_\tau \delta_\tau^+$ has a binomial distribution, $B(\sum_\tau \gamma_\tau^+, G(\hat{\theta}+\omega))$, where $\sum_\tau \gamma_\tau^+$ is the number of times that $F_{L\tau} > F_{m\tau}$; and $\sum_\tau \delta_\tau^-$ has a binomial distribution, $B(\sum_\tau \gamma_\tau^-, 1 - G(\hat{\theta}+\omega))$, where $\sum_\tau \gamma_\tau^-$ is the number of times that $F_{L\tau} < F_{m\tau}$. Hence, under the null, $\mathbf{S}'$ has a mean of

$$\frac{\sum_\tau \gamma_\tau^+}{N} G(\hat{\theta}+\omega) + \frac{\sum_\tau \gamma_\tau^-}{N}(1 - G(\hat{\theta}+\omega)),$$

which varies with ω and $g(\cdot)$ unless the last forecast exceeds the consensus as often as it falls short, in which case it has mean 0.5. In contrast, under the null, **S** has mean,

$$\frac{N}{2\sum_\tau \gamma_\tau^+}\frac{\sum_\tau \gamma_\tau^+}{N}G(\hat{\theta}+\omega)+\frac{N}{2\sum_\tau \gamma_\tau^-}\frac{\sum_\tau \gamma_\tau^-}{N}(1-G(\hat{\theta}+\omega))=\frac{1}{2}G(\hat{\theta}+\omega)+\frac{1}{2}(1-G(\hat{\theta}+\omega))=\frac{1}{2},$$

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

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

Thus, under the null of median unbiasedness, asymptotically,

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

Importantly, the variance of **S** *is bounded from above by a number* $\frac{1}{16}\left[\frac{1}{\sum_\tau \gamma_\tau^+}+\frac{1}{\sum_\tau \gamma_\tau^-}\right]$*, that does not depend on* ω *or* $g(\cdot)$*. Hence, neither the mean of* **S** *nor the upper bound on its variance are affected by variations in* ω_τ *and* $g_\tau(\cdot)$ *across firm quarters. Thus, with regard to hypothesis testing, if we reject the null hypothesis of unbiasedness using the upper bound on the variance, then we can always reject the null.*

Our frequency test for forecast bias has other appealing features:

- It is non-parametric and unrelated to the scale of errors across firms: no assumption is made about the structure of the relationship between the error in the last forecast and the difference in F_L and F_m.
- Outliers do not have disproportionate effects on frequencies. Consequently, the tests are robust to failing to exclude unusual "one-time" events such as large discretionary write-downs of assets that analysts do not seek to predict (see, for example, the concerns detailed in Keane and Runkle (1998) and Lim (2001)).

Under the alternative hypothesis that compensation is a convex in relative performance, $\mathbf{S} > 0.5$. That is, were $R(\cdot)$ convex, the last forecaster would bias his forecast past E away from F_m. Hence, if $F_L > F_m$, then F_L would exceed the last forecaster's expectation of E so that more often than not, E should fall short of F_L. Conversely, if $F_L < F_m$, then more often than not, E should exceed F_L. Therefore, convex relative compensation implies $\mathbf{S} > 0.5$. An analogous argument shows that if $R(\cdot)$ is a concave function of relative performance, then $\mathbf{S} < 0.5$.

The results, reported in Table 2, show that $\mathbf{S} = 0.608$: The analyst overshoots earnings in the direction away from the consensus about 60% of the time. Were the last forecast unbiased, and unaffected by the consensus, then **S** should be one-half.[12] The fact that $\delta_\tau^+ \cup \delta_\tau^-$ occurs so frequently

[12]Our test of median unbiasedness becomes a test of mean unbiasedness if the last analyst's posterior is symmetrically distributed so that the median corresponds to the mean. In the data, the distribution of forecast errors is approximately symmetrically distributed.

Table 2: Frequency Test of Bias in the Last Forecast

Samples	# of observations	Prob($F_L > F_m$)	**S**
All	52,794	0.418	0.608
			[0.604, 0.613]
Segmented by Number of Analysts (*i.e*, cover)[a]			
$2 \leq$ cover ≤ 3	12,837	0.440	0.624
			[0.616, 0.633]
$4 \leq$ cover ≤ 6	15,266	0.420	0.601
			[0.581, 0.604]
$7 \leq$ cover ≤ 10	12,958	0.410	0.608
			[0.599, 0.617]
cover ≥ 11	11,733	0.396	0.598
			[0.589, 0.607]
Segmented by **days** from *last* forecast EPS announcement date[b]			
days ≥ 25	13,934	0.433	0.631
			[0.623, 0.622]
$14 \leq$ days ≤ 24	12,134	0.411	0.608
			[0.600, 0.618]
$7 \leq$ days ≤ 13	13,724	0.412	0.598
			[0.590, 0.607]
days ≤ 6	13,002	0.412	0.593
			[0.585, 0.602]
Segmented by Calendar Year[c]			
1989	3,316	0.428	0.622
			[0.605,0.640]
1990	3,367	0.433	0.641
			[0.625, 0.659]
1991	3,946	0.392	0.612
			[0.596, 0.628]
1992	4,810	0.402	0.601
			[0.586, 0.615]
1993	4,338	0.432	0.593
			[0.578, 0.608]
1994	5,758	0.446	0.618
			[0.605, 0.630]
1995	6,357	0.437	0.593
			[0.581, 0.605]
1996	6,273	0.420	0.606
			[0.594, 0.618]
1997	6,408	0.422	0.612
			[0.600, 0.623]
1998	6,723	0.368	0.604
			[0.592, 0.616]
1999	1,498	0.428	0.592
			[0.567, 0.615]

NOTES:–
95% confidence intervals are reported in square brackets; confidence intervals for **S** are computed using the upper bound on the variance.
a. **cover** is the number of analysts providing a forecast for a firm in a quarter.
b. **days** is the number of days between the last forecast's publication date in the I/B/E/S database and the earnings announcement date.

Table 3: Conditional and Unconditional Probability of *Last* Forecast Overshooting Earnings

Sample	Prob($F_L > E$)	Prob($F_L < E \mid F_L < F_m$)	Prob($F_L > E \mid F_L > F_m$)
All	0.423	0.667	0.549
	Segmented by Number of Analysts (*i.e*, cover)[a]		
$2 \leq$ cover ≤ 3	0.454	0.655	0.594
$4 \leq$ cover ≤ 6	0.416	0.669	0.533
$7 \leq$ cover ≤ 10	0.423	0.665	0.552
cover ≥ 11	0.397	0.680	0.516
	Segmented by **days** from *last* forecast **So** EPS announcement date[b]		
days ≥ 25	0.439	0.675	0.587
$14 \leq$ days ≤ 24	0.432	0.657	0.560
$7 \leq$ days ≤ 13	0.418	0.662	0.534
days ≤ 6	0.403	0.674	0.513
	Segmented by Calendar Year		
1989	0.544	0.561	0.684
1990	0.543	0.580	0.705
1991	0.519	0.569	0.655
1992	0.492	0.588	0.613
1993	0.452	0.628	0.557
1994	0.415	0.690	0.546
1995	0.408	0.673	0.513
1996	0.377	0.712	0.499
1997	0.351	0.743	0.480
1998	0.347	0.730	0.479
1999	0.268	0.810	0.373

NOTES:–
This table reports the unconditional and conditional frequencies with which the *last* forecast exceeds the actual EPS, conditional on the whether the *last* forecast exceeds or falls short of the consensus forecast. 95% confidence intervals are reported in square brackets.
a. **cover** denotes the number of analysts issuing forecasts for a firm in that quarter.
b. **days** is the number of days between the last forecast's publication date in the I/B/E/S database and the earnings announcement date.

indicates that the last analyst "over-emphasizes" his own information in his forecast. Table 2 also reveals that our frequency test is *extremely* robust: There is **remarkably little variation** across sub–samples in our estimates of the probability that the last forecast overshoots earnings away from the consensus: **S** varies from a high of 64.1% in 1990 to a low of 59.2% in 1999. Table 2 also reveals that the the last analyst is slightly less likely to overshoot both when the last forecast is closer to the earnings announcement, and when the firm is followed by more analysts.

Table 3 reports the values of **S** for various sub–samples according to analyst–coverage (cover), the lag between actual earnings announcement date and the last forecast's date (**days**), and the calendar year. They show that despite the countervailing effects of common shocks the last analyst's forecast exhibits a contrarian bias **both** when he has positive and negative private information. Our estimates of each conditional probability of overshooting both exceed one–half: if his forecast exceeds the consensus, then it exceeds actual earnings 54.9% of the time, and if his forecast is less than the consensus then it falls short of actual earnings 66.7% of the time.[13] Simply put, the evidence against both the null of unbiasedness, and the alternative of herding is overwhelming.

Comparing Tables 2 and 3, we see that our estimates are robust to very large common unforecasted earnings shocks. Table 3 highlights the large impact of common unforecasted earnings shocks from year to year. The conditional probability that the last forecast falls short of earnings given that it fell short of the consensus varies from a high of 81% in 1999, when earnings were unexpectedly high, to a low of 56.1% in 1989 when the economy was in an unexpected downturn. So, too, the conditional probability that the last forecast exceeded earnings given that it exceeded the consensus varies from a high of 70.5% in 1990 to a low of 37.3% in 1999, (where unconditionally the last forecast exceeded earnings only 27% of the time). These results document both the importance of designing a test that is robust to common shocks, and our success in doing so. That is, while variations across different sub–samples shift up or down the mean probability that the last forecast measures exceed true earnings (see Table 3), and hence the frequency with which δ^+ occurs (see Figure 1), there is a corresponding offsetting shift in the frequency with which δ^- occurs.

Table 3 also documents that the last analyst is systematically more pessimistic than earlier analysts — issuing a forecast that exceeds the consensus only 42% of the time, a level of relative pessimism that does not vary across years. The last analyst is especially likely to be pessimistic relative to the consensus if there are many analysts, or if he issues a forecast close to the earnings announcement date. This may reflect that analysts with more positive initial signals tend to report earlier, so that their clients can buy on these recommendations. Alternatively, firms may 'manage' later analysts' forecasts down so as to generate 'positive' earnings surprises. These issues are explored in Bernhardt and Campello (2001).

[13]The difference in the frequencies of over-shooting, 54.9% versus 66.7%, reflect the unanticipated generally positive earnings shocks over the sample period, and highlight the importance of controlling for such shocks.

3.2.1 Economic Impact

Our frequency analysis robustly reveals that the last analyst's forecast is biased, but does not contain information about the economic magnitude of the bias. Accordingly, we now estimate the strategic bias introduced by the last analyst in his forecast. We first detail the regressions that we run of the error in the last forecast on the difference between the last forecast and the consensus, and then offer a structural interpretation.

