

Analysts' Weighting of Private and Public Information

Qi Chen

Fuqua School of Business, Duke University

Wei Jiang

Columbia Business School

Using both a linear regression method and a probability-based method, we find that on average, analysts place larger than efficient weights on (i.e., they overweight) their private information when they forecast corporate earnings. We also find that analysts overweight more when issuing forecasts more favorable than the consensus, and overweight less, and may even underweight, private information when issuing forecasts less favorable than the consensus. Further, the deviation from efficient weighting increases when the benefits from doing so are high or when the costs of doing so are low. These results suggest that analysts' incentives play a larger role in misweighting than their behavioral biases.

This article provides evidence on how analysts weight private and public information in forecasting corporate earnings. Specifically, we examine whether sell-side analysts misweight information, and if so, the potential explanations for such misweighting. We define misweighting as analysts weighting information differently from the efficient benchmark weights that minimize forecast errors (*FE*) (i.e., the optimal statistical weights in forming rational Bayesian expectation).

Evidence on weighting behavior is related to, but distinct from, evidence on the properties of realized forecasts such as bias and accuracy. Prior literature has primarily focused on forecast properties and documented evidence that suggests inefficient analyst forecast behaviors. However, forecast properties are a function of both information precision and analysts' forecast behavior, and, as such, cannot provide unambiguous

We benefit from helpful comments from Alon Brav, Charles Calomiris, Jennifer Francis, Itay Goldstein, David Hirshleifer, Pete Kyle, Terence Lim, Cade Massey, Maureen O'Hara (the editor), Per Olsson, Michael Roberts, Jeremy Stein, Kent Womack, Luigi Zingales, Eric Zitzewitz, and workshop participants at Carnegie Mellon, Columbia, Duke, London School of Economics, University of Chicago, University of Wisconsin at Madison, the 39th Federal Reserve Bank of Chicago Annual Conference on Banking Structure and Competition, and 2003 Western Finance Association Conference. We especially appreciate the constructive and detailed comments from two anonymous referees. We thank Tom Denton Jr, Xin Wang, Li Xu, and Peng Yin for research assistance. All errors are our own. Address correspondence to: Qi Chen, Fuqua School of Business, Duke University, Durham, NC 27708, USA, e-mail: qc2@duke.edu; or Wei Jiang, Columbia Business School, 3022 Broadway, New York, NY 10027, USA, e-mail: wj2006@columbia.edu.

inference about the latter.¹ Tests on weighting behavior provide more direct evidence of forecast inefficiency and can help discriminate competing explanations. Different explanations exist for forecast inefficiencies, with some attributing them to analysts' incentives and others to analysts' cognitive bias (such as overconfidence) in processing information. Since analysts are perceived to represent sophisticated investors and their forecasts are an important input to market expectations, understanding the source of forecast inefficiencies has important implications for studies on investor rationality and market efficiency.

Using a large sample of earnings forecasts over 1985–2001, we find that analysts overweight private information on average. We also document asymmetric weighting behavior in that analysts overweight private information when issuing forecasts more favorable than the consensus, and underweight private information when issuing forecasts less favorable than the consensus. We refer to this asymmetric weighting as optimistic weighting, that is, analysts tend to overweight the more favorable of the public and private signals.

We further explore the potential explanations for overweighting and optimistic weighting. For overweighting, we consider two competing hypotheses: one is based on analysts' incentives to signal ability or to generate trading commissions [Prendergast and Stole (1996), Ehrbeck and Waldmann (1996), and Zitzewitz (2001b)], and the other is based on analysts' overconfidence about their own ability due to attribution bias in learning [Griffin and Tversky (1992) and Gervais and Odean (2001)]. The incentive-based hypothesis presumes that analysts are informed about their own forecast ability; it predicts that analysts strategically overweight private information to exaggerate the news content of their private information in order to signal their forecast ability or to generate trading commissions. Because the costs and benefits of overweighting vary for analysts of different abilities, this hypothesis predicts that the degree of overweighting is related to analysts' ability. In contrast, the overconfidence hypothesis presumes that analysts are ignorant of their ability, and become overconfident (i.e., overestimate their own ability) after a run of good performance due to their attribution bias in learning. Thus, the overconfidence hypothesis predicts a positive relation between overweighting and analysts' prior forecast accuracy [i.e., track records (*TR*)], and no relation between overweighting and ability once controlling for *TR*.

To test the above hypotheses, we measure an analyst's ability as the frequency of his forecasts, which if added to the consensus would move the consensus in the direction of reported earnings. We prove that this measure of ability depends only on the relative precision of an analyst's

¹ An analyst may apply efficient Bayesian weights to his information but still have lower accuracy than another analyst who weights information inefficiently, if the latter analyst has more precise information.

private signals and not on his weighting strategy conditional on his ability.² This property is important as it allows us to identify the relation between analysts' forecast behavior (strategy) and ability, which serves as an identifying constraint to distinguish strategic overweighting from overconfidence. Using this ability measure (and its counterpart for track records), we find that the degree of overweighting is negatively related to analysts' ability, and is positively related to analysts' track records only after controlling for ability. Further, we find that analysts overweight private information more when they cover heavily traded stocks (when the potential trading commissions are high), and more so when they have good track records (when their perceived abilities are high). These results lend incremental support to the incentive-based hypothesis relative to the overconfidence hypothesis.

We also consider two hypotheses about optimistic weighting. The cognitive bias hypothesis argues that optimistic weighting arises because analysts hold optimistic priors about the firms they choose to cover and therefore, over (under) estimate the precision of their favorable (unfavorable) information [McNichols and O'Brien (1997)]. Alternatively, the incentive hypothesis argues that analysts intentionally weight information optimistically to placate the management of the firms they cover [Francis and Philbrick (1993) and Lim (2001)]. Since optimistic weighting hurts forecast accuracy and analysts' reputation, the incentive hypothesis predicts that the degree of optimistic weighting increases (decreases) with the potential benefits (costs) of doing so. In contrast, the cognitive bias hypothesis presumes that analysts are unaware of their misweighting, and therefore predicts no systematic relation between optimistic weighting and its costs and benefits.

Using investment banking affiliation (IB) as a proxy for the benefits of optimistic weighting, we find that analysts are more likely to weight information optimistically, if their employers will (in the future) underwrite for the firms they cover. We use analysts' experience and forecast timing as proxies for the costs of optimistic weighting. The argument is that mistakes early on in an analyst's career cause more reputation damage [Holmstrom (1999)], and intuitively, errors are more costly for forecasts issued closer to earnings release date because they are more likely to be remembered. We find that analysts misweight less early in their careers and when they forecast closer to earnings release dates. These findings are more consistent with the incentive hypothesis than with the cognitive bias hypothesis. Further, we find that past optimistic biases do not lead to less optimistic weighting in the future, suggesting either that analysts do not learn from their past misweighting or that the misweighting is mostly intentional.

² The standard measures of ability such as absolute or relative forecast accuracy do not satisfy this requirement.

Our results in general suggest that analysts' incentives play a larger role in misweighting than behavioral biases. Our results do not imply that all analysts are immune to behavioral biases (such as overconfidence or optimistic bias, or other forms of cognitive biases that we do not explicitly consider). The discriminating power of our tests is limited by the extent to which our empirical proxies capture the heterogeneity among analysts.

Our article makes the following contributions to the literature. First, it is the first to document optimistic weighting and its association with analysts' investment banking affiliation, experience, and forecast timing. Optimistic weighting, which we show is distinguishable from added bias, has several implications. Because it implies that analysts under (over) react to unfavorable (favorable) public information, optimistic weighting offers a potential reconciliation of the conflicting results on whether the consensus forecast (some weighted average of individual forecasts) under or overreacts to past public information.³ It also suggests that favorable news is more likely to be reflected in market expectations, supporting Hong and Stein's (1999) argument that bad news travels slowly. Further, as analysts deviate more from consensus when they have favorable news, higher forecast dispersion may be associated with exaggerated favorable news, causing unreasonably high current prices, consistent with Diether, Malloy, and Scherbina's (2001) finding that firms with high forecast dispersion earn lower future returns. Lastly, optimistic weighting in earnings forecasts complements Welch's (2000) finding that analysts herd in stock recommendations, especially when the consensus recommendation is overly optimistic.

Our article also contributes by offering tests to distinguish competing hypotheses for overweighting as well as by offering alternative explanation to prior findings on analysts' herding. While other studies find evidence of overweighting [e.g., Zitzewitz (2001a)], they do not further identify strategic overweighting from overconfidence. As discussed earlier, our ability measure increases the power of our tests to discriminate between these two possibilities. We find a positive relation between ability and distance from the consensus, and a negative relation between ability and overweighting. These relations question the appropriateness of using forecast distance to measure analyst herding as is commonly done in the literature.⁴

³ Specifically, DeBondt and Thaler (1990) find overreaction, Abarbanell and Bernard (1992) find underreaction, and Easterwood and Nutt (1999) find overreaction in firms experiencing increasing earnings but underreaction in firms experiencing decreasing earnings.

⁴ Prior studies find a positive relation between forecasts' distance from consensus and analysts' ranking, prior forecast accuracy, or experience [e.g., Hong, Kubik, and Solomon (2000) and Clement and Tse (2004)]. These results have been interpreted as indicating that low-ability, inexperienced analysts herd to the consensus. Our ability measure shows that low-ability analysts actually overweight private information (i.e., anti-herd to consensus) while high-ability analysts underweight. We also find that compared with inexperienced analysts, experienced analysts do not overweight more their favorable private information, but rather tend to underweight less their unfavorable private information.

A third contribution of our article is the generality of our model specification. We use both a linear regression method and a probability-based method. Together, they identify analysts' weighting strategy under general information structures (such as nonnormal distributions) and general (such as nonlinear) forecasting strategies. This article is the first to introduce the probability method which is more robust to outliers and measurement errors—two of the major empirical difficulties in analyzing analyst forecast data. Without the probability method, it is difficult to distinguish certain explanations from spurious correlations due to data irregularities. The consistency of our results from both the regression- and probability-based methods increases our confidence in the conclusions drawn.

The rest of the article is organized as follows. Section 1 formally defines weighting behavior and outlines test methods. Section 2 provides a description of the data and an overview of analysts' weighting behavior. Section 3 discusses alternative explanations for misweighting and tests competing hypotheses. Section 4 presents robustness checks and Section 5 concludes.

1. Definitions and Testing Methods

1.1 Definitions

Let z denote a firm's earnings that an analyst forecasts. Without loss of generality, we assume that z follows a diffuse zero-mean normal distribution. For notational parsimony, we represent the analyst's information set at the time of forecasting as the union of two sets: one for public information that is observable by both the analyst and investors and one for the analyst's private information. Let c be a sufficient statistics for all public information for z , expressed as

$$c = z + \varepsilon_c, \quad (1)$$

where $\varepsilon_c \sim N\left(0, \frac{1}{p_c}\right)$ and independent of z . We refer to c as the market consensus about z which is observed by investors and the analyst. Let y be the analyst's private signal about z with

$$y = z + \varepsilon_y, \quad (2)$$

where $\varepsilon_y \sim N\left(0, \frac{1}{p_y}\right)$ and independent of z and ε_c .⁵ The analyst's best conditional estimate of z given y and c is formed by Bayes' rule as follows:

$$E[z | y, c] = hy + (1 - h)c, \quad (3)$$

⁵ Assuming mutual independence between ε_c and ε_y is for tractability and is not crucial for our analysis. As long as y and c are not perfectly correlated, they can be orthogonalized in this manner.

where $h \equiv \frac{p_y}{p_c + p_y} \in [0, 1]$ is the precision of analysts' private signals relative to the consensus. Because (3) is the mean-squared error (MSE) minimizing estimate of z , we hereafter refer to h as the efficient or Bayesian weight. Practically, each y signal may have different precision level. On average, analysts of higher ability have signals with higher precision. For simplicity, we index analyst i 's ability by h_i .

In making forecasts, analysts may not apply the efficient weight h . Let

$$f = ky + (1 - k)c \quad (4)$$

be an analyst's forecasting strategy, where k is the actual weight the analyst places on his private signal. Although we assume a linear forecasting strategy for simplicity, our empirical tests do not rely on this assumption. (Section 4 discusses how our tests can be applied under more general assumptions about information structures and forecast strategies.) Throughout the article, we say an analyst *misweights* information if $k \neq h$. We define over or underweighting more precisely as follows:

Definition 1. Let h_i and k_i be analyst i 's efficient weight and actual weight on private information, given by (3) and (4), respectively. We say that analyst i *overweights* (*underweights*) private information if $k_i > h_i$ ($k_i < h_i$). Further, analyst i *overweights more*, or *underweights less*, private information than analyst j if $\frac{k_i}{h_i} > \frac{k_j}{h_j}$.

Definition 1 defines an analyst's degree of over or underweighting as his actual weights scaled by the efficient weights, similar to Gervais and Odean's (2001) definition of over or underconfidence.⁶ Therefore, while an analyst with higher forecast ability may place higher absolute weight on private information than an analyst of lower ability (i.e., $k_i > k_j$), he does not necessarily overweight more than the lower-ability analyst (i.e., $\frac{k_i}{h_i}$ may be higher or lower than $\frac{k_j}{h_j}$).

1.2 Testing methods

Researchers observe an analyst's forecast (f) and some measure of the consensus (c), but not his private signal (y), his precision (h), or his weighting on y (k). The relation between the analyst's forecast error and his forecast's deviation from the consensus, however, reveals information about how the analyst weights his signals. We use two methods to extract this information.

The first method is *regression based*. It builds on the idea that FE should not be predictable by available information, if the analyst

⁶ An alternative definition is in terms of the difference between h_i and k_i . Our tests, discussed next, cannot measure the unscaled difference between h_i and k_i . This is because both a forecast's distance from consensus, and the relation between its deviation from the consensus and its error, are related to the ratio of h_i and k_i , not the difference between h_i and k_i .

efficiently weights information. To see this, note that given the analyst's forecasting strategy (4), we can express the analyst's expected forecast error (FE) as a function of his forecast's deviation from the consensus, that is,

$$E((f - z)|y, c) \equiv E(FE|y, c) = \frac{(k - h)}{k} (f - c) = \beta_0 \cdot Dev,$$

where $\beta_0 = 1 - \frac{h}{k}$, $FE = f - z$, and $Dev = f - c$. This equation implies that if the analyst uses efficient weights (i.e., $k = h$), his forecast's deviation from the consensus should have no predictive power for his FE . In the regression

$$FE = \alpha + \beta_0 \cdot Dev + \varepsilon, \quad (5)$$

the coefficient estimate $\hat{\beta}_0$ converges in probability to $(1 - \frac{h}{k})$. Therefore, a positive (negative) $\hat{\beta}_0$ suggests overweighting (underweighting).⁷

The second method is *probability based*. It builds on the idea that if an analyst uses efficient weights, it is equally likely that his forecasts "overshoot" true earnings (i.e., the forecasts err on the same side as their deviation from the consensus forecast) or "undershoot" true earnings. More precisely, define

$$\pi = \Pr(\text{sign}(FE) = \text{sign}(Dev)). \quad (6)$$

Then the expected value of π under the null hypothesis of efficient weighting is 0.5, and higher values of π suggest more overweighting of private information. The empirical analogue to (6) is $\hat{\pi} = \frac{1}{J} \sum_{j=1}^J 1(\text{sign}(FE_j) = \text{sign}(Dev_j))$, where j is the subscript for the forecast and J is the number of forecasts in the sample. $1(\cdot)$ is an indicator function that equals 1 if the argument is true and 0 otherwise.

The following Proposition formalizes the above discussion of the two methods [proof of part (ii) is provided in Appendix A].

Proposition 1. *Suppose y and f are generated according to (2) and (4). Then: (i) β_0 in (5) is zero if and only if the analyst uses efficient weights. Further, $\beta_0 > 0$ if $k > h$ and $\beta_0 < 0$ if $k < h$. (ii) π in (6) is $\frac{1}{2}$ if and only if the analyst uses efficient weights. Further, $\pi > \frac{1}{2}$ if $k > h$; and $\pi < \frac{1}{2}$ if $k < h$.*

When $\frac{h}{k}$ varies across analysts, β_0 and π could capture different aspects of misweighting behavior. Specifically, β_0 measures the average *magnitude* of under or overweighting, while π measures the median

⁷ The regression test is similar to the test in Zitzewitz (2001a) which regresses the forecast error of the consensus on Dev . He interprets a coefficient estimate on Dev greater (less) than one as analysts exaggerating (compressing) private news.

tendency to under or overweight by the sample analysts. Moreover, the π -statistic is less sensitive to outliers because each pair of $\{FE, Dev\}$ contributes equally to the statistic regardless of the magnitudes of FE and Dev .

Both the regression and the probability methods can be adapted to examine cross-sectional variation in the magnitude of misweighting. For the regression method, we estimate the following equation:

$$FE = \alpha + \sum_{m=1}^M \gamma_m \cdot X_m \cdot Dev + \varepsilon, \quad (7)$$

where $\mathbf{X}$ is a vector of M factors (including a constant) affecting analysts' weighting and $\boldsymbol{\gamma}$ is a conformable vector of coefficients. A positive (negative) $\hat{\gamma}_m$ indicates that a higher value of X_m is associated with more overweighting (underweighting). Section 3 discusses possible $\mathbf{X}$ covariates.

The probability test can be performed on subsamples sorted by different X_m variables. It can also be extended using the following specification:

$$\mathbf{1}(\text{sign}(FE) = \text{sign}(Dev)) = \alpha + \sum_{m=1}^M \gamma_m \cdot X_m + \varepsilon, \quad (8)$$

where $\mathbf{1}(\cdot)$ is the indicator function. Coefficient $\hat{\gamma}_m$ can be estimated via maximum likelihood methods such as probit.

We note that both methods better capture analysts' weighting behavior than the measured distance between analysts' forecasts and the consensus. Previous studies used the forecast distance from consensus as a proxy for forecast "boldness" or herding tendency. Such an approach does not separate forecast strategy from forecast ability because the forecast distance from consensus depends on both. To see this point, we derive the mean-squared deviation from consensus as follows:

$$E(f - c)^2 = k^2 E(y - c)^2 = h^2 \left(\frac{1}{p_y} + \frac{1}{p_c} \right) \left(\frac{k}{h} \right)^2 = \frac{h}{p_c} \left(\frac{k}{h} \right)^2. \quad (9)$$

Equation (9) shows that the distance from consensus is a function of both the relative precision of an analyst's signals (captured by $\frac{h}{p_c}$) and his forecast strategy [captured by $\left(\frac{k}{h}\right)^2$]. Thus, while analyst i may intentionally be bold by overweighting private information (i.e., $k_i > h_i$), his forecast may still be closer to the consensus, if his true ability is lower than another analyst j (i.e., $h_j > h_i$) who herds by underweighting (i.e., $k_j < h_j$). Findings in Section 3 show that this is indeed the case.