Since earnings are reported on a per share basis, we express the error and difference as a percent of price of the share as at the end of the previous quarter.[14] We use the previous quarter to remove the contemporaneous effect of recent forecasts on the stock's price. The forecast error is then:

$$\text{ERROR}_{L(it)} = \frac{F_{L(it)} - E_{it}}{P_{i(t-1)}},$$

where $P_{i(t-1)}$ is firm i's share price at the *end of the previous* quarter $(t-1)$; and the difference between the last forecast and the consensus, expressed as a percentage of share price[15], is

$$SFD_{it} = \frac{F_{L(it)} - F_{m(it)}}{P_{i(t-1)}}.$$

Our benchmark OLS regression with firm–fixed effects is:

$$\text{ERROR}_{L(it)} = \beta_0 + \sum_i \text{firm}_i + \beta_1 \,\text{SFD}_{it} + \epsilon_{it} \tag{3}$$

Under the null hypothesis that the last forecast is unbiased, $\text{ERROR}_{L(it)}$, is essentially a normalized forecast error from the last analyst's forecasting regression: $\text{ERROR}_{L(it)}$ should be orthogonal to the everything in the last analyst's information set, including all of the independent variables in our regression, so that β_1 should be zero. However, in light of our frequency findings, we expect $\beta_1 > 0$.

To glean more information about the behavior of the last analyst, we then re-estimate equation (3) controlling for measures of how uncertain the last analyst is about what earnings per share will be: with convex compensation the strategic bias is predicted to rise if the analyst is more uncertain about earnings per share. Hence, the bias that we uncover is predicted to fall with the number of analysts covering the firm (Proposition 4, Corollary 1), or firm size.

Our second regression controls for analyst coverage, by adding the logarithm of analyst coverage and its interaction with SFD to equation (3);

$$\text{ERROR}_{L(it)} = \beta_0 + \sum_i \text{firm}\ _i + \beta_1 \text{SFD}_{it} + \beta_2 ln(L_{it}) + \beta_3 ln(L_{it}) {*} \text{SFD}_{it} + \epsilon_{it}. \tag{4}$$

[14]This method of standardizing forecasts error and forecast differences is standard, ensuring that our results do not vary with the firm's choice of shares outstanding. See, for example, Lys and Sohn (1990), Alexander and Ang (1997) and Lim (2001).

[15]Note that normalizing by share price causes our measures to be more sensitive to deviations for value stocks and stocks with declining earning, where current earnings are a greater portion of firm value, than they are for growth stocks. See notes under Table 1 on how we filter.

We also estimate equation (3) with using the firm–size (market capitalization) rather than analyst coverage. If the last forecast is unbiased, then $\mathrm{ERROR}_{L(it)}$ should not be systematically related to SFD_{it}. The interaction between SFD_{it} and analyst coverage should not be significantly different from zero. If, instead, the last analyst *strategically* biases his forecast away from the consensus due to convex relative performance incentives, then the coefficient on SFD_{it} should be positive, but the coefficient on interaction between SFD_{it} and the logarithm of the number of analysts should be negative, since the bias is predicted to fall as the last analyst's uncertainty is resolved. We also test whether the results could be driven by the timing of the *last* forecast relative to the actual earnings announcement date. That is, if the last forecast is issued closer to the earnings date, the analyst may have better information, and hence smaller forecast errors. We control for this by including $lndays = \ln(1 + days)$ in the regression. The results are reported in Table 4.

Structural interpretations of the parameter estimates can be obtained under the assumption that the last analyst's strategically–chosen bias is a linear function of the difference between the last analyst's (unobserved) true posterior estimate of earnings and the consensus forecast. In that case, he issues a forecast equal to:

$$F_{L\tau} = \hat{\theta}_{L\tau} + a_L(\hat{\theta}_{L\tau} - F_{m\tau}), \tag{5}$$

where $\hat{\theta}_{L\tau}$ is the last analyst's true posterior estimate of earnings in a firm–quarter τ, given all available information. We index the strategically chosen bias because it should rise with the amount of uncertainty the last analyst faces about earnings (and hence fall with the number of other forecasts/information at his disposal).

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

$$\begin{aligned} F_{L\tau} - F_{m\tau} &= \hat{\theta}_{L\tau} + a_L(\hat{\theta}_{L\tau} - F_{m\tau}) - F_{m\tau} \\ &= (1 + a_L)(\hat{\theta}_{L\tau} - F_{m\tau}) \end{aligned}$$

which implies that

$$(\hat{\theta}_{L\tau} - F_{m\tau}) = \frac{F_{L\tau} - F_{m\tau}}{1 + a_L}. \tag{6}$$

We can use this relationship to express $\mathrm{ERROR}_{L(it)}$ as the sum of a true (unobserved) forecasting error for the last analyst, $\frac{\hat{\theta}_{L\tau} - E_\tau}{P_{\tau-1}}$, which should be orthogonal to everything in the last analyst's information set, plus a bias term that we can write in terms of SFD_{it}, the difference between the last forecast and the consensus forecast:

$$\mathrm{ERROR}_{L(it)} = \frac{\hat{\theta}_{L\tau} - E_\tau}{P_{\tau-1}} + \left(\frac{a_L}{1 + a_L}\right) \mathrm{SFD}_{it} \tag{7}$$

Subtracting $\left(\frac{a_L}{1+a_L}\right) \mathrm{SFD}_{it}$ 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 in the analyst's information set. Hence, the sum of the OLS coefficients on SFD_τ, $\beta_1 + \beta_3 ln(L_\tau)$, provide an unbiased estimate of $\frac{a_L}{1+a_L}$. Thus, an unbiased estimate of the strategic bias coefficient when there are L analysts is

$$\hat{a}_L = \frac{\beta_1 + \beta_3 ln(L)}{1 - (\beta_1 + \beta_3 ln(L))}.$$

In turn, given $F_{L\tau} - F_{m\tau}$, we can back out the expected dollar bias in the forecast as

$$(\beta_1 + \beta_3 ln(L))(F_{L\tau} - F_{m\tau}).$$

To see this, note that the expected bias is

$$a_L(\hat{\theta}_{L\tau} - F_{m\tau}) = \frac{a_L}{1 + a_L}(F_{L\tau} - F_{m\tau}).$$

Substituting for our estimate of a_L yields

$$\frac{\hat{a}_L}{1 + \hat{a}_L} = \beta_1 + \beta_3 ln(L).$$

Table 4: Economic Impact

This table reports the results from firm–fixed effects regressions of ERROR in last forecast expressed as a percentage of last quarter's stock price, P_{t-1}) on the difference in the last and outstanding mean forecasts, SFD $= \frac{F_L - F_m}{P_{t-1}}$. Robust standard errors (in parentheses) were computed using the Hubert-White method:(*), (**) and (***) indicate significant at 90%, 95% and 99% levels.

Dependent variable: $\text{ERROR}_{L(it)} = \frac{F_{L(it)} - E_{it}}{P_{i(t-1)}}$

SFD	$ln(L)$	$ln(L)*SFD$	$lnsize$	$lnsize*SFD$	$lndays$	$lndays*SFD$	R^2
0.833	0.047	−0.067					0.443
(0.045)***	(0.008)***	(0.028)**					
1.245			−0.061	−0.041			0.446
(0.176)***			(0.005)***	(0.013)***			
0.815	0.049	−0.065			0.004	0.005	0.443
(0.102)***	(0.008)***	(0.031)**			(0.003)	(0.025)	

Table 4 presents the results from these regression analyses. To account for possible correlation in the forecast errors, standard errors are computed using the Hubert-White (robust cluster) method. The results strongly support the hypothesis that the bias in the last analyst's forecast is due to some form of convex relative performance compensation and is strategically chosen. That is, if the last analyst's forecast overshoots the consensus, he tends to overshoot earnings, but the amount by which he overshoots falls almost uniformly with measures of the amount of information at his disposal (number of analysts or firm size).

The amount by which the analyst overshoots is economically very large. For example, the second row (in Panel A) indicates that if two analysts follow the firm, then for every one percent of the stock price that the last forecast overshoots the consensus, on average it overshoots true EPS by 0.82%, implying that the last analyst introduces a strategic forecast bias of about **3.6** times the difference between *last* analyst's (unobserved) true posterior estimate of earnings and the consensus forecast; while if 20 analysts follow the firm, the last forecast overshoots true EPS by 0.77%, implying a strategic bias that is about 1.7 times the difference in the analyst's (unobserved) true estimate of earnings and the consensus forecast.

The last row Table 4 (Panel B) shows that the forecast bias is qualitatively unaffected by the number of days between the last analyst's forecast and the earnings announcement date. Our findings are not due to the fact that the last analyst may have more information at his disposal just because he issues a forecast closer to the announcement date.

3.2.2 Does the market unravel the bias?

Our analysis reveals that the last analyst's reported forecast over–emphasizes his private information: his forecast tends to over–shoot EPS away from the outstanding consensus. This raises the question: Are investors systematically fooled by the bias in the last analyst's forecast? For example, a large positive deviation in the last forecast away from the consensus, $(F_L - F_m)$, may mislead investors who do not account for the bias in this forecast into over–estimating by how much earnings will go up. Investors who treat the last report as unbiased will over–react to their forecasts of earnings. Then, if the 'correct' information does not leak out until near the earnings announcement date, a large positive $(F_L - F_m)$ should be associated with subsequent disappointment and a downward revision in prices around the earnings announcement, and hence negative excess returns. Conversely, a large negative $F_L - F_m$ should be associated with positive excess returns around the earnings announcement. In this section, we explore whether investors are systematically fooled by the forecasts, by determining if there is a systematic relationship between the degree to which the last forecast overshoots the consensus and excess returns around the earnings announcement date.

To obtain excess returns, we first estimate a market model for each firm–quarter:

$$r_{it} = \gamma_0 + \gamma_1 r_{mt} + \epsilon_t, \tag{8}$$

where r_{it} is the daily return on security i and r_{mt} is the return on the `CRSP` value–weighted market index. For each firm–quarter, the market model regression (equation (8)) uses returns for the period $t = -270$ through to $t = -31$, where time $t = 0$ denotes the day of actual earnings announcement by the firm. The least squares estimates, $\hat{\gamma}_0$ and $\hat{\gamma}_1$, are then used to calculate the cumulative abnormal returns (CAR) around the earnings announcement date for each firm–quarter τ:

$$CAR_{[\underline{t},\bar{t}]} = \sum_{t=\underline{t}}^{\bar{t}} \left[r_{\tau t} - (\hat{\gamma}_0 + \hat{\gamma}_1 r_{mt}) \right], \tag{9}$$

where $\underline{t}$ and $\bar{t}$ are starting and end points of the period over which abnormal returns are cumulated.

In our analysis, it is important to exclude the market's reaction to the information in the last forecast. For instance, if the last forecast exceeds the consensus, then this good news will be reflected in the price and lead to immediate positive excess returns; cumulating those immediate returns may cause us to fail to find evidence of bias, since the market's failure to incorporate the bias leads to predicted negative excess returns.

Table 5: Market reaction: Do investors unravel bias?