2. Data Sample and Overview

2.1 Sample description

We obtain analysts' quarterly earnings forecasts from the Zacks Investment Research database and firms' earnings and stock price data from COMPUSTAT and CRSP. Information about firms' underwriters is obtained from Thompson Financial's SDC database, and information about All-American Research Teams is hand-collected from the *Institutional Investor* magazine. We begin by identifying analysts whose first forecasts (for any firm and any type of forecast) recorded by Zacks were made on or after January 1, 1985. This constraint ensures that the experience and track records of our sample analysts can be measured accurately. Following the standard practice in the literature, we eliminate forecasts for firms whose average share prices are less than \$5 and average market capitalizations are less than \$100 million (both in 2001 CPI-deflated dollars) to mitigate the influence of extreme outliers. Our sample contains 1,367,599 forecasts, representing 3195 firms, 5306 analysts, and 51,200 analyst-firm pairs.⁸

FE is the difference between the forecasted and the realized earnings, and the forecast's deviation from consensus (Dev) is the difference between the forecast and the consensus. An ideal measure for the consensus c is the best predictor of earnings z using all public information at the time of the forecast. For brevity and following the standard practice in the literature, our main analyses measure c as a weighted average of all prevailing forecasts. In sensitivity checks, we repeat all analyses measuring c as the predicted earnings from time-adapted public information and find qualitatively similar results. To ensure that the consensus is observed by an analyst when he forecasts, the consensus includes all forecasts for the same firm-quarter issued at least one day before the current forecast. We use an inverse-weighting scheme that assigns higher weights to more recent forecasts as they contain more updated information. Specifically, if there are $n = 1, \dots, N$ prevailing forecasts, F_n , that are issued d_n days before the current forecast date, with $d_N > d_{N-1} > \dots > d_1$, the inverse weighting assigns weight $w_n = \frac{1/d_n}{\sum_{n=1}^N 1/d_n}$. We obtain qualitatively similar results (not reported) when we use two other weighting schemes: equal weighting (which assigns weight $w_n = \frac{1}{N}$ to each F_n) and linear forecast order weighting (which assigns $w_n = \frac{N-n+1}{\sum_{n=1}^N (N-n+1)}$). Table 1 lists the definitions and summary statistics for the main variables used in this article.

⁸ The number of observations in individual regressions varies because of additional information requirements. The final forecast in our sample was issued in March 2001, before Regulation FD took effect.

Table 1
Variable definitions and summary statistics

	Number of observations	Mean	Standard deviation	1st percentile	25th percentile	Median	75th percentile	99th percentile
<i>Ability(Dir)</i>	52,569	0.247	0.489	-1.00	0	0.267	0.543	1.00
<i>Ability(Beat)</i>	52,569	0.141	0.491	-1.00	-0.109	0.143	0.429	1.00
<i>Days</i>	1,346,238	1.910	1.5477	0.03	0.71	1.52	2.81	6.38
<i>Dev</i>	1,367,599	-0.019	0.129	-0.472	-0.046	-0.007	0.019	0.333
<i>Exp</i>	1,367,599	0.277	0.381	0	0.03	0.13	0.37	1.78
<i>FE</i>	1,367,599	0.162	0.638	-0.850	-0.030	0.030	0.200	2.790
<i>IB</i>	1,367,599	0.041	0.199	0	0	0	0	1
<i>TR(Dir)</i>	1,192,201	0.278	0.507	-1.00	0	0.329	0.632	1.00
<i>TR(Beat)</i>	1,192,201	0.068	0.485	-1.00	-0.200	0.059	0.333	1.00
<i>Vol</i>	1,367,599	0.0037	0.1595	0.0001	0.0017	0.0056	0.0191	0.5817

Sample consists of forecasts by analysts whose first appearances in Zacks occurred on or after January 1, 1985, for firms whose average stock prices are at least \$5 and average market capitalizations at least \$100 millions (both in 2001 CPI-deflated values) in the sample period. *Ability* is calculated using all forecasts in an analyst-firm pair. The other variables are calculated for each forecast. *Ability(Dir)*, negative of the average of sign indicators for all forecasts in an analyst-firm pair. The sign indicator equals 1, 0, or -1 if the product of the forecast's error and the error of its corresponding *Consensus* has a positive, zero, or negative sign. *Consensus* is calculated for each forecast as a weighted average of all outstanding forecasts with weights inversely proportional to days to this forecast date; *Ability(Beat)*, average of sign indicators for all forecasts in an analyst-firm pair. The sign indicator equals 1, 0, or -1 if a forecast is more accurate than, equal to, or less accurate than its corresponding consensus; *Days*, number of days, in 100s, between the forecast date and the earnings release date; *Dev*, difference between the current forecast and its corresponding consensus forecast; *Exp*, number of past tested forecasts, in 100s, by the analyst up to the date of the forecast; *FE*, forecast error, the difference between the forecast and the realized earnings as reported on COMPUSTAT (item number 19, EPS excluding extraordinary items); *IB*, dummy variable equal to one if the analyst is affiliated with a bank that is a lead or co-underwriter for the security issuance of the firm during the five-year period centered on the year of the forecast date; *TR(Dir)*, constructed the same way as *Ability(Dir)*, using only information up to the time of the forecast; *TR(Beat)*, constructed the same way as *Ability(Beat)*, using only information up to the time of the forecast; *Vol*, average daily trading volume (in \$ billions) during the 50 trading days preceding the current forecast date.

2.2 Overview of analysts' weighting behavior

Figure 1 plots a nonparametric kernel regression of $FE = f(Dev)$ without imposing a structural model. An overall tendency of $f(\cdot)$ increasing with *Dev* suggests overweighting and $f(\cdot)$ decreasing with *Dev* suggests underweighting. Figure 1(a) uses *FE* and *Dev* in dollar units, and Figure 1(b) scales both variables by the stock price five days before the forecast. Both figures show a "V"-shaped relation between *FE* and *Dev* in that *FE* increases with *Dev* when *Dev* is positive; and decreases with *Dev* when *Dev* is negative. The Ellison and Ellison (2000) specification test does not reject linearity for each segment at less than the 5% significance level, but rejects linearity for the whole sample (i.e., no kink) at less than the 1% level.⁹ Juxtaposed in Figure 1 with the kernel regression plots are the estimated sample densities of *Dev*, providing an overview of the data

⁹ The Ellison and Ellison (2000) test statistics are based on quadratic forms in the null (linear) model's residuals. The idea is that quadratic forms of residuals (with some kernel weighting matrix) can detect a spatial correlation in the residuals due to nonlinearity.




Figure 1
Kernel regressions of forecast error versus deviation from consensus

The “V”-shaped solid lines are the kernel regressions of forecast errors (*FE*) against forecast deviation from the consensus (*Dev*). The dotted lines around the V-shape represent the 95% confidence intervals for the predicted *FE* from the nonparametric regression. *FE* and *Dev* are in dollar units in (a) and are normalized as percentages of share prices five days before the forecast dates in (b). The solid bell-curves are the kernel-based estimated density functions of *Dev*. Bartlett kernel is used, and window widths are adapted to data density.

concentration. Together they show that the increasing (decreasing) relation between *FE* and *Dev* when *Dev* is positive (negative) holds in ranges of moderate *Dev* values as well as in more extreme ranges.

The average magnitude of misweighting by our sample analysts is estimated by the coefficient β_0 from the following regression:

$$FE_{i,j,t} = \alpha_j + \beta_0 Dev_{i,j,t} + \varepsilon_{i,j,t}, \quad (10)$$

where the subscripts i , j , and t indicate that the variable is related to analyst i 's forecast for firm j made at time t , and α_j is a firm-specific intercept. Since most variables are specific to each forecast, we omit all i, j, t subscripts when there is no confusion. Given that β_0 in (5) is a ratio of weights and is independent of the scales in FE and Dev , we estimate (10) using FE and Dev in dollar units. This is preferred to using scaled versions of FE and Dev , because scaling introduces potential bias if analysts' misweighting is correlated with the scaling variable. Sensitivity checks show that scaling FE and Dev by stock price produces similar, but noisier, results.

After factoring out the firm-specific component, the error disturbances in (10) could be correlated for two reasons. First, common shocks to firm earnings may result in forecasts for different firms on the same quarter to systematically err on the same side. Second, same analysts' FE may be serially correlated. To accommodate such error correlations, we adjust standard errors for arbitrary within-cluster error correlations as well as heteroskedasticity in all estimations. Each cluster includes all forecasts for firms in the same two-digit SIC industry and for the same calendar quarter. The effective sample size for estimating standard errors is on the order of the number of clusters (which is about 0.33% of the total number of observations in our sample), and inference obtained from such large-cluster adjusted standard errors is conservative if there is no cluster effect among some subgroups of observations in the same cluster [Wooldridge (2003)].¹⁰ For the probability method, we calculate the t -statistics (for whether the statistic $\hat{\pi}$ differs from the null value of 0.5) based on block-resampling bootstrap standard errors that adjust for the same clustering.¹¹

Estimating (10) on the whole sample yields a coefficient $\hat{\beta} = 0.19$ (t -statistic = 9.70), which indicates that analysts overweight on average. However, separate regressions on subsamples partitioned by the sign of Dev show strong evidence of asymmetric misweighting. For the positive Dev subsample, $\hat{\beta}$ is 0.75 (t -statistic = 18.78); and for the negative Dev subsample, $\hat{\beta}$ is -0.14 (t -statistic = -4.78). These estimates imply that analysts overweight private information by 75% when their forecasts are more favorable than the consensus,

¹⁰ We believe clustering at the industry-quarter level sufficiently addresses the correlations among error disturbances. We do not find significant serial correlation for forecasts made by the same analysts for earnings for different quarters (net of firm fixed effects), neither do we find significant correlation of residual forecast errors of firms in different two-digit SIC industries.

¹¹ Block-resampling bootstrap (where all observations belong to the same cluster are automatically resampled if one observation in that group is resampled) adjusts for within-cluster correlations in estimating the standard errors of the frequency statistic π . For a reference on block-resampling bootstrap, see Davison and Hinkley (1997), chapter 8.2.3.

and underweight by 14% when their forecasts are less favorable than the consensus. We obtain qualitatively similar results for the subsample of forecasts issued during the stock market boom period from 1993 to 1999. These results are consistent with the V-shaped graphs shown in Figure 1.

The probability method measures analysts' tendency (rather than magnitude) to misweight information. For the whole sample, the $\hat{\pi}$ -statistic is 0.506. Though statistically significantly different from the null value of 0.5 at the 5% level, the difference is negligible in economic terms. However, the differences are substantial for subsamples partitioned on the sign of *Dev*.¹² The $\hat{\pi}$ -statistic for the positive *Dev* subsample is 0.674 (*t*-statistic = 26.8) and 0.412 (*t*-statistic = -25.7) for the negative *Dev* subsample, indicating that the probability of overweighting is 67.4% for favorable forecasts and 41.2% for unfavorable forecasts (a probability less than 50% indicates underweighting).

In summary, *aggregate* results indicate some evidence of analysts overweighting. Results on subsamples partitioned on the sign of *Dev* show that analysts' forecasts deviate too much from the consensus (i.e., overweight private information) when they exceed the consensus, but deviate too little (underweight private information) when they are below the consensus. Assuming no perversion in forecasting [i.e., assuming analysts will not forecast lower (higher) than the consensus when their private signals indicate more (less) favorable information than the consensus], these results suggest that analysts tend to overweight the more favorable information. We term this pattern "optimistic weighting" to distinguish it from the pattern where analysts place too much weight on their private information irrespective of whether the private information is favorable or unfavorable ("overweighting").

3. Why Do Analysts Misweight Information?

Both overweighting and optimistic weighting indicate inefficient forecast behaviors. This section first discusses the hypotheses that predict these two types of misweighting, and then presents empirical evidence to distinguish the hypotheses.

3.1 Hypotheses and empirical design

3.1.1 Hypotheses for overweighting and optimistic weighting. We consider two hypotheses for overweighting. The first is the overconfidence hypothesis which argues that analysts do not know their true ability and

¹² Observations with *Dev* = 0 are treated as underweighting in calculating the statistics for the whole sample, and are randomly split into subsamples of positive and negative *Dev* in proportion to their sample sizes. Due to their infrequency (2.4% of the whole sample), our results are almost identical if observations with *Dev* = 0 are excluded (potentially biasing the results toward overweighting).

overweight private information because they overestimate (i.e., are overconfident about) the precision of their private information. Overconfidence can arise because of analysts' attribution bias in learning about their ability from past performance. That is, they are likely to attribute past successes to their own ability, but blame past failures on external forces [Gervais and Odean (2001)].¹³ Attribution bias implies a positive correlation between overweighting and analysts' past forecast accuracy (i.e., TR). Further, the overconfidence hypothesis presumes that analysts are uninformed about their ability, and thus predicts no relation between overweighting and analysts' ability after controlling for TR.¹⁴

The second hypothesis argues that overweighting is due to analysts' strategic incentives to generate large forecast deviations from the consensus. One such incentive is to obtain high perceived ability, as suggested by theoretical models by Prendergast and Stole (1996), Ehrbeck and Waldmann (1996), and Avery and Chevalier (1999). The general intuition is the following. For a given weighting strategy, forecasts by higher-ability analysts will, on average, deviate more from the consensus [see Equation (9)]. Thus, the market rationally uses information in forecast deviations to form beliefs about an analyst's ability, particularly if direct performance measures (i.e., FE) are noisy and take time to accumulate. Consequently, analysts may overweight private information in order to generate large forecast deviations, a pattern typical of high-ability analysts.

Generating trading commissions provides another incentive to produce large forecast deviations. Press articles report that analysts' compensations, including those at independent research firms, are often tied to the trading commissions they generate for their firms' broker/dealer arms.¹⁵ Such incentive structures potentially encourage analysts to make bold statements or to exaggerate unsubstantiated news (both implying overweighting). Lys and Sohn (1990) and Irvine (2004) find that analyst forecasts have larger stock price and volume impacts when their forecasts deviate more from the consensus. According to the incentive hypothesis, an analyst is likely to overweight more when he covers heavily traded stocks and when he has high perceived ability. This is because more trading volume is generated by a given amount of news in his forecasts [Cooper, Day, and Lewis (2001)]. Since investors form perceptions about

¹³ Theoretically "under-confidence" is also a possibility after the agent experiences a run of bad performance. However, the literature on learning has shown a prevalent attribution bias in learning in that people are prone to attributing success to their own dispositions and failure to external forces [e.g., Hastorf, Schneider, and Polefka (1970) and Griffin and Tversky (1992)].

¹⁴ To illustrate this point, suppose an analyst is uninformed about his ability (h). He weights his private information based on his self-assessment $\hat{h} = E(h|Z) + \delta$, where Z represents information available to the analyst at the time, including his own track records. δ is a random error which has positive mean if the analyst is overconfident. It is straightforward to show that the magnitude of misweighting, $E(\hat{h} - h|z) = \delta$, is uncorrelated with h controlling for Z .

¹⁵ See "The Myth of Independence," *Business Week*, September 8, 2003.

an analyst's ability from his track records [Chen, Francis, and Jiang (2005)], the incentive hypothesis also predicts a positive relation between overweighting and track records, similar to the prediction from the overconfidence hypothesis.

The overconfidence and incentive hypotheses differ in their predictions about the relation between analysts' true ability and their weighting behaviors. While overconfidence predicts no systematic relation once track records are controlled for, the incentive hypothesis predicts otherwise. Specifically, overweighting increases the probability of having large forecast error relative to peers (i.e., the consensus), which puts analysts' future reputation at risk. Findings in prior literature indicate that higher-ability analysts are likely to stay in the profession longer.¹⁶ Thus, the incentive hypothesis implies that higher-ability analysts overweight less because their continuation value in the profession is higher and reputation is more valuable for them. We summarize the above discussion in the following hypothesis:

Hypothesis 1. *If analysts overweight due to overconfidence, then the degree of overweighting is positively related to their past forecast accuracy and unrelated to their forecast ability once track records are controlled for. If analysts overweight strategically (to signal ability or to generate trading commissions), then the degree of overweighting is positively related to their past forecast accuracy, conditional on true ability; and negatively related to true ability, conditional on their past forecast accuracy. Further, analysts are more likely to overweight when the potential for generating trading volumes is higher, and more so when their track records are better.*

We next consider two explanations for optimistic weighting. The first is the *optimistic bias* hypothesis drawn from psychology research. It suggests that individuals tend to rationalize their behaviors, and are often too optimistic in that they assess a higher probability (than the truth) to outcomes favorable to themselves [e.g., Festinger (1957) and Kahneman, Slovic, and Tversky (1982)]. Analysts have been shown to initiate (terminate) the coverage of a firm whose prospects they view as favorable (unfavorable) [McNichols and O'Brien (1997) and Rajan and Servaes (1997)]. To rationalize his prior favorable views about the firms he chooses to cover, an analyst may overestimate the precision of (and thus overweight) favorable signals about such firms; similarly, he may underestimate the precision of (and underweight) unfavorable signals.

Alternatively, the *incentive* hypothesis conjectures that analysts intentionally over (under) weight favorable (unfavorable) information to placate the management of the firms they cover. Two oft-discussed benefits from catering to client management are to obtain investment banking

¹⁶ See e.g., Stickel (1992), Mikhail, Walther, and Willis (1999), Hong, Kubik, and Solomon (2000), Zitzewitz (2001b), and Hong and Kubik (2003).

business [Dugar and Nathan (1995), Lin and McNichols (1998), and Michaely and Womack (1999)] and to gain better access to the firm's information [Francis and Philbrick (1993) and Lim (2001)].¹⁷

The two hypotheses have different predictions about whether optimistic weighting is related to the costs and benefits of doing so. The optimistic bias hypothesis presumes that analysts are unaware of their bias in interpreting information, and hence are unaware of their misweighting. This hypothesis, therefore, predicts no systematic relation between optimistic weighting and the benefits or costs of doing so. In contrast, the incentive hypothesis predicts that analysts engage in more optimistic weighting when the benefits (costs) of optimistic weighting are relatively high (low).