Robust standard errors (in parentheses) were computed using the Hubert-White method: (*), (**) and (***) indicate significant at 90%, 95% and 99% levels. Regressions have R^2 of about 1%. There are 56,967 observations.

Dependent variable: $\mathrm{CAR}_{[-2,1]}$

SFD	$ln(L)$	$ln(L)*SFD$	$lnsize$	$lnsize*SFD$	$lndays$	$lndays*SFD$
0.546	−0.257	−0.411				
(0.203)***	(0.141)*	(0.137)***				
−0.429	0.012	−0.441	−0.618	0.085		
(0.787)	(0.143)	(0.135)***	(0.103)***	(0.064)		
0.398	−0.248	−0.399			0.021	0.046
(0.426)	(0.148)*	(0.139)***			(0.065)	(0.122)

We then regress $\mathrm{CAR}_{[2,-1]}$ on SFD, using analyst coverage to control for information available to the last analyst. Table 5 presents the findings. To disentangle the impact of new information from the reaction to the last forecast, we drop all observations for which the actual earnings per share was reported less than three days after publication of the last forecast. Our findings are not systematically affected by other choices. Qualitatively, the following empirical regularities emerge[16]:

- Mean abnormal returns, cummulated over three–day period around earnings announcement date, are generally negatively correlated with SFD *when* more analysts follow the firm (and hence, greater investor interest).
- The negative impact of SFD on mean cumulative abnormal returns for greater analyst–coverage holds regardless of firm–size and number of days since the release of the *last* forecast.

The data indicate that, for firms followed by more analysts, if the last forecast overshoots the consensus then the stock market price response to the actual earnings announcement is in the opposite direction. Further, the greater is the difference between the last forecast and the consensus forecast, the greater is the expected change in stock price. In summary, the data suggest that investors appear to be fooled by the last analyst's forecast, missing the bias in his report and consequently attaching an undue weight to the last forecast relative to the consensus.

[16] Similar empirical regularities regarding percentage mean cumulative abnormal returns emerge for other windows.

3.3 Discussion

Our frequency findings that the last forecast tends to overshoot earnings in the direction away from the consensus are consistent with both

(i) our strategic model of analyst forecasting and convex relative performance compensation; and

(ii) a myopic *last* analyst who fails to incorporate the information contained in the consensus forecast into his own forecast. If the last analyst ignores the information contained in earlier forecasts, and myopically issues an unbiased forecast based only on his private information, then, true earnings will tend to be between the consensus and the last forecast (see Ehrbeck and Waldmann (1996)).

Hence, both our strategic model with convex relative performance compensation and analyst myopia predict that the analyst should overshoot true earnings in the direction away from the consensus with a probability exceeding 0.5.

However, our regressions do shed some light on these alternative explanations. If the bias is due to strategic behavior by a rational last analyst with convex relative performance compensation, then for any given difference between the consensus and his forecast, $F_L - F_m$, the error in his forecast should be less on average if there are more analysts. In contrast, if the last analyst is myopic, and completely ignores all information in the consensus, the opposite prediction obtains: If the consensus forecast accuracy rises with analyst coverage, then, for any given difference between the consensus and the last analyst's forecast, the expected bias in the forecast of a myopic last analyst who ignores all information in the consensus rises with the number of analysts: $E[(F_L - \hat{\theta}_L)|F_L - F_m]$ rises with L.

If the last analyst ignores the information contained in the consensus forecast, as the number of analysts rises, the consensus becomes more accurate so that fixing the difference between the myopic analyst's forecast and the consensus, a myopic last analyst, on average, will overshoot EPS by more. Empirically, however, we find the bias is about 50% higher when there are two analysts than when there are more that 10 analysts. This is strong evidence against the hypothesis that the last analyst ignores the information in the consensus (see also Welch's (2000) findings).

However, weaker versions of analyst myopia could be consistent with our empirical findings. For example, the last analyst may only be slightly over-confident, believing that his private signal is more accurate than it is. Then, his forecast will only slightly underweight the information in the consensus, so that the bias in his forecast will fall as the information content in the consensus rises, *i.e,* as the number of analysts rises. Still, the bias we find is economically so large, that it suggests slight analyst over-confidence does not underlie our results. Even when there are more than 10 analysts, on average, for every one percent of price that the last forecast overshoots the consensus, the last forecast also overshoots EPS by more than $\frac{3}{4}$ of one percent. Also, the days between the last analyst's forecast and the earnings announcement, which may capture the quality of the last analyst's information relative to the consensus (and hence the weight an over-confident analyst should place on the consensus), does not significantly affect the forecast bias.

Finally, analysts are experienced and have huge financial stakes in generating better forecasts. If ever one should place weight on rational explanations of outcomes/decision-making, it should be in environments such as this, where agents are well–trained and have strong economic incentives.

4 Conclusion

This paper presents a compensation–based model to explain the bias in individual analysts' forecast of earnings per share. We derive the implications of different forms of relative performance compensation for the pattern of forecasts.

We then develop a frequency test for forecast bias that is robust to correlated signals among analysts, common unforecasted shocks to firms earnings, and information arrival. We find overwhelming evidence that the last forecast is not an unbiased reflection of that analyst's information: The last analyst strategically reports a strongly contrarian forecast, biasing his forecast *away* from the consensus forecast — 60% of the time, the last forecast overshoots actual EPS in the direction away from the outstanding mean forecast.

We then derive structural estimates of the economic magnitude of the bias. The bias is large: on average, for every one percent of price that the last analyst's forecast overshoots the consensus, it overshoots earnings by between one half to three quarters of a percent, depending on the analyst following. The forecast bias falls with analyst following, suggesting that the bias is chosen strategically, and is due to some form of convex relative performance compensation for analysts. Our estimates imply that a better estimate of EPS than either the last forecast or the average of all forecasts could be obtained by accounting for the bias in the last forecast (even just by averaging the last and consensus forecasts). Finally, we find that investors fail to completely unravel the bias in the last forecast.

Appendix

Proof of Proposition 1: If relative performance compensation is linear then an analyst's compensation becomes
$$w_j = \overline{w} - (1+\lambda)|F_j - E| + |F_m - E|.$$
Consider the last analyst to report his forecast. Since the $(L-1)$ preceding forecasts have been announced, F_L will not affect F_m, the mean of all other forecasts. Hence, maximizing $E(w_L|\hat{\theta}_L)$ is equivalent to minimizing
$$E\left[|F_L - E| \mid \hat{\theta}_L\right] = \int_E (|F_L - E|)dG(E|\hat{\theta}_L)$$
To show that the optimal forecast is $F_L = \hat{\theta}_L$, we prove that the expected forecast error is greater if the last analyst reports any other forecast.

Suppose he reports $F_L = \hat{\theta}_L - \xi$ (a lesser forecast), where $\xi > 0$. Then the difference in expected errors is
$$\begin{aligned}
\int_E (|\hat{\theta}_L - E| - |\hat{\theta}_L - \xi - E|)dG(E|\hat{\theta}_L) &= \int_{E \leq \hat{\theta}_L - \xi} \xi \, dG(E|\hat{\theta}_L) + \int_{\hat{\theta}_L - \xi}^{\hat{\theta}_L} (2\hat{\theta}_L - 2E - \xi)dG(E|\hat{\theta}_L) \\
&- \int_{E > \hat{\theta}_L} \xi \, dG(E|\hat{\theta}_L) \\
&\leq \int_{E \leq \hat{\theta}_L} \xi \, dG(E|\hat{\theta}_L) - \int_{E > \hat{\theta}_L} \xi \, dG(E|\hat{\theta}_L) \qquad (10)
\end{aligned}$$
since $E > \hat{\theta}_L - \xi$ in the second term on the RHS of (10), which implies that $2\hat{\theta}_L - 2E - \xi \leq \xi$. By the symmetry of $g(E|\hat{\theta}_L)$ about $\hat{\theta}_L$, the last term in (10) vanishes. Hence,
$$\int_E |\hat{\theta}_L - E|dG(E|\hat{\theta}_L) \leq \int_E |\hat{\theta}_L - \xi - E|dG(E|\hat{\theta}_L)$$
Similarly, $\forall \xi > 0$, we get $\int_E (|\hat{\theta}_L - E|)dG(E|\hat{\theta}_L) < \int_E (|\hat{\theta}_L + \xi - E|)dG(E|\hat{\theta}_L)$. The last analyst has a larger expected forecast error if he announces a forecast greater than his posterior expectation. Hence the optimal forecast is $F_L = \hat{\theta}_L$.

Proof of Lemma 1: Given his posterior beliefs, the last analyst's expected payoff is
$$E\left[w_L(F_L)| \, \hat{\theta}_L\right] = \overline{w} - \lambda \int_E (|F_L - E|)dG(E|\hat{\theta}_L) + \int_E R(|F_m - E| - |F_L - E|)dG(E|\hat{\theta}_L).$$
Suppose $F_m < \hat{\theta}_L$. We first show that $F_L \not< F_m$. To do this, we prove that for every forecast $F_L < F_m$, there exists some other $F_L > F_m$ that yields a higher payoff. Consider the following two reports by the last analyst: $F_L = F_m + \xi$ and $F_L = F_m - \xi$, where $\xi > 0$. The difference in payoffs will be given by
$$\begin{aligned}
E\left[w_L(F_m + \xi) - w_L(F_m - \xi) \mid \hat{\theta}_L\right] &= \int_{E \leq F_m - \xi} (-2\lambda\xi + R(-\xi) - \\
R(\xi)\ ) \, dG(E|\hat{\theta}_L) &+ \int_{F_m - \xi}^{F_m} (2\lambda(E - F_m) + R(-\xi) - R(2F_m - 2E - \xi)) \, dG(E|\hat{\theta}_L) \\
&+ \int_{F_m}^{F_m + \xi} (2\lambda(E - F_m) + R(2E - 2F_m - \xi) - R(-\xi)) \, dG(E|\hat{\theta}_L) \\
&+ \int_{E \geq F_m + \xi} (2\lambda\xi + R(\xi) - R(-\xi)) \, dG(E|\hat{\theta}_L) \qquad (11)
\end{aligned}$$

Clearly, the sum of the first and fourth terms on the RHS of (11) is positive by the symmetry of $g(E|\hat{\theta}_L)$ about $\hat{\theta}_L$. Secondly, the third term is at least equal to the second term since $g(E|\hat{\theta}_L)$ is single–peaked at $\hat{\theta}_L$.

Therefore,

$$E\left[w_L(F_m+\xi) - w_L(F_m-\xi)|\,\hat{\theta}_L\right] > 0$$

and $F_m + \xi$ dominates $F_m - \xi$, $\forall \xi > 0$.