Further, under the optimistic bias hypothesis, if analysts learn to correct their bias, optimistic weighting should diminish as they observe more of their prior biases. In contrast, the incentive hypothesis predicts that misweighting increases with analysts' experience, as studies on reputation concerns [e.g., Holmstrom (1999)] suggest that the costs of misweighting are higher early in an analyst's career for two reasons. First, analysts would benefit more from good starting track records, since their expected tenure in the profession is longer. Second, early in their career, the market's assessment of analysts' true ability is diffuse because of the lack of track records; consequently, large early forecast errors inflict the most damage to their reputation. We summarize the above discussion in the following hypothesis:

Hypothesis 2. *If analysts optimistically weight information for incentive reasons, then the magnitude of optimistic weighting should be positively (negatively) related to the benefits (costs) of optimistic weighting. If optimistic weighting is due to optimistic bias, then there should be no such relations. Further, if optimistic weighting is unintentional and analysts learn over time, then optimistic weighting should diminish as analysts observe more of their forecast biases.*

3.1.2 Empirical specification. To test the hypotheses, we estimate the vector of $\hat{\gamma}$ coefficients in (7) and (8). The independent variables (the X covariates) are motivated from Hypotheses 1 and 2. Specifically, to test Hypothesis 1, we include variables measuring analysts' true ability (*Ability*) and track records (*TR*), as well as a variable proxying the trading commission incentive. We postpone the discussion of the construction of *Ability* and *TR* to the next section. We proxy for the trading commission incentive

¹⁷ On March 29, 2001, Laura Martin, a CSFB research official, was quoted saying "I will NOT lower numbers on AOL, even though they can't make them. . . . I don't think you do investors a favor if you so irritate a company that they stop talking to you." (The *Washington Post*, October 22, 2002.)

with the trading volume (*Vol*), measured in billions of dollars, of the stock being covered during the 50 trading days prior to the forecast. The idea is that the potential for trading commissions is higher for stocks with high trading volume.

For Hypothesis 2, we include independent variables proxying for the costs and benefits of misweighting. We use *IB* to proxy for the benefit of optimistic weighting, where *IB* is a dummy variable capturing the underwriting relation between an analyst's employer and the firm he covers during the five-year window centered on the year of the current forecast. Analysts could exaggerate favorable news and compress unfavorable news either to help their employers win future investment banking business or to help promote firms they underwrote in the past.

We use two measures as inverse proxies for the costs of FE: the number of days (in 100s) between the forecast date and the actual earnings release date (*Days*), and analyst's experience (*Exp*), measured as the number of realized forecasts (in 100s) that the analyst has made before the current forecast date. As discussed earlier, theory suggests that forecast errors are less costly for a more experienced analyst because the marginal forecast errors has a smaller impact on his perceived ability. We further conjecture that forecast errors are more costly for forecasts issued closer to earnings release dates because they are more likely to be remembered and penalized by investors, and because the benefit of misleading the market within a short period of time is small.

We estimate both (7) and (8) separately on subsamples partitioned by the type of news (sign of *Dev*). Throughout, in the rest of the article, we use $\hat{\gamma}_x^+(\hat{\gamma}_x^-)$ to denote the coefficient estimate for a x covariate on the subsample of positive (negative) *Dev* and $\hat{\gamma}_x$ for the coefficients on x for both subsamples. Given the overall optimistic weighting documented earlier, a positive $\hat{\gamma}_x^+$ or a negative $\hat{\gamma}_x^-$ implies greater optimistic weighting when x is higher.

Hypothesis 1 implies that regardless of the type of news, a finding of $\hat{\gamma}_{TR} > 0$ and $\hat{\gamma}_{Ability} = \hat{\gamma}_{Vol} = 0$ is consistent with the overconfidence hypothesis, while a finding of $\hat{\gamma}_{TR} > 0$, $\hat{\gamma}_{Ability} \neq 0$, and $\hat{\gamma}_{Vol} > 0$ provides support for the incentive hypothesis. According to Hypothesis 2, a finding of $\hat{\gamma}_{IB}^+ > 0$ and $\hat{\gamma}_{IB}^- < 0$ indicates that affiliated analysts misweight more than unaffiliated analysts. Further, the incentive hypothesis predicts (i) $\hat{\gamma}_{Days}^+ > 0$ and $\hat{\gamma}_{Days}^- < 0$, and (ii) $\hat{\gamma}_{Exp}^+ > 0$ and $\hat{\gamma}_{Exp}^- < 0$. The optimistic bias hypothesis, on the other hand, predicts (i) $\hat{\gamma}_{Days}^+ = \hat{\gamma}_{Days}^- = 0$ and (ii) $\hat{\gamma}_{Exp}^+ \leq 0$ and $\hat{\gamma}_{Exp}^- \geq 0$ if analysts learn about, and correct, their misweighting overtime.

3.1.3 Measuring analyst ability and track records. Analysts' true ability (*Ability*) and *TR* are two key variables in testing our hypotheses. Given the available information, our best effort to capture an analyst's ability is

to use all his forecast observations. Since Hypothesis 1 predicts that analysts' ability affects their weighting strategy, it is important that the *Ability* measure be a function only of the precision of the analyst's private signals [i.e., h in (3)] but not of his weighting strategy [i.e., k in (4) for a given h]. The following measure meets these conditions and is used as our main ability measure:

$$Ability(Dir)_{ij} = -\frac{1}{N_{ij}} \sum_{t=1}^{N_{ij}} \text{sign}(FE_Con_{ij,t} \cdot Dev_{ij,t}), \quad (11)$$

where $\text{sign}(\cdot)$ is the sign function and $FE_Con_{ij,t}$ is the forecast error of the consensus. N_{ij} is the total number of tested forecasts (tested forecasts are those whose errors are observed in the sample period) by analyst i for firm j .¹⁸ Intuitively, $Ability(Dir)$ measures the frequency that an analyst's forecasts move the new consensus (after incorporating his forecasts) in the direction of reported earnings. $Ability(Dir)$, bounded between -1 and 1 , is the empirical analogue to (using the notations in Section 1)

$$2 \cdot \Pr[\text{sign}(f - c) = \text{sign}(z - c)] - 1. \quad (12)$$

The following proposition shows that this measure is independent of analysts' forecast strategy k given their ability (the proof is in Appendix A).

Proposition 2. *Suppose y and f are generated according to (2) and (4). Assume there is no perversion in analysts' forecasts (i.e., an analyst would not issue a forecast higher (lower) than the consensus if his private signal is lower (higher) than the consensus). Then (12) is monotonically increasing in h , and is independent of k for any given h .*

As a robustness check, we also construct a performance-based ability measure calculated as the frequency of an analyst's forecasts being more accurate than the consensus:

$$Ability(Beat)_{ij} = \frac{1}{N_{ij}} \sum_{t=1}^{N_{ij}} \text{sign}(|FE_Con_{ij,t}| - |FE_{ij,t}|). \quad (13)$$

Both $Ability(Dir)$ and $Ability(Beat)$ are scale-free, and both control for the accuracy of public information (consensus). Such properties are desirable because they allow for cross-sectional comparisons, and adjust both for the differences in the predictability of firms' earnings and for the fact that forecasts become more precise as the earnings release date approaches. $Ability(Beat)$ is a less reliable ability measure because it

¹⁸ In our analysis, $Ability_{ij}$ and $TR_{ij,t}$ are both analyst-firm specific rather than analyst specific to account for Park and Stice's (2000) finding that analysts' ability is firm specific.

measures an analyst's ex post forecast accuracy, which could be affected by both ability and weighting strategy. The correlation coefficient between $Ability(Dir)$ and $Ability(Beat)$ is 0.88, indicating that analysts with more precise signals do achieve better forecast accuracy.

We measure analyst i 's TR for firm j at forecast date t , $TR(Dir)_{i,j,t}$ or $TR(Beat)_{i,j,t}$, analogously to the corresponding ability measures but using only information available at or before t . That is,

$$TR(Dir)_{i,j,t} = -\frac{1}{N_{i,j,t}} \sum_{k=1}^{N_{i,j,t}} \text{sign}(FE_Con_{i,j,k} \cdot Dev_{i,j,k}),$$

where $N_{i,j,t}$ is the number of realized forecasts made by analyst i on firm j up to time t . $TR_{i,j,t}$ is also a proxy for investors' perception of an analyst's ability and the analyst's updated estimate of his own ability if he was not informed of it. $Ability_{i,j}$ and $TR_{i,j,t}$ are correlated by construction, and the correlation is high when an analyst has a long track record. To delineate the separate effects of ability and track record in regressions where both measures appear, we orthogonalize them by replacing TR with the residual from a linear regression of TR on $Ability$. This residual represents an analyst's prior performance that is not explained by his true ability.

We use $Ability(Dir)$ to verify an important premise underlying the incentive hypothesis: forecasts by higher-ability analysts, on average, deviate more from the consensus than lower-ability analysts' forecasts in equilibrium. For each analyst–firm, we measure the average forecast deviation as follows:

$$Disp_{i,j} = \frac{\text{average}|Dev| \text{ for all forecasts in an analyst-firm pair}}{\text{average}|Dev| \text{ for all analysts covering the same firm}}.$$

The correlation between $Disp$ and $Ability(Dir)$ is 0.18, significant at less than the 1% level.¹⁹ This result indicates that more deviation from consensus is indeed a pattern typical of high-ability analysts, and therefore, overweighting private information can be viewed as a mechanism to signal ability.

3.2 Empirical results

Table 2 reports the main results testing the two hypotheses. Specification (1) estimates (7) using the linear regression to identify the average effect of covariates on misweighting. Specification (2) complements (1) with a median regression, which identifies the effects for forecasts with median

¹⁹ We also obtain qualitatively similar results when we replace $Disp_{i,j}$ with $Disp_Ind_{i,j}$, calculated as $Disp_Ind_{i,j} = \frac{\text{average}|Dev/P_{-5}| \text{ for analyst-firm pair } i,j}{\text{average}|Dev/P_{-5}| \text{ for all analyst-firm pairs in the same 2-digit SIC}}$ where P_{-5} is the stock price five days before the forecast date.