It is also straightforward to show that $F_L \neq F_m$: that is, it is not optimal for the last analyst to mimic the mean of earlier forecasts by reporting $F_L = F_m$. Because,

$$\begin{aligned} E\left[w_L(F_m+\epsilon) - w_L(F_m)|\,\hat{\theta}_L\right] &= \int_{E\le F_m} (-\lambda\epsilon + R(-\epsilon))\, dG(E|\hat{\theta}_L) \\ &+ \int_{F_m}^{F_m+\epsilon} (\lambda(2E-2F_m-\epsilon) + R(2E-2F_m-\epsilon))\, dG(E|\hat{\theta}_L) \\ &+ \int_{E>F_m+\epsilon} (\lambda\epsilon + R(\epsilon))\, dG(E|\hat{\theta}_L). \end{aligned}$$

Dividing through by ϵ and taking the limit as $\epsilon \to 0$, we get the derivative of $w_L(F_L)$ at $F_L = F_m$ to be

$$\lim_{\epsilon\to 0} E\left[\frac{w_L(F_m+\epsilon) - w_L(F_m)}{\epsilon}\middle|\,\hat{\theta}_L\right] = -\int_{E\le F_m} (\lambda + R'(0))dG(E|\hat{\theta}_L) + \int_{E>F_m} (\lambda + R'(0))dG(E|\hat{\theta}_L)$$

Using the symmetry of $g(E|\hat{\theta}_L)$ about $\hat{\theta}_L$, we get derivative of $w_L(F_L)$ w.r.t F_L at $F_L = F_m$ to be

$$\lim_{\epsilon\to 0} E\left[\frac{w_L(F_m+\epsilon) - w_L(F_m)}{\epsilon}\middle|\,\hat{\theta}_L\right] = 2\int_{F_m}^{\hat{\theta}_L} (\lambda + R'(0))dG(E|\hat{\theta}_L) > 0.$$

That is $w_L(F_L)$ is increasing at $F_L = F_m$. Hence, $F_L > F_m$ if $\hat{\theta}_L > F_m$. The proof that $F_L < F_m$ if $\hat{\theta}_L < F_m$ is analogous.

Proof of Proposition 2: (i) Suppose $F_m < \hat{\theta}_L$, then by Lemma 1, $F_L > F_m$. So it remains to show that $F_L \notin [\hat{\theta}_L, \infty)$. To do this, we show that $w_L(F_L)$ is decreasing in F_L over $[\hat{\theta}_L, \infty)$. Consider $F_L = \hat{\theta}_L + \xi$, where $\xi \ge 0$. The derivative of $w_L(F_L)$ at this point is given by

$$\begin{aligned} \lim_{\epsilon\to 0} E\left[\frac{w_L(\hat{\theta}_L+\xi+\epsilon) - w_L(\hat{\theta}_L+\xi)}{\epsilon}\middle|\,\hat{\theta}_L\right] &= -\int_{E\le F_m} (\lambda + R'(F_m-\hat{\theta}_L-\xi))dG(E|\hat{\theta}_L) \\ &- \int_{F_m}^{\hat{\theta}_L+\xi} (\lambda + R'(2E-\hat{\theta}_L-\xi-F_m))dG(E|\hat{\theta}_L) \\ &+ \int_{E>\hat{\theta}_L+\xi} (\lambda + R'(\hat{\theta}_L+\xi-F_m))dG(E|\hat{\theta}_L). \end{aligned} \tag{12}$$

Consider the second term on the RHS of (12). Since $E < (\hat{\theta}_L + \xi)$, we have $R'(2E - \hat{\theta}_L - \xi - F_m) \ge R'(\hat{\theta}_L + \xi - F_m)$ by the concavity of $R(\cdot)$. Also $R'(F_m - \hat{\theta}_L - \xi) > R'(\hat{\theta}_L + \xi - F_m)$. Therefore,

$$\begin{aligned} \lim_{\epsilon\to 0} E\left[\frac{w_L(\hat{\theta}_L+\xi+\epsilon) - w_L(\hat{\theta}_L+\xi)}{\epsilon}\middle|\,\hat{\theta}_L\right] &< -\int_{E\le\hat{\theta}_L+\xi} (\lambda + R'(\hat{\theta}_L+\xi-F_m))dG(E|\hat{\theta}_L) \\ &+ \int_{E>\hat{\theta}_L+\xi} (\lambda + R'(\hat{\theta}_L+\xi-F_m))dG(E|\hat{\theta}_L). \end{aligned}$$

Exploiting the symmetry of $g(E|\hat{\theta}_L)$ about $\hat{\theta}_L$, we get

$$\lim_{\epsilon\to 0} E\left[\frac{w_L(\hat{\theta}_L+\xi+\epsilon) - w_L(\hat{\theta}_L+\xi)}{\epsilon}\middle|\,\hat{\theta}_L\right] < -\int_{\hat{\theta}_L-\xi}^{\hat{\theta}_L+\xi} (\lambda + R'(\hat{\theta}_L+\xi-F_m))dG(E|\hat{\theta}_L) < 0,$$

$\forall \xi \geq 0$. So the last analyst's payoff is decreasing in $F_L \in [\hat{\theta}_L, \infty)$. Combining this result with Lemma 1, it follows that the last analyst will "locate" between F_m and $\hat{\theta}_L$.

The proof that $\hat{\theta}_L < F_L < F_m$ if $\hat{\theta}_L < F_m$ is analogous.

(ii) Let $F_m < \hat{\theta}_L$ and F_L be the optimal forecast of the second analyst given $\hat{\theta}_L$ and F_m. Then it will suffice to show that if the consensus forecast were to be revised to $F_m^* = F_m - \delta$, $(\delta > 0)$, then the last analyst will report $F_L^* \in (F_L - \delta, F_L)$.

By Proposition 2, we know that $F_L \in (F_m, \hat{\theta}_L)$, and by the optimality of $F_L(\hat{\theta}_L)$, $\frac{\partial w_L}{\partial F_L} = 0$ at $F_L(\hat{\theta}_L)$. But

$$\begin{aligned}\lim_{\epsilon \to 0} E\left[\frac{w_L(F_L+\epsilon) - w_L(F_L)}{\epsilon}\middle| \hat{\theta}_L\right] &= -\int_{E \leq F_m} (\lambda + R'(F_m - F_L))dG(E|\hat{\theta}_L) \\ &- \int_{F_m}^{F_L} (\lambda + R'(2E - F_m - F_L))dG(E|\hat{\theta}_L) \\ &+ \int_{E > F_L} (\lambda + R'(F_L - F_m))dG(E|\hat{\theta}_L) = 0.\end{aligned}$$

Let $H(F_L)$ be defined as

$$\begin{aligned}H(F_L) &= -\int_{E \leq F_m} (\lambda + R'(F_m - F_L))dG(E|\hat{\theta}_L) - \int_{F_m}^{F_L} (\lambda + R'(2E - F_L - F_m))dG(E|\hat{\theta}_L) \\ &+ \int_{E > F_L} (\lambda + R'(F_L - F_m)dG(E|\hat{\theta}_L)) = 0. \end{aligned} \tag{13}$$

Now, suppose the consensus forecast were $F_m^* = F_m - \delta$ $(\delta > 0)$. Then the derivative of w_L at F_L will be

$$\begin{aligned}\lim_{\epsilon \to 0} E\left[\frac{w_L(F_L+\epsilon) - w_L(F_L)}{\epsilon}\middle| \hat{\theta}_L, F_m^*\right] &= -\int_{E \leq F_m - \delta} (\lambda + R'(F_m - \delta - F_L))dG(E|\hat{\theta}_L) \\ - \int_{F_m - \delta}^{F_L} (\lambda + R'(2E - F_m + \delta - F_L))dG(E|\hat{\theta}_L) &+ \int_{E > F_L} (\lambda + R'(F_L - F_m + \delta))dG(E|\hat{\theta}_L). \end{aligned} \tag{14}$$

By a simple change of variable, we can write the RHS of equation (14) as

$$\begin{aligned}-\int_{E \leq F_m} (\lambda + R'(F_m - F_L - \delta))dG(E|\hat{\theta}_L) &- \int_{F_m}^{F_L+\delta} (\lambda + R'(2E - F_m - F_L - \delta))dG(E|\hat{\theta}_L) \\ &+ \int_{E > F_L + \delta} (\lambda + R'(F_L + \delta - F_m))dG(E|\hat{\theta}_L)\end{aligned}$$

which is $H(F_L + \delta)$. By the optimality of F_L,

$$\lim_{\delta \to 0} E\left[\frac{H(F_L+\delta) - H(F_L)}{\delta}\right] < 0$$

which implies that $H(F_L + \delta) < H(F_L) = 0$. Hence,

$$\lim_{\epsilon \to 0} E\left[\frac{w_L(F_L+\epsilon) - w_L(F_L)}{\epsilon}\middle| \hat{\theta}_L, F_m^*\right] < 0$$

and the last analyst will shade his forecast in the direction of revision in the outstanding consensus.

Finally, it is straightforward to show that

$$\lim_{\epsilon \to 0} E\left[\frac{w_L(F_L - \delta + \epsilon) - w_L(F_L - \delta)}{\epsilon}\middle| \hat{\theta}_L, F_m^*\right] > 0$$

which implies that the last analyst will bias his forecast by less than δ in response to the new mean. That is, $0 < \frac{\partial(\hat{\theta}_L - F_L)}{\partial(\hat{\theta}_L - F_m)} < 1$.

Proof of Proposition 3: (i) Suppose $F_m < \hat{\theta}_L$, then by Lemma 1, $F_L > F_m$. So we need to show that $F_L \notin (F_m, \hat{\theta}_L]$. To do this, we prove that $w_L(F_L)$ is increasing in F_L over $(F_m, \hat{\theta}_L]$. Consider two reports by the last analyst: $F_L = \hat{\theta}_L - \xi + \epsilon$ and $F_L = \hat{\theta}_L - \xi$, where $0 \leq \xi < \hat{\theta}_L - F_m$. The payoffs from these two reports can be used to compute the derivative of $w_L(F_L)$ at $F_L = \hat{\theta}_L - \xi$. This is given by

$$\begin{aligned} \lim_{\epsilon \to 0} E\left[\frac{w_L(\hat{\theta}_L - \xi + \epsilon) - w_L(\hat{\theta}_L - \xi)}{\epsilon} \middle| \hat{\theta}_L \right] &= -\int_{E \leq F_m} (\lambda + R'(F_m - \hat{\theta}_L + \xi)) dG(E|\hat{\theta}_L) \\ &- \int_{F_m}^{\hat{\theta}_L - \xi} (\lambda + R'(2E - \hat{\theta}_L + \xi - F_m)) dG(E|\hat{\theta}_L) \\ &+ \int_{E > \hat{\theta}_L - \xi} (\lambda + R'(\hat{\theta}_L - \xi - F_m)) dG(E|\hat{\theta}_L). \end{aligned} \tag{15}$$

Consider the second term on the RHS of (15). Since $E < (\hat{\theta}_L - \xi)$, we have $R'(2E - \hat{\theta}_L + \xi - F_m) \leq R'(\hat{\theta}_L - \xi - F_m)$ by the convexity of $R(\cdot)$. Also $R'(F_m - \hat{\theta}_L + \xi) \leq R'(\hat{\theta}_L - \xi - F_m)$. Hence,

$$\begin{aligned} \lim_{\epsilon \to 0} E\left[\frac{w_L(\hat{\theta}_L - \xi + \epsilon) - w_L(\hat{\theta}_L - \xi)}{\epsilon} \middle| \hat{\theta}_L \right] &> -\int_{E \leq \hat{\theta}_L - \xi} (\lambda + R'(\hat{\theta}_L - \xi - F_m)) dG(E|\hat{\theta}_L) \\ &+ \int_{E > \hat{\theta}_L - \xi} (\lambda + R'(\hat{\theta}_L - \xi - F_m)) dG(E|\hat{\theta}_L). \end{aligned}$$

Again, by the symmetry of $g(E|\hat{\theta}_L)$ about $\hat{\theta}_L$, we get

$$\lim_{\epsilon \to 0} E\left[\frac{w_L(\hat{\theta}_L - \xi + \epsilon) - w_L(\hat{\theta}_L - \xi)}{\epsilon} \middle| \hat{\theta}_L \right] > \int_{\hat{\theta}_L - \xi}^{\hat{\theta}_L + \xi} (\lambda + R'(\hat{\theta}_L - \xi - F_m)) dG(E|\hat{\theta}_L) \geq 0,$$

$\forall\, 0 \leq \xi < \hat{\theta}_L - F_m$. The last analyst's payoff is therefore increasing in F_L over $(F_m, \hat{\theta}_L]$. Hence, the optimal F_L will be such that $F_m < \hat{\theta}_L < F_L$.