Table 2
Testing Hypotheses 1 and 2—main results

	(1)		(2)		(3)	
	Linear regression		Median regression		Maximum likelihood (probit)	
	<i>Dev</i> > 0	<i>Dev</i> < 0	<i>Dev</i> > 0	<i>Dev</i> < 0	<i>Dev</i> > 0	<i>Dev</i> < 0
<i>Ability</i> (* <i>Dev</i>)	−1.452 (−16.71)	−1.172 (−17.71)	−1.267 (−22.02)	−0.770 (−17.86)	−0.804 (−35.04)	−1.140 (−71.51)
<i>Tr</i> (* <i>Dev</i>)	0.805 (8.35)	0.474 (6.91)	0.737 (13.54)	0.417 (12.32)	0.439 (24.67)	0.501 (29.75)
<i>Vol</i> (* <i>Dev</i>)	0.726 (3.32)	0.104 (1.27)	1.159 (3.34)	−0.005 (−0.22)	0.216 (2.52)	−0.126 (−1.78)
<i>IB</i> (* <i>Dev</i>)	0.164 (1.74)	0.045 (0.65)	0.027 (0.52)	−0.005 (−0.18)	0.121 (1.34)	−0.063s (−0.56)
<i>Days</i> (* <i>Dev</i>)	0.354 (16.99)	−0.150 (−10.95)	0.455 (28.34)	−0.216 (−22.25)	0.152 (28.80)	−0.153 (−36.12)
<i>Exp</i> (* <i>Dev</i>)	0.020 (0.27)	−0.188 (−4.25)	−0.027 (−0.78)	−0.156 (−5.12)	−0.052 (−4.16)	−0.021 (−1.74)
Constant(* <i>Dev</i>)	0.759 (19.53)	−0.072 (−2.22)	0.865 (37.55)	−0.037 (−2.06)	0.349 (27.31)	−0.220 (−19.51)
Adj. <i>R</i> ²	0.121	0.105	0.034	0.007	0.042	0.061
Number of observations	472,285	674,180	472,285	674,180	447,907	633,830

Specifications (1) and (2) estimate

$$FE = \alpha_j + \gamma_0 Dev + \gamma_1 X \cdot Dev + \varepsilon$$

using linear regression and median regression, respectively, including firm fixed effects. Specification (3) estimates

$$1[\text{sign}(FE) = \text{sign}(Dev)] = \alpha + \gamma X + \varepsilon$$

using probit. *FE* is the forecast error and *Dev* is the forecast distance from the consensus. Each specification is estimated separately on subsamples of *Dev* > 0 and *Dev* < 0, with results reported under separate column headings. *X* is a vector of covariates that include the following: *Ability*, the true ability measure using all observations in an analyst–firm pair; *TR*, the track records measure that uses only past forecasts whose errors have been realized by the time of the current forecast, *TR*, and *Ability* are orthogonalized; *Vol*, the firm’s average daily trading volumes, in billions of dollars, over the 50-trading-days’ period preceding the forecast date; *IB*, dummy variable for investment banking affiliation; *Days*, number of days, in 100s, between forecast and earnings announcement; and *Exp*, number of tested forecasts, in 100s, issued by an analyst before the forecast date. See Table 1 for detailed variable definitions. Reported in the parentheses are *t*-statistics based on standard errors that adjust for heteroskedasticity as well as arbitrary correlation among observations clustered at the same quarter and same SIC two-digit industry.

FE conditional on all the covariates. The median regression is less influenced by outliers and is more robust to undue influence from right-skewness of the dependent variable (*FE*). Lastly, Specification (3) estimates (8) using probit, which reveals the effect of covariates on the tendency to misweight. The probability method is robust to outliers in both *FE* and *Dev*,²⁰ and is consistent under more general assumptions

²⁰ This is because the dependent variable is $1[\text{sign}(FE) = \text{sign}(Dev)]$, where only the relative signs of *FE* and *Dev* matter, but not their magnitudes.

about the information structure and forecasting strategy (see Section 4). Overall, the three specifications offer consistent results.²¹

3.2.1 Hypothesis 1: overconfidence or strategic overweighting. Hypothesis 1 predicts that analysts' overweighting is related to their ability, track records, and potential trading commissions. These predictions hold regardless of whether analysts receive favorable or unfavorable news. In Table 2, the coefficients of interest for Hypothesis 1 are $\hat{\gamma}_{TR}$, $\hat{\gamma}_{Ability}$, and $\hat{\gamma}_{Vol}$. Results across the different specifications are similar, so we discuss specification (1) only. The coefficients on $TR \cdot Dev$ are significantly positive at less than the 1% level on both the positive and negative Dev subsamples, and $\hat{\gamma}_{TR}^+$ is significantly different from $\hat{\gamma}_{TR}^-$ at the 10% level.²² The positive values of right-skewness $\hat{\gamma}_{TR}^+$ and $\hat{\gamma}_{TR}^-$ indicate that analysts overweight both types of news more when they have better track records, conditional on their true ability. Given that the mid-50% range (i.e., from the 25th to the 75th percentiles) of analysts' TR in our sample is 0.63 (see Table 1), a value of $\hat{\gamma}_{TR}^+ = 0.805$ implies that between two analysts with similar ability, the analyst at the 75th percentile TR overweights 51% ($= 0.63 \cdot 0.805$) more than the analyst at the 25th percentile TR when the private information is more favorable than the consensus. Similarly, for negative news, the analyst at the 75th percentile TR overweights 30% ($= 0.63 \cdot 0.474$) more than the analyst at the 25th percentile TR . The positive relation between TR and overweighting is consistent with both the overconfidence and the incentive hypotheses.

The coefficients on $Ability \cdot Dev$ are significantly negative (at less than 1%) on both subsamples (and not significantly different from each other at the 10% level), indicating that among analysts with similar track records, those with higher true ability overweight private information less. Specifically, $\hat{\gamma}_{Ability}^+ = -1.452$ ($\hat{\gamma}_{Ability}^- = -1.172$) implies that between two analysts with similar TR , the analyst at the 25th percentile of $Ability$ overweights positive (negative) news 79% (64%) more than the analyst at the 75th percentile of $Ability$. This result supports only the strategic overweighting hypothesis, and does not support the overconfidence explanation (see footnote 14). We interpret these results as suggesting that analysts know their intrinsic ability, with the less-skilled analysts bet more to stand out from the crowd and hope to be right rather than to be efficient in their forecasting.

It is important to note that the relation between misweighting and $ability/TR$ is identified only by controlling for each other. For comparison, in results not tabulated, we find that when TR is excluded, the coefficient on

²¹ In unreported sensitivity checks, we show that winsorizing (or trimming) both FE and Dev at 0.5% extremes (or at 1% extreme based on absolute values) yield similar coefficient estimates with overall higher significance levels.

²² Comparison of coefficients estimated from different subsample adjusts for cross-sample correlation of error disturbances due to industry-quarter clustering.

Ability · *Dev* is significantly negative, indicating that higher-ability analysts unconditionally overweight private information less. On the other hand, when *Ability* is excluded, the coefficient on *TR* is negative.²³ Given that *TR* is highly correlated with *Ability*, this result is consistent with analysts knowing their own ability.

The coefficient on *Vol* · *Dev* demonstrates the impact of trading volume on analysts' weighting behavior. Column 1 in Table 2 shows a significantly positive $\hat{\gamma}_{Vol}^+$ of 0.726 (*t*-statistic = 3.32) and an insignificant $\hat{\gamma}_{Vol}^-$. The difference between $\hat{\gamma}_{Vol}^+$ and $\hat{\gamma}_{Vol}^-$ is significant at the 5% level. Assuming that a given amount of news in a forecast (captured by *Dev*) generates more trading commissions for stocks with higher trading volumes, this result is consistent with the incentive hypothesis. It indicates that analysts overweight positive news about 72.6% more when the average trading volume during the recent 10 weeks is one billion dollars higher. The absence of a similar pattern for negative news is consistent with the fact that positive news spurs more trading than negative news, possibly because there are fewer potential sellers than potential buyers [due to restrictions of short-sell, see, e.g., Hong, Lim, and Stein (2000)]. Accordingly, analysts have weaker trading-based incentives to overweight unfavorable private information than they do favorable information.

Given the importance of the *Ability* and *TR* variables in our analyses, the first two columns of Table 3 provide sensitivity checks using alternative measures. Column 1 repeats the regression using *Ability*(*Beat*) and *TR*(*Beat*) as defined in (13). Results are qualitatively similar. Since ability measures based on a short history are noisy, column 2 of Table 3 reports results for the subsample of analysts with at least 30 tested forecasts. Results are similar. (We do not use this subsample in our main tests because of the potential selection bias for analysts who survived to make at least 30 forecasts.)

To increase the power of the test, we also construct an out-of-sample ability measure using the independent ranking of analysts by *Institutional Investor*. We hand-collect *Institutional Investor's* annual ranking of analysts from 1988 to 2000 and create a dummy variable *II_STAR_i* that equals 1 if analyst *i* is selected as a member of the All-American Research team in any year from 1988 to 2000. If we replace *Ability*(*Dir*) with *II_STAR* without the *TR* variable, the coefficient (not tabulated) on *II_STAR* is significantly negative (at less than 1%). If we include our

²³ The negative relation between overweighting and *TR* is also consistent with the following alternative explanation: Suppose analysts do not know their own ability but rationally update from their track records. Those with good track records would like to build on and improve their reputation; in the meantime, good track records make it less important for them to use deviation from the consensus to signal their ability. The full specification with both *TR* and *Ability* included as regressors reveals that the marginal effect of *TR* is actually positive, indicating that the negative effect from the alternative explanation, if it exists, does not dominate the positive effect predicted by the strategic overweighting hypothesis.