The proof that $F_L < \hat{\theta}_L < F_m$ if $\hat{\theta}_L < F_m$ is analogous.

Proof of Proposition 4: Consider two forecasts, $F_L, \tilde{F}_L$, where $F_L, \tilde{F}_L$ are feasible forecasts given convex relative compensation: $\tilde{F}_L > F_L > \hat{\theta} > F_m$ or $\tilde{F}_L < F_L < \hat{\theta} < F_m$. Then F_L is a relatively more profitable forecast than $\tilde{F}_L$ when g_L than when $\tilde{g}_L$:

$$\begin{aligned} &E\Big[-\lambda|F_L - E| + R(|F_m - E| - |F_L - E|) - \Big[-\lambda|\tilde{F}_L - E| + R(|F_m - E| - |\tilde{F}_L - E|)\Big] \mid g_L\Big] \\ &-E\Big[-\lambda|F_L - E| + R(|F_m - E| - |F_L - E|) - \Big[-\lambda|\tilde{F}_L - E| + R(|F_m - E| - |\tilde{F}_L - E|)\Big] \mid \tilde{g}_L\Big] \\ &= \int_E \big[R(|F_m - E| - |F_L - E|) - R(|F_m - E| - |\bar{F}_L - E|)\big](g_L(E) - \tilde{g}_L(E)) dE > 0. \end{aligned}$$

To see this last inequality, note that for any $a < b$, the convexity of $R(\cdot)$ ensures that given equally likely realizations, $E - x$ and $\theta + x$, $x = a, b$, that

$$\begin{aligned} &R(|F_m - E - a| - |F_L - E - a|) + R(|F_m - E + a| - |F_L - E + a|) - \big[R(|F_m - E - a| - |\tilde{F}_L - E - a|) \\ &\qquad + R(|F_m - E + a| - |\tilde{F}_L - E + a|)\big] \\ &> R(|F_m - E - b| - |F_L - E - b|) + R(|F_m - E + b| - |F_L - E + b|) - \big[R(|F_m - E - b| - |\tilde{F}_L - E - b|) \\ &\qquad + R(|F_m - E + b| - |\tilde{F}_L - E + b|)\big] \end{aligned}$$

References

[1] Alexander, John C. and James S. Ang (1985), "The Market Response to Low–Balling Earnings Expectations," *Working paper,* Florida State University.

[2] Bernhardt, Dan, Murillo Campello and Edward Kutsoati, (2001) "Who Herds?", *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] Chevalier, Judith and Glenn Ellison (1997), "Risk Taking by Mutual Fund as a Response to Incentives," *Journal of Political Economy,* 106(6), 1167–1200

[5] Dugar, Amitabh and Siva Nathan (1995), "The Effects of Investment Banking Relationships on Financial Earnings Forecasts and Investment Recommendations," *Contemporary Accounting Research,* 12, 131–60.

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

[7] Foster, George, Chris Olsen and Terry Shevlin (1984), "Earnings Releases, Anomalies, and the Behaviour of Security Returns," *Accounting Review,* 56, 574–603.

[8] Gallo, Gampiero, Clive Granger, and Yongil Jeon (2001), "The Impact of the Use of Forecasts in Infomration Sets" *mimeo* University of San Diego.

[9] Henry, George (1989), "Wall Street Economists: Are they Worth Their Salt?" *Business Economics,* Oct., 44–48.

[10] Holmstrom, Bengt (1982), "Managerial Incentives Problems — A Dynamic Perspective," in *Essays in Economics and Management in Honor of Lars Wahlbeck.* Helsinki: Swedish School of Economics.

[11] Hong, Harrison, Jeffrey Kubik and Amit Solomon (1998), "Security Analysts' Career Concerns and Herding of Earnings Forecasts," *mimeo*, Stanford University.

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

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

[14] Kutsoati (1998) "Dodos, Contrarians and Central Bankers: Essays on Incentives and Contracts", PhD Thesis, Queen's University, Canada.

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

[16] Laster, David, Paul Bennett and In Sun Geoum (1996), "Rational Bias in Macroeconomic Forecasts," Federal Reserve Bank, New York, *Staff Report,* #21

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

[18] Lys, Thomas and Songkyu Sohn (1990), "The Association between Revisions of Financial Analysts' Earnings Forecasts and Security–Price Changes," *Journal of Accounting and Economics,* 13, 341–363.

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

[20] Scharfstein, David and Jeremy C. Stein (1990), "Herd Behavior and Investment", *American Economic Review,* June 1990, 465–79.

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

[22] — (1992) "Reputation and Performance among Security Analysts," *Journal of Finance,* 47, 1811–1836.

[23] Trueman, Brett (1990), "On the Incentives for Security Analysts to Revise Their earnings Forecasts," *Contemporary Accounting Research,* 7(1), 203 –222.

[24] — (1994), "Analyst Forecasts and Herding Behavior," *Review of Financial Studies,* 7(1), 97–124.

[25] Welch, Ivo (2000), "Herding Among Security Analysts", *Journal of Financial Economics*, 58 (3), 369–96.

[26] Zweibel, Jeffrey (1995), "Corporate Conservatism and Relative Compensation", *Journal of Political Economy,* 103(1), 1–25.

Who Herds?*

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

Murillo Campello
The Eli Broad College of Business
Michigan State University
East Lansing, MI 48824

Edward Kutsoati
Department of Economics
Tufts University
Medford, MA 02155

September 11, 2001
Very preliminary

Abstract

We test for forecast bias in earnings forecasts by developing a test statistic that is robust to (i) correlated information amongst analysts; (ii) common industry-wide earnings shocks; and (iii) information arrival at different points during the forecast horizon. We find no empirical support for herding. On the contrary, analysts systematically issue *contrarian* forecasts that overshoot the publicly–available consensus forecast in the direction of their private information. Analysts that issue revised forecasts that are closer to the consensus than their original forecasts are particularly likely to overshoot. We also find that the forecast bias is economically large and declines with the amount of information at the analyst's disposal. The magnitude of the bias, its systematic variation with analyst following, and the pattern of bias in forecast revisions strongly suggest that the bias is strategically chosen.

*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.

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 to the consensus forecast, issuing forecasts that ignore or under-weight their own information (see, for example, Trueman (1990 and 1994) and Hong, Kubik and Solomon (2000), amongst others.) In large part, these herding stories have arisen because forecasts seem to be "more clustered than they should be." 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 the final forecasts are either uniformly too low or too high; and then conclude that the forecasters are herding.

Of course, clustered forecasts need not imply that analysts herd. Firstly, earlier forecasts may contain information that subsequent analysts, as good Bayesians, incorporate into their forecasts. Secondly, 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, then their forecasts will reflect this common (perhaps mis) information, so that forecasts will tend all to be too high or too low relative to the earnings realization. In the extreme case where analysts have identical information, each unbiased analyst will make the same forecast, which will invariably differ from realized earnings. Thirdly, common unanticipated market-wide earnings shocks can cause most, if not all, forecasts to be too low or too high relative to actual earnings. For example, in 2000 almost all earnings forecasts for technology stocks greatly exceeded actual earnings. Thus, the fact that over a period of time forecasts systematically exceed or fall short of earnings, again does not imply that forecasts are biased. 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 (i) correlated signals or common information sources; (ii) common, unforecasted shocks; and (iii) information arrival at different points during the forecast horizon. 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. Finally, the opposite behavior is predicted if an analyst seeks to issue forecasts that distinguishes himself from other analysts.

The key feature that these observations share is that no assumptions are made about how an analyst's information set is formed. Thus, 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 unforecasted earnings shocks, our test statistic then 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 how relative performance compensation affects analyst forecasts. Their 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, or recent forecasts, (2) whether the forecast was a revision, (3) analyst coverage, (4) the number of prior forecasts, (5) the time since the previous forecast, 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). That is, analysts systematically issue *contrarian* forecasts that overshoot the consensus forecast in the direction of their private information. The conditional probability that an analyst's forecast overshoots 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 and information arrival. These findings are all the opposite of those predicted by herding.

We also explore the bias in revised forecasts. Forecast bias is *particularly* high for analysts who adjust their forecasts to be *closer* to the consensus than their original forecasts. Such forecasts *overshoot* earnings 68.5% of the time! Apparently, these analysts initially staked out extreme positions, and upon receiving new information, found their positions too outlandish, revised their forecasts closer to their estimate earnings, but still maintained extreme positions. In contrast, revised forecasts that moved further from the consensus were essentially unbiased, suggesting that these analysts were able to separate sufficiently from others without having to bias their forecasts. Thus it appears that analysts choose their contrarian forecasts strategically.

We then show that the forecast bias is *economically large*. Under the assumption that forecast bias is a linear function of the expected error in the consensus forecast given the *analyst*'s information, we can derive estimates of both the strategic bias chosen and its expected dollar impact, as a function of the observable difference between the last forecast and the consensus forecast. We find that if 20 analysts follow a firm in a particular quarter, then for every one percent of stock price that the second analyst's forecast exceeds the first forecast, on average, his forecast overshoots actual earnings by about 1.1%, *ex post*. However, for the same difference between the forecast and the outstanding consensus, the 20th analyst to report a forecast, on average, overshoots earnings by only 0.77% of the stock's price. This implies a bias of about 2.58 times the difference between the analyst's (unobserved) true estimate of earnings and the consensus forecast if he were the 2nd to report, while if he reports last (20th), then the implied bias is approximately 0.99 times. Thus, even though the forecast bias falls with the amount of information at the analyst's disposal, it remains economically significant.