Table 3
Results from alternative measures of key variables

	(1)		(2)		(3)		(4)		(5)	
	<i>Ability/TR</i> based on <i>Beat</i>		Firm-analyst pair w/30 forecasts or more		<i>Past_IB</i>		<i>Future_IB</i>		<i>Future_IB (IV)</i>	
	<i>Dev</i> > 0	<i>Dev</i> < 0	<i>Dev</i> > 0	<i>Dev</i> < 0	<i>Dev</i> > 0	<i>Dev</i> < 0	<i>Dev</i> > 0	<i>Dev</i> < 0	<i>Dev</i> > 0	<i>Dev</i> < 0
<i>Ability^a Dev</i>	-1.284 (-17.78)	-1.393 (-22.14)	-1.531 (-13.20)	-1.074 (-13.66)	-1.452 (-16.71)	-1.172 (-17.71)	-1.442 (-16.61)	-1.172 (-17.68)	-1.745 (-10.57)	-1.026 (-7.95)
<i>TR^a Dev</i>	0.652 (7.26)	0.459 (7.48)	0.630 (4.82)	0.437 (5.29)	0.805 (8.35)	0.474 (6.91)	0.797 (8.30)	0.474 (6.90)	0.894 (3.97)	0.517 (3.73)
<i>Vol^a Dev</i>	0.735 (3.57)	0.103 (1.36)	0.853 (4.03)	0.074 (0.93)	0.726 (3.32)	0.104 (1.27)	0.731 (3.34)	0.105 (1.27)	1.096 (2.68)	0.541 (2.41)
<i>IB^a Dev</i>	0.178 (1.88)	0.048 (0.69)	0.232 (2.00)	0.007 (0.09)	0.164 (1.74)	0.045 (0.65)	1.680 (2.92)	0.109 (0.34)	2.453 (1.82)	-1.235 (-1.01)
<i>Days^a Dev</i>	0.360 (22.78)	-0.145 (-12.18)	0.372 (21.92)	-0.142 (-11.01)	0.354 (16.99)	-0.150 (-10.95)	0.353 (17.00)	-0.150 (-10.95)	0.409 (10.12)	-0.123 (-4.82)
<i>Exp^a Dev</i>	0.051 (0.75)	-0.184 (-4.59)	0.061 (0.88)	-0.171 (-4.15)	0.020 (0.27)	-0.188 (-4.25)	0.024 (0.32)	-0.188 (-4.24)	-0.331 (-1.92)	-0.184 (-1.38)
<i>Dev</i>	0.713 (20.78)	-0.083 (-2.89)	0.743 (18.72)	-0.075 (-2.38)	0.759 (19.53)	-0.072 (-2.22)	0.763 (20.03)	-0.070 (-2.17)	0.441 (2.22)	-0.025 (-0.15)
Adj. <i>R</i> ²	0.120	0.107	0.118	0.101	0.121	0.105	0.122	0.105	0.127	0.112
Number of observations	472,285	674,180	373,614	534,550	472,285	674,180	472,285	674,180	91,285	124,168

This table repeats specification (1) in Table 1 using alternative measures for *Ability/TR* and *IB*. In column 1, *Ability(Beat)* and *TR(Beat)* replace the default ability and track records measures (see Table 1 for detailed definition). Column 2 only includes observations from analyst-firm pairs that have at least 30 observations in the full sample (so that analysts' ability could be measured reliably). Columns 3 and 4 decompose *IB* into investment banking affiliation during the two-year window before and after the forecast date year. Finally, column 5 uses an instrumented *IB* that is the predicted future investment banking possibility using time-adapted information (such as brokerage size and analysts' experience) on the subsample of underwriting banks only. Reported in the parentheses are *t*-statistics based on standard errors that adjust for heteroskedasticity as well as arbitrary correlation among observations clustered at the same quarter and same SIC two-digit industry.

main $TR(Dir)$ measure, the coefficient on II_STAR remains significantly negative at less than 1%, and the coefficient on $TR(Dir)$ is significantly positive at less than the 1% level.

In summary, the above results indicate that higher-ability analysts are less likely to overweight their private information than lower-ability analysts. Our results further suggest that the positive correlation between ability and distance (documented earlier) does not provide unambiguous inference on analysts' weighting behavior, and as such, is not a good measure of whether analysts' are bold or herd. Further, our findings suggest that analysts condition their weighting on ability (because otherwise, we would expect a negative correlation between ability and deviation).²⁴

3.2.2 Hypothesis 2: optimistic bias or strategic misweighting. The variables of interest to Hypothesis 2 are investment banking affiliation (IB), analyst experience (Exp), and the distance between the forecast date and the earnings release date ($Days$).

Investment banking incentives. The IB variable in Table 2 is a dummy variable equal to one, if the analyst's employer has an investment banking relationship with the firm that the analyst covers during the five-year period centered on the year of the forecast. We find no evidence that affiliated analysts engage in more optimistic weighting than unaffiliated analysts, as neither $\hat{\gamma}_{IB}^+$ nor $\hat{\gamma}_{IB}^-$ is distinguishable from zero at the 5% significance level. In unreported results, we restrict the analysis to only forecasts made by analysts from banks that participated in at least one underwriting deal during the sample period and find no significant coefficients. In addition, to capture a brokerage house's status in investment banking business, we replace IB with a dummy variable that equals one, if the brokerage house is one of the elite investment banks per Hong and Kubik (2003). The coefficient estimates are, again, statistically insignificant.

In Table 3, columns 3 and 4, we decompose the IB dummy into $Past_IB$ and $Future_IB$. $Past_IB$ ($Future_IB$) is a dummy variable equal to one, if the forecast is for a firm for which the analyst's employer underwrote security issuances in the two-year period before (after) the forecast. If analysts' optimistic bias is unintentional and thus does not vanish after the deals are completed, then we should not expect the coefficient for $Past_IB$ to be significantly different from that for $Future_IB$. Column 3 shows that $\hat{\gamma}_{Past_IB}^+$ and $\hat{\gamma}_{Past_IB}^-$ are not significantly different from 0 at the 5% level, suggesting that analysts employed by firms with prior underwriting relationships do not misweight information more than unaffiliated analysts. However, when we use $Future_IB$, the coefficient $\hat{\gamma}_{Future_IB}^+$ is 1.65 and significant at less than the 1% level, whereas

²⁴ If the actual weight k is constant or uncorrelated with ability h , then the deviation, $E(f - c)^2 = \frac{h}{\rho_c} \left(\frac{k}{h}\right)^2$, should be negatively correlated with h .

$\hat{\gamma}_{\text{Future_IB}}^-$ is indistinguishable from zero. Assuming that analysts are informed about potential underwriting relationships, these results suggest that analysts affiliated with banks seeking future underwriting business tend to weight favorable news by 165% more relative to unaffiliated analysts. In contrast, they do not underweight unfavorable news more than unaffiliated analysts. We repeat the above analyses allowing year-specific effects on misweighting (adding year dummies interactive with *Dev*) to account for the possible effects of investment banking business cycles, results (not tabulated) are qualitatively similar.

Using *Future_IB* as a regressor suffers from a potential endogeneity problem. If an analyst's optimism increases the probability that his employer obtains the underwriting business, then a positive relation between overweighting favorable news and future underwriting business will arise. And, since the analyst's optimism can be either unintentional or strategic, a positive $\hat{\gamma}_{\text{Future_IB}}^+$ cannot rule out the behavioral bias hypothesis.

We use an instrumental variable approach to probe this issue. Specifically, we try to filter out the endogenous component in *Future_IB* using a predicted value of *Future_IB* from instruments. We obtain the projected variable, $\widehat{\text{Future_IB}}$, using the subsample of forecasts made by analysts affiliated with underwriting banks only, from a set of time-adapted exogenous variables: (i) total underwriting deals in which the bank participated during the past five years; (ii) total number of analysts employed by the bank in the current year; (iii) total number of analysts employed by the bank which covers the firm; (iv) total number of other banks with analysts covering the firm during the past three years; (v) a dummy variable if the bank is an elite bank per Hong and Kubik (2003); (vi) the market capitalization of the firm at the beginning of the current year; and (vii) the return of the firm's stock over the past two years. Column 5 of Table 3 reports results from using the instrumented $\widehat{\text{Future_IB}}$. The coefficient estimate $\hat{\gamma}_{\text{Future_IB}}^+ \searrow$ remains positive with an estimate of 2.45 (significant at the 10% level) and $\hat{\gamma}_{\text{Future_IB}}^- \searrow$ is not significant at conventional levels.

These results suggest that analysts overweight more on positive news when the conditional probability of future underwriting deals is higher. We hesitate to emphasize these results too strongly, because the extra variance introduced by instrument variable estimation renders the significance level low.

Overall, affiliated analysts' earnings forecasts may not be as optimistically biased as their stock recommendations or their long-term growth projections as found in prior studies (or as reported in the popular press). Two reasons may explain the difference. First, unlike stock recommendation or long-term growth forecasts, the accuracy of earnings forecasts is clearly established on earnings release dates. Second, barring firms managing earnings to meet or beat analysts' forecasts, earnings are more or less

exogenous to analysts' forecasts. In contrast, stock prices can be endogenously affected by analysts' recommendations, making it more difficult to assess the accuracy of recommendations, at least in the short run. Together, our results (as well as findings in prior studies) imply that analysts are more disciplined when the performance metric for their forecast accuracy is more clearly defined.