Most earlier attempts at detecting herding among security analysts 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 to the consensus forecast, and are more likely to lose their jobs for making forecasts that deviate from the consensus.[1] A concern with this testing strategy is that it does not account for the possible correlation of information or for information arrival over time. Welch (2000) takes a related approach. Using a dataset of analysts' recommendations (on whether to buy, hold or sell a stock), 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 well reflect the fact the 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 per share." As in the current paper, Keane and Runkle attempt to control for correlated signals, but their methodology causes them to restrict their analysis to a set of only 21 very heavily followed firms. We find that forecasts biases are smallest for these types of firms, and conjecture that sample selection biases are likely to explain Keane and Runkle's failure to uncover the contrarian bias. Our findings, on the

[1] Lamont (1995) offers similar finding for forecasters of GNP and other macroeconomic indicators. See also Laster, Bennett and Geoum (1999), and Ehrbeck and Waldmann (1996).

other hand, support the theory developed in Bernhardt and Kutsoati (2001) that analysts care about both absolute and relative forecast accuracy, and that compensation is a convex function of relative forecast accuracy — analysts gain more from outperforming other analysts than they lose from underperforming. Such relative forecast compensation arises naturally if the forecast industry uses relative forecast accuracy to gauge analyst ability.[2] Finally, our results complement those of Zitzewitz (2001), who finds that analysts may 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. In Section 3, we explain in detail our sample selection procedure, present our empirical findings, and explore their economic significance. Section 4 concludes.

2 Do analysts herd?

This paper defines herding to be a choice by an analyst to announce a forecast of EPS that is closer to a *publicly known consensus* than his own information suggests. His information includes both any signals that his own analysis uncovers, as well as public information releases by the firm, and the forecasts of earlier analysts. For example, herding may arise if the market uses relative forecast accuracy to judge ability, and the analyst is risk averse about the possibility of being fired for poor relative performance. If, instead, the rewards from better relative forecast accuracy exceed the losses from bad performance, analysts may adopt *anti-herding* behavior, reporting contrarian forecasts that overshoot their private estimate of EPS in the direction away from the consensus.

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 beliefs about EPS given his information, 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. If an analyst reports an unbiased forecast, then $F_{ijk} = \hat{E}_{ijk}$. While we do not observe analyst i's posterior forecast of earnings, $\hat{E}_{ijk}$, if a forecast is unbiased, then it should be as likely to exceed earnings per share as fall short, both unconditionally, and conditional on anything in his information set. In particular, the likelihood with which a forecast exceeds actual earnings should not vary with the amount by which the forecast exceeded or fell short of the extant consensus forecast.

If, instead, the analyst 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

[2] 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.

realized earnings more than half of the time. A herding analyst will trade off absolute and relative forecast accuracy, so that the greater is the (percentage of price) difference between $\overline{F}_{ijk}$ and F_{ijk}, the more pronounced should be this relationship.

Finally, if an analyst tries to distinguish himself from the others, then if his private estimate of earnings exceeds the consensus, his forecast will 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 his forecast exceeds the consensus, then the forecast should exceed realized earnings more than half of the time; and if his forecast is less than the consensus, the forecast should fall short of realized earnings more than half of the time. Again the greater is the difference between $\overline{F}_{ijk}$ and F_{ijk}, the more pronounced this relationship should be.

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 outstanding consensus. For example, z_τ^+ might be the event that (i) forecast F_τ exceeded the outstanding consensus, $\overline{F}_\tau$, by at least x percent of the stock price at the end of the previous quarter, *i.e.* $\frac{F_\tau - \overline{F}_\tau}{P_{\tau-1}} > x$, (ii) the analyst was the third analyst to report, and (iii) the analyst had issued a prior forecast, so that F_τ represented a revised forecast. Similarly, z_τ^- might be the event that $\frac{\overline{F}_\tau - F_\tau}{P_{\tau-1}} > x$, and F_τ was a revised forecast issued by the third analyst.

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

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

and 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 of unbiasedness

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

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

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

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,

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

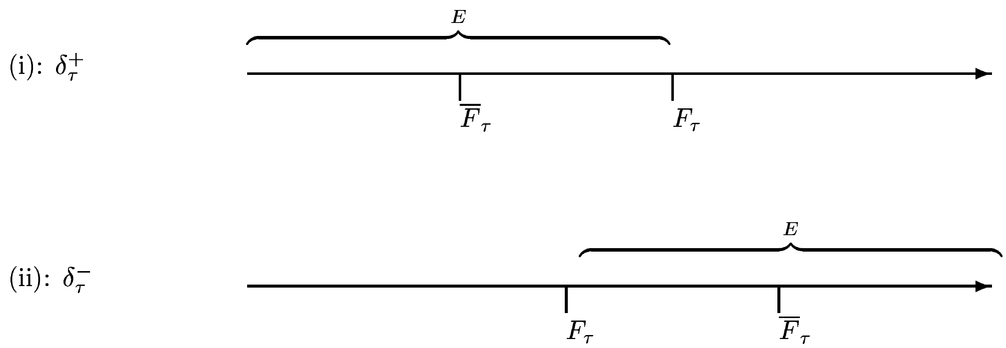

Figure 1: Realization of actual EPS relative to F_τ and $\overline{F}_\tau$

Figure 1 illustrates the events δ_τ^+ and δ_τ^-. If the 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, the 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 the sample averages to estimate the population probability of overshooting,

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

Here $\frac{\sum_\tau \delta_\tau^+}{\sum_\tau \gamma_\tau^+}$ is our estimate of the conditional probability of overshooting actual earnings given z^+; while $\frac{\sum_\tau \delta_\tau^-}{\sum_\tau \gamma_\tau^-}$ is our estimate of the conditional probability of falling short of true earnings given z^-. $\mathbf{S}$ is therefore the average of our estimates of the two conditional overshooting probabilities.

Note that the properties of S do not depend on how analysts forms 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, $\mathbf{S}$ 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, $\mathbf{S}$ 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

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

Thus, under the null of median unbiasedness, asymptotically,

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

Importantly, 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* $g(\cdot)$. *Hence, neither the mean of* $\mathbf{S}$ *nor the bound on its variance are affected by variations in* ω_τ *and* $g_\tau(\cdot)$ *across firm quarters. Thus, with regard to hypothesis testing, if the null hypothesis of unbiasedness is rejected using the upper bound on the variance, then it can always be rejected.*

In sum, $\mathbf{S}$ is a sufficient statistic for test of bias in analysts's forecasts. If forecasts are unbiased, then $\mathbf{S} = 0.5$. The alternative of herding implies that $\mathbf{S} < 0.5$; and anti-herding implies that $\mathbf{S} > 0.5$.

Finally, we 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, for example, 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.[3] After 1988, forecasts 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. Finally, 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 first filter the data for likely data entry errors, deleting any forecast with an absolute error value exceeding \$10.[4] Some forecasts are dated over 360 days before the actual earnings report, and are likely 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 the total forecasts, including revisions, reported in each firm-quarter.

It is common practice for analysts to revise their old forecasts. Consequently, we only focus on the most recent forecast reported by each analyst when processing the consensus forecast. For every firm-quarter, we sort all forecasts by the order (*i.e.*, date) of release, from the oldest to the most recent forecast. 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 hypothetical 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.

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

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

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

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 27 April 1992 (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 15 May 1992. 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 other forecasts made within the previous ℓ days, we set the lag to be $\ell = 3$ days.[5]

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

[5]The results are qualitatively similar for $\ell = 5$ and $\ell = 7$.

Table 2: Sample Statistics

	min	max	Percentiles 25	50	75	Sample moments Mean	Std. dev.	Skewness	Kurtosis
Number of analysts (`cover`)	2	38	6	9	14	10.12	5.75	0.88	3.72
Forecasts–Analysts ratio	1	4	1.14	1.33	1.57	1.39	0.34	1.19	4.92
Consensus *error* (as a percent of stock price)	-48.5	160.4	-0.07	0.001	0.18	0.23	2.0	24.48	1212
Forecast *error* (as a percent of stock price)	-72.6	159.4	-0.09	0	0.11	0.14	1.87	24.59	1357
Forecast minus *consensus* (as a percent of stock price)	-100.58	59.80	-0.10	-0.01	0.04	-0.08	0.94	-19.71	1320
Days between *any* two forecasts	0	116	1	6	13	9.58	12.46	2.42	10.4
Days between *own* forecasts (revised forecasts only)	0	117	21	36	56	38.59	22.08	0.32	2.34

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* in the last and consensus forecasts. To remove the effects of outliers in percentage errors and differences, we delete observations with stock price less than \$5.00 (since we do not want small denominators (see also, Lim (2001)). The mean stock price in our sample is about \$35.00, suggesting that analysts on average overshot earnings by about \$0.05. The summary statistics shows that the distribution of the percentage forecasts error has extreme fat tails. Such outliers may have excessive impact on estimates of mean regressions using these variables. We address this problem in section 3.3.

3.2 Results of Frequency Test

We first report the frequency test of herding conditional on the events that the forecast exceeded the consensus or not. 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 on average, $\mathbf{S}(0)$ equals 0.60 for $\ell = 3$. That is, analysts tend to overshoot earnings in the direction away from the consensus about 60% of the time. Figure 2 shows that $\mathbf{S}(x)$ rises rapidly as we increase x: larger deviations from the consensus substantially raise the likelihood of overshooting.

The qualitative pattern of overshooting varies little with analyst–coverage (number of analysts following firm in that quarter), the *order* of release of forecast given analyst following, or the forecast-year. The probability of overshooting is slightly higher for firms followed by fewer analysts,

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

	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 [0.599, 0.602]	0.464	0.579	0.622
1989 – 1992	88,890	0.607 [0.604, 0.611]	0.551	0.673	0.542
1993 – 1995	82,180	0.603 [0.599, 0.606]	0.440	0.555	0.650
1996 – 1999	101,802	0.590 [0.587, 0.593]	0.406	0.514	0.667
Segmented by aAnalysts following (`cover`)					
`cover` ≤ 5	60,672	0.605 [0.601, 0.609]	0.450	0.571	0.639
$6 \leq$ `cover` ≤ 14	150,248	0.605 [0.602, 0.607]	0.470	0.590	0.620
`cover` ≥ 15	61,952	0.586 [0.581, 0.590]	0.459	0.561	0.611
Segmented by timing of forecast (`order`)					
`order` $= 3$	34, 864	0.609 [0.603, 0.614]	0.479	0.599	0.619
`cover` = 2nd-to-last	49,855	0.593 [0.589, 0.598]	0.431	0.541	0.648
`cover` = last	58,339	0.592 [0.588, 0.596]	0.421	0.528	0.657
Segmented by timing `order` in small/large group `cover`					
`order` $= 2$ if `cover` ≤ 5	23,889	0.617 [0.611, 0.623]	0.480	0.608	0.626
`order` $= last$ if `cover` ≤ 5	29,875	0.598 [0.592, 0.604]	0.438	0.550	0.648
`order` $= 2$ if $6 \leq$ `cover` ≤ 14	11,965	0.613 [0.604, 0.622]	0.544	0.661	0.564
`order` $= last$ if $6 \leq$ `cover` ≤ 14	24,045	0.587 [0.581, 0.593]	0.407	0.509	0.664
`order` $= 2$ if `cover` ≥ 15	1,949	0.581 [0.559, 0.603]	0.567	0.652	0.509
`order` $= last$ if `cover` ≥ 15	4,419	0.581 [0.565, 0.596]	0.377	0.469	0.685
Segmented by days between forecasts (`days`)					
$3 \leq$ `days` ≤ 7	80,525	0.611 [0.607, 0.614]	0.467	0.591	0.631
`days` ≥ 8	99,074	0.609 [0.606, 0.612]	0.448	0.569	0.648

NOTES:–
Table reports average conditional probability, **S**, developed in the text. 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 are reported in square brackets.

and earlier analysts are slightly more likely to overshoot than later analysts. Also note that the forecast bias of the second analyst to report falls with the number of analysts who follow him. Finally, there is almost no variation in the nature of bias in forecasts released within 3 to 7 days after the preceding forecast and those that were issued at least a week later.

Notice that $\mathbf{S}(0)$ varies little with the *order* of release in a large group (Panel C), even though the mean probability that a forecast exceeds true earnings falls if he reports late in the forecast period. Table 3 also shows the large impact of common unforecasted earnings shocks from year to year — the conditional probability that a forecast falls short of earnings given that it fell 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 shocks, and our success in doing so.

The frequency tests are inconsistent with herding among analysts. Rather, the evidence suggests analysts exaggerate their private information. An analyst's forecast exhibits a contrarian bias **both** when they have positive and negative private information, relative to the consensus (see the last 3 columns). That is, although, on average, forecasts fell short of earnings 54% of the time in the full sample period, a forecast tends to exceed earnings about 58% of the time when it exceeds the consensus: $\Pr(F_\tau > E | F_\tau > \overline{F}_{\tau-3})$ equals 0.579 when $\Pr(F_\tau > E) = 0.46$. So, too, when the forecast is less than the consensus then it tends to fall short of actual earnings, $\Pr(F_\tau < E | F_\tau < \overline{F}_{\tau-3}) = 0.622$.

Next we explore the forecast bias in *revised* forecasts. In particular, we look at the pattern in analysts' revised forecasts, relative to the outstanding consensus both at the time they issue their older and revised forecasts. This pattern sheds more light on how analysts strategically bias their forecasts, even after it has been revised. An interesting feature of forecast revisions is that analysts are far more likely to revise forecasts down than up — 66.5% of the time, the revised forecast offers a reduced/more pessimistic earnings forecast, suggesting that forecasts are issued strategically. Bernhardt and Campello (2001) provide evidence that these downward forecast revisions may be induced by firms seeking to reduce the average analyst forecast, in order to generate a 'positive earnings surprise.'

We decompose revised forecasts in two distinct ways. We first decomposed revised forecasts according to whether the forecast was revised in the *same* or *opposite* direction as the consensus change; and whether the forecast–consensus difference is *smaller* or *larger* after the revision. A revised forecast is said to be *trend–breaking* if it is revised in the direction *opposite* of the change in the consensus forecast since the analyst's older forecast was reported. A revised forecast is *trend–following* if the revision is in the same direction as change in the consensus forecast. About 40% of all revisions are in the opposite direction of the change in consensus (*trend–breakers*).

We also decompose revised forecasts according whether the revised forecast, F^r_τ, is closer to (or further away from) the most recent consensus, than the initial forecast, F^o_τ, is: *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. And, finally, we group the sample into those revisions that were in a "direction" of the *old* consensus and those "away" from it, regardless of the direction of revision in the consensus. That is, in one group we have forecasts that were revised *down* but still exceed the consensus or revised *up* but still lower that the consensus; and the second group has forecasts that were revised *up* and exceeded the consensus or those that were revised *down* and lower than the consensus. About 75% of all forecast revisions are in a direction *further away* from the *old* consensus. This latter decomposition, shown in the last Panel of Table 4, in part reveals those extreme forecasts that necessarily needed to be revised, and what the over-shooting probability would have been had they not been revised.

On first pass, Table 4 suggests that revised forecasts are significantly less biased than those that are not revised — for the entire sample of revised forecasts, revised forecasts overshoot earnings only 54% of the time, far less than the 60% rate at which all forecasts overshoot. Note that this reduced overshooting rate is *not* due to the fact that the revised forecasts are issued 'later', as Table 3 indicated that the last analyst to report was only slightly less biased than earlier forecasts, overshooting 59.2% of the time.

Table 4 also indicates that there is no significant difference in the frequency with which trend–breaking and trend–following forecasts overshoot. More interesting, revisions that *raise* the difference between the consensus forecast and the analyst's forecast are *not* more likely to overshoot, in contrast to what analyst myopia theories would predict.

However, the results on forecast bias when we decompose forecast revisions according to whether revisions were (i) toward the *old* consensus, or "away" from it (regardless of the direction of revision in the consensus); or (ii) toward the *new* consensus, or away from it; reveals that the bias characterization is far more subtle. First, note that if a forecast was revised, then the original forecasts are truly awful, suggesting that analysts are reluctant to revise forecasts. The revised forecasts are significantly better than the original forecasts. Next, observe that forecast revisions that move *away* from the old consensus, or *away* from the new consensus are essentially median unbiased. overshooting true earnings 48% of the time. In very sharp contrast, forecast revisions that move toward the consensus measures, especially the old consensus, while more accurate than the original forecasts, do not move nearly far enough. In particular, forecasts that move toward the old consensus remain extreme, overshooting earnings in the direction away from the consensus 68.5% of the time, an overshooting rate that substantially exceeds that for the entire population of forecasts.

Table 4: Test of Bias when Analyst Revises Forecasts

Sample	# obs.	**S**	Pr($F_\tau > E$)	Pr($F_\tau > E \mid z_\tau^+$)	Pr($F_\tau < E \mid z_\tau^-$)
			Conditional Frequencies		
Revised forecasts	81,494	0.539	0.446	0.496	0.582
		[0.535, 0.542]			
Direction of revision relative to consensus revision					
Trend-breakers	32,751	0.552	0.452	0.519	0.585
		[0.546, 0.558]			
Trend-followers	48,743	0.530	0.442	0.480	0.579
		[0.525, 0.534]			
Forecast–consensus difference after revisions					
Difference smaller (than before)	44,871	0.549	0.442	0.498	0.599
($\lvert F_\tau^r - \overline{F}_\tau^r\rvert < \lvert F_\tau^o - \overline{F}_\tau^o\rvert$)		[0.544, 0.553]			
Difference larger (than before)	36,623	0.528	0.452	0.492	0.564
($\lvert F_\tau^r - \overline{F}_\tau^r\rvert > \lvert F_\tau^o - \overline{F}_\tau^o\rvert$)		[0.522, 0.535]			
"Location" of old & revised forecast, relative to *current* consensus					
Revision *closer* to current consensus:	39,979	0.568	0.450	0.525	0.611
($\lvert F_\tau^r - \overline{F}_\tau^r\rvert < \lvert F_\tau^o - \overline{F}_\tau^r\rvert$)		[0.563, 0.573]			
original forecast		0.845	0.562	0.853	0.837
		[0.841, 0.850]			
Revision *further* from current consensus	41,515	0.505	0.443	0.451	0.560
($\lvert F_\tau^r - \overline{F}_\tau^r\rvert > \lvert F_\tau^o - \overline{F}_\tau^r\rvert$)		[0.500, 0.511]			
original forecast		0.571	0.609	0.724	0.417
		[0.566, 0.575]			
Revisions toward/away from old consensus					
Revision *toward* old consensus	20,591	0.685	0.456	0.624	0.747
		[0.679, 0.692]			
original forecast		0.824	0.535	0.828	0.821
		[0.819, 0.830]			
Revision *away* from old consensus	60,903	0.481	0.443	0.416	0.546
		[0.476, 0.486]			
original forecast		0.346	0.659	0.411	0.281
		[0.338, 0.354]			

NOTES:–
Table reports **S** values for revised forecasts, and also when grouped in distinct ways (see text). 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 are reported in square brackets.

How might we interpret these results? One possibility is that analysts care about both absolute forecast accuracy, and forecast accuracy relative to the consensus. Consider first the median unbiased result for revised forecasts that move further away from the consensus. 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 remains median unbiased. In contrast, an analyst who revises his forecast toward the consensus only does so when his original forecast is 'too' extreme given his updated information, but, seeking to distinguish his forecast from the consensus, his revised forecast does not go far enough.

3.3 Economic Significance of Forecast Bias

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

We first estimate the relationship between forecast error and the difference between the forecast and the 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,[6]

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

The equation we estimate is

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

Under the null hypothesis that an analyst's forecast is unbiased, one can interpret the dependent variable $Error_\tau = \frac{F_\tau - E_\tau}{P_{\tau-1}}$ 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

[6]Note that normalizing by share price causes our measures to be more sensitive to deviations by 'value' stocks and stocks with declining earnings than they are for growth stocks. To reduce this sensitivity and mitigate the effects of outliers on our results, we drop firm-quarters with stock price less that $5.00.

estimates should be zero, since the forecast error would be orthogonal to everything in the analyst's information set, including then 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`, *lnorder*, and the log of the number of analysts following the firm, *lncov*. To understand why we do this, notice that if the consensus contains more information but an analyst ignores this, then his error will be larger for every percentage deviation from the consensus. And to the extent that a greater analyst following implies a more precise consensus, the error in forecasts will vary with `order` and `cover`: `order` x `cover`.

If a forecast is unbiased, then Error_τ should not be systematically related to SFD_τ. The interaction between both SFD_τ and analyst coverage, and SFD_τ and the timing of forecast should not be significantly different from zero. If analysts *herd* towards the consensus, then the coefficient on SFD_τ should be negative. On the other hand, if analysts are exaggerating their private information, then we expect the coefficient to be positive. Finally, if analysts bias their forecasts to separate from other forecasts, the coefficients on interaction between SFD_τ and the logarithm of the number of analysts and timing of forecasts should be negative, since the bias should fall if an analyst is less uncertain about earnings.

We first estimate the relationship in equation (2) using ordinary least squares (fixed-effects) regression. With OLS fixed-effects, we control for possible correlation of forecast errors within firms and hence obtain robust standard errors computed using the Hubert-White (robust cluster) method in STATA (see Rogers (1993)). However, a significant concern with OLS regression is that given the massively fat tails and left-skewness in the distribution of forecast errors (see Table 2), outliers in the tail of especially those in the left tail will have an excessive impact on estimates. As a result, OLS fixed-effects may give rise to estimates that do not fairly represent the relationships between the dependent and independent variables.

We address this concern in our estimations in a number ways. First, before running the OLS regression, we delete all observations in the 2.5% tails of the distribution of *coonsensus forecast* error (see Lim (2001)). Qualitatively similar findings emerge if we truncate 1% or 5% of the tails. Second, we estimate equation (2) with quantile (*i.e.*, median) and robust regressions, using the entire sample but without fixed–effects. In contrast to OLS regression which minimizes mean squared deviations, quantile (median) regression minimizes the mean absolute value of deviations. Hence, estimates from a quantile (median) regression are less sensitive to outliers. Finally we deal with the problem of extreme outliers via a robust regression technique. This consist of a series of iterations to detect and assign smaller weights to gross outliers (*e.g.,* Cook's distance).[7]

[7]A drawback with these two regressions using STATA is that, unlike OLS fixed–effects regression, STATA has

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

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

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 $\beta_{n,t}(F_\tau - \overline{F}_\tau)$.

no features that can allow us to run fixed–effects quantile or robust regression or control for non–independence of forecast erros within firms; an important issue in estimation of equation (2). The estimates using OLS (fixed–effects) regression are generally higher, but the qualitative nature of our results was unaltered.

Table 5: Economic Significance of Bias in Analysts' Forecasts

Panel A reports results of the regression for all forecasts and results for revised forecasts only are reported in Panel B. In fixed–effects regressions, robust standard errors (in parentheses) were computed using the Hubert-White method to account for possible correlation of forecast errors within firms: (*), (**) and (***) indicate significance at 90%, 95% and 99% levels.

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

	SFD	$lnorder$	$lncov$	$lnorder * lncov$	$lnorder * SFD$	$lncov * SFD$	$lnorder * lncov * SFD$
Panel A: Entire sample							
Fixed–effects	0.542	−0.074	0.140	0.009	0.286	0.083	−0.129
	(0.099)***	(0.016)***	(0.010)***	(0.006)	(0.111)***	(0.046)*	(0.041)***
Quantile	0.186	−0.010	0.015	−0.001	−0.061	0.103	−0.034
	(0.001)***	(0.0008)***	(0.0005)***	(0.0003)***	(0.001)***	(0.0005)***	(0.0003)***
Robust	0.049	−0.022	0.032	−0.001	−0.184	0.208	−0.011
	(0.003)***	(0.002)***	(0.001)***	(0.0007)*	(0.002)***	(0.001)***	(0.001)***
Panel B: Revised forecasts only							
Fixed–effects	−0.085	−0.039	0.240	−0.023	0.743	0.147	−0.242
	(0.139)	(0.039)	(0.023)***	(0.015)	(0.209)***	(0.095)	(0.075)***
Quantile	−0.033	−0.010	0.015	−0.002	0.034	0.022	−0.014
	(0.002)***	(0.002)***	(0.001)***	(0.0006)***	(0.002)***	(0.001)***	(0.0006)***
Robust	−0.036	−0.028	0.040	−0.002	0.005	0.053	−0.018
	(0.005)***	(0.004)***	(0.003)***	(0.001)*	(0.004)	(0.003)***	(0.001)***

Our findings are reported in Table 5. The results displayed in Panel A are for the entire sample. Panel B reports results from the regressions of percentage forecast errors on the difference between *revised* forecast and the consensus formed three days prior to the release of an analyst's revised forecast. The positive coefficient on `lncov` indicates that forecast errors are *larger* for firms followed by more analysts. The negative coefficient on `lnorder` indicates that that analysts who report later in the forecast-horizon make smaller errors for each percentage deviation from the consensus. Note that higher values of `order` imply higher values of `cover` (but not conversely), so that those reporting later will have more information at their disposal. Since they make lesser errors, it must be that they take this information into account.

It is easier to interpret the coefficient estimates on interaction terms within a particular context. The estimates imply that, for firms with higher analyst–coverage, analysts reporting later (rather than earlier) make significantly less errors for every percentage deviation from the consensus. Consequently the bias in later forecasts is significantly less if more analysts are following the firm. For example, using the estimates from the OLS fixed–effects regression in Panel A, 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 1% of the stock's price, then his forecast overshoots actual earnings on average by about 1.1%, *ex post*. However, for the same difference between the forecast of 20th reporting analyst and the consensus, he only overshoots earnings by 0.77% of the stock's price. Using our structural estimates, these imply an approximate bias of 2.58 times the difference between the analyst's (unobserved) true estimate of earnings and the consensus forecast if he were the 2nd to report (out of 20 analysts following the firm); while if he reports last (the 20th), then the implied bias is approximately 0.99 times.[8] The results also suggests that analysts who revise their old forecasts make use of relevant their information, issuing less biased forecasts.

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. If the forecast overshoots the consensus, it also tend to overshoot earnings, but the amount by which the analyst overshoots falls with the amount of information at his disposal. The large economic magnitudes of the forecast overshooting strongly suggest that it is not just due to analysts *myopically* ignoring information in the consensus.

[8]The implied bias in the estimates from the quantile regression is 0.62 (when analyst reports 2nd) and 0.01 (when he is the 20th analyst to issue a forecast), while robust regression estimates imply biases of 1.08 and 0.022, respectively.

4 Conclusion

This paper develops a simple frequency test for bias in the forecasts of security analysts. Detecting forecast bias 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 firms earnings, and information arrival.

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

Overshooting is greater for forecasts that are not revised, slightly greater for firms followed by fewer analysts and for analysts who report earlier in the forecast–horizon. The latter results demonstrates that this forecast bias is not driven by analysts' failure to incorporate useful information into their forecasts. We also find that forecast revisions that move away from the consensus are essentially unbiased, while those that move toward the consensus are particularly biased. That is, analysts are reluctant to revised their forecasts toward the consensus, and revisions in the direction of the consensus do not go far enough.

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) reporting biases his forecast by about 2.58 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 0.99 times the difference. Thus, the bias falls as uncertainty around earnings is resolved — with large analyst-coverage, those reporting late have access to more information — but the economic magnitude of the bias remains very large.

References

[1] Bernhardt, Dan, and Edward Kutsoati (2001), "Analyst's Compensation and Forecasts: theory, tests and evidence", *mimeo*, University of Illinois.

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

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

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

[5] Hong, Harrison, Jeffrey Kubik and Amit Solomon (2000), "Security Analysts' Career Concerns and Herding of Earnings Forecasts," *Rand Journal of Economics*, 31 (1), 121–144.

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

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

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

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

[10] Lim, Terence (2001), "Rationality and Analysts' Forecasts Bias," *Journal of Finance*, 56 (1), 369-385.

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

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

[13] Stickel, Scott (1990), "Predicting Individual Analyst Earnings Forecasts," *Journal of Accounting Research,* 28(2), 409-417.

[14] Stickel, Scott (1992) "Reputation and Performance among Security Analysts," *Journal of Finance,* 47, 1811-1836.

[15] Trueman, Brett (1990), "On the Incentives for Security Analysts to Revise Their earnings Forecasts," *Contemporary Accounting Research,* 7(1), 203-222.

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

[17] Welch, Ivo (2000), "Herding Among Security Analysts, " *Journal of Financial Economics*, 58 (3), 369-96.

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

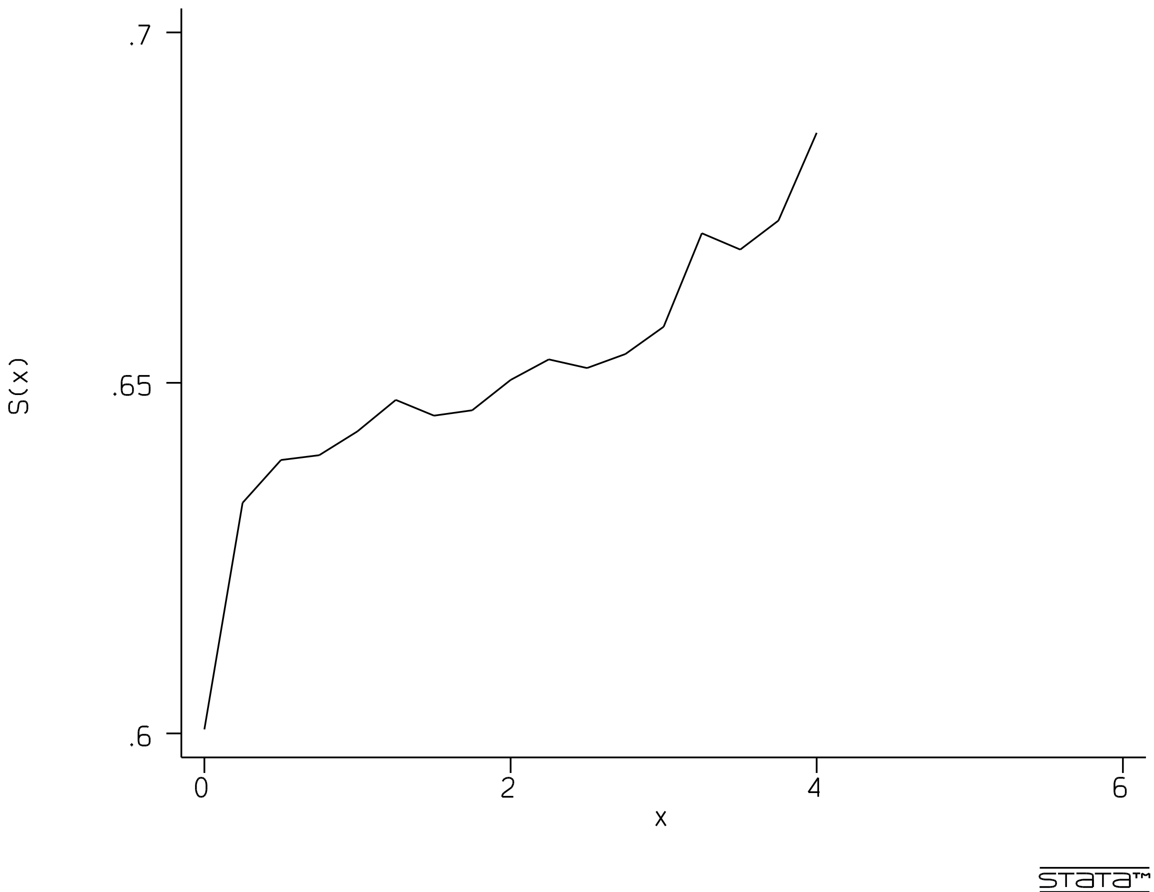

Figure 2: Average conditional probability that forecast overshoots earnings away from consensus, $\mathbf{S}(x)$, given that it exceeds or falls short of the consensus by at least $x\%$ of the stock's price. Higher values of x naturally pick more extreme forecasts. In order to eliminate the impact of such outliers, we delete all observations for which $x = 20\%$ in the plot of $\mathbf{S}(x)$ against x.