Forecast timing and analyst experience. Tables 2 and 3 also report the effects of forecast timing (*Days*) and experience (*Exp*) on analysts' misweighting. Across all specifications, the coefficient estimates for *Days* · *Dev* in the positive *Dev* subsample ($\hat{\gamma}_{\text{Days}}^+$) are positive and those in the negative *Dev* subsample ($\hat{\gamma}_{\text{Days}}^-$) are negative, all significant at less than the 1% level. These results indicate that analysts overweight positive news more, and underweight negative news more, when the forecasts are made farther away from the earnings release date. For example, Table 2 column 1 shows that $\hat{\gamma}_{\text{Days}}^+ = 0.354$, suggesting that for a favorable forecast, analysts would overweight 35.4% more if the forecast is issued 100 days earlier. Similarly, $\hat{\gamma}_{\text{Days}}^- = -0.150$ implies that for a forecast of negative news, analysts would underweight 15% more if they issued the forecast 100 days earlier.²⁵

As a robustness check, we also calculate the probability statistic $\hat{\pi}$ for subsamples of forecasts sorted by *Days* (results not tabulated). We find that the probability of overweighting positive news decreases approximately linearly from about 0.8 to 0.5 (the neutral value), when the interval between the forecast date and the earnings announcement declines from two years to one month. The same probability for negative news for the same interval increases from about 0.25 (indicating underweighting) to 0.5. That is, analysts' weightings are close to the efficient weightings for forecasts made within a month of earnings announcement. To the extent that the cost of *FE* is higher when earnings realization dates are closer, these results are consistent with the incentive hypothesis.

Table 2 shows that $\hat{\gamma}_{\text{Exp}}^-$ is negative and significant at the 1% level in columns 1 and 2 and at the 10% level in column 3. These results indicate that analysts underweight negative news more as they become more experienced. $\hat{\gamma}_{\text{Exp}}^- = -0.188$ (in column 1) indicates that when issuing negative forecasts, an analyst underweights 18.8% more than a similar analyst with 100 fewer prior-tested forecasts. Evidence is mixed on the effect of experience on analysts' weighting of positive news: $\hat{\gamma}_{\text{Exp}}^+$ is indistinguishable from zero in columns 1 and 2 and significantly negative in column 3 at the 1% level. Results are similar (not reported) when we

²⁵ These results are not driven by the fact that forecasts closer to the earnings release dates are, on average, more accurate. Recall that the slope coefficient β_0 in (5) measures the weight analysts place on private signals relative to the efficient weight (i.e., $1 - \frac{h}{k}$), not the precision of information (i.e., h) or the actual weights (i.e., k).

measure Exp with the number of years since an analyst's first appearance in the Zacks database.

The fact that we do not find strong evidence for misweighting diminishing with experience (negative $\hat{\gamma}_{Exp}^-$ and mostly insignificant $\hat{\gamma}_{Exp}^+$) seems to suggest that analysts either do not learn from their experience or that they misweight intentionally. Learning may not occur if the optimism in an analyst's past forecasts is not salient. If learning takes place, we expect its effect to be more pronounced when an analyst's previous forecasts indicate stronger evidence of bias. To examine this conjecture, we replace $Exp \cdot Dev$ with $Bias \cdot Dev$, where $Bias_{i,j,t}$ is the sum of $B_{i,j,t_q} = \text{sign}(FE_{i,j,t_q})$, $t_q = 1, \dots, Q$, for all Q forecasts by analyst i for firm j whose errors have been observed by t . If an analyst is not systematically biased, $Bias_{i,j,t}$ should be close to zero. We use the sum (instead of the average) of $B_{i,j,t}$ to account for the accuracy of $Bias_{i,j,t}$ as an estimate of the true bias.²⁶

The optimistic bias hypothesis predicts that if analysts correct their bias from learning, then $\hat{\gamma}_{Bias}^+ < 0$ and $\hat{\gamma}_{Bias}^- > 0$. On the other hand, if no learning takes place, then past bias predicts future bias, and the predictive power is higher with longer track record. Accordingly, $\hat{\gamma}_{Bias}^+ > 0$ and $\hat{\gamma}_{Bias}^- < 0$. The estimated coefficients (not tabulated) are $\hat{\gamma}_{Bias}^+ = 0.032$ and $\hat{\gamma}_{Bias}^- = -0.015$, both significant at less than 1%. These estimates are consistent with the hypothesis of strategic misweighting, or with the hypothesis that no learning takes place.

3.2.3 Alternative explanations. Our empirical analysis indicates evidence of optimistic weighting. We note that optimistic weighting will result in positively biased forecasts, similar to what one would observe if analysts added a positive bias to their otherwise efficiently weighted forecasts (the "added-bias" hypothesis). For illustration, we use the same notation as in (4) but modify the analyst's forecasting strategy to allow him to add a bias to his forecast:

$$f = b + ky + (1 - k)c, \quad \text{where } k = \bar{k}(\underline{k}) \text{ for } y > (<)c, \quad (14)$$

and b is the added bias. Then, in regression (5), the bias b is subsumed in the intercept, while the slope estimate will capture the average misweighting. In addition, if misweighting is symmetric (i.e., $\bar{k} = \underline{k} \neq h$), then $E(FE) = b$; if optimistic weighting exists (i.e., $\bar{k} > h > \underline{k}$), then $E(FE) > b$.

When we regress (5) on the positive and negative Dev subsamples separately (with firm fixed effects), the average firm-specific intercepts for both subsamples are nearly identical at \$0.128. The slope coefficients, on the other hand, are very different (0.75 and -0.14, respectively, as

²⁶ An analyst with one previous optimistic forecast is not expected to become aware of the bias; whereas an analyst with eight out of ten past forecasts being positively biased should be aware of the bias.

discussed in Section 2.2), indicating optimistic weighting. Also consistent with optimistic weighting, we find that the average *FE* for the whole sample is \$0.165, higher than the intercepts in both subsamples (\$0.128).²⁷ Lastly, Figure 1 offers similar inference in that the V-shaped graph intersects the vertical axis at a positive value. Together, data indicates that optimistic weighting exists separately from added bias.

We also note that certain statistical properties in the data can only result from misweighting, and not from added bias. If $\bar{k} > h > \underline{k}$ in (14), then the density function of *Dev* should have a fatter tail for the *Dev* > 0 region (analysts exaggerate good news and issue forecasts more dispersed from the consensus) than for the *Dev* < 0 region (analysts compress negative news and issue forecasts closer to the consensus). The density function of *Dev* plotted in Figure 1 evidences this pattern. Note that this pattern could not result from added bias alone because added bias would only cause a parallel shift in the distribution.²⁸

Under certain assumptions about how analysts adjust their added bias, the added-bias hypothesis provides scenarios that are observationally similar to misweighting. The first scenario concerns an alternative explanation to underweighting negative news. Richardson, Teoh, and Wysocki (2001) find that some analysts start off a fiscal year with a high bias and revise their forecasts down toward the earnings release date. Such downward revisions can generate forecast patterns similar to underreaction to negative private news. To assess the sensitivity of our results to this possibility, we exclude from our sample “sequential” revising-down analyst-firm pairings, defined as analyst-firm pairs with more than 75% of the forecasts revised down (from the consensus). This eliminates about 15% of the observations. We then estimate (7) on the negative *Dev* subsample, and obtain qualitatively similar results to those in Table 2 for all covariates. Most importantly, the coefficient for *Dev* is -0.087, significantly negative at less than the 1% level.

A second scenario concerns an alternative explanation to overweighting positive private signals. It is possible that an experienced analyst is more likely to receive early news about the firm, to be the first to revise, and to deviate the most from the consensus. A more experienced analyst is also likely to be freer to add a higher bias to her private estimate (because, as theory suggests, the penalty of being wrong is lower for

²⁷ While the total bias cannot be exactly decomposed into added bias and bias attributable to optimistic weighting, the difference between \$0.165 and \$0.128 is a lower-bound estimate for the optimistic weighting induced bias. Because the magnitude of misweighting varies cross-sectionally, the intercept also contains part of the optimistic weighting induced bias that is not captured by the coefficient estimates from the linear regression.

²⁸ Note that the asymmetry of the distribution of *Dev* around zero is not a result of the skewness in forecasts or in the underlying earnings. *Dev* is the difference between the forecast and the consensus in which the skewness of the underlying earnings offsets. In our sample, the skewness of individual forecasts and their corresponding consensus is very close (2.7 versus 2.8).

more established analysts). The less experienced analysts who follow make smaller deviations from the consensus forecast, but their optimal bias is also lower. In this scenario, FE and deviations from the consensus will be positively correlated when the private signal is favorable, creating the impression that analysts on average overweight positive private signals. To assess the sensitivity of our results to this scenario, we note first that while our sample experienced analysts engage in more optimistic weighting, they primarily do so by compressing negative news (as opposed to exaggerating good news). (Table 2 shows that $\hat{\gamma}_{Exp}^-$ is significantly negative while $\hat{\gamma}_{Exp}^+$ mostly insignificant.) Further, on the positive Dev subsample, experienced analysts do not deviate more from consensus than inexperienced analysts: the correlation between Dev and Exp is 0.02.

4. Robustness Checks

4.1 Extension to general information structures

The tests developed in Section 1 can be generalized to situations where the underlying variables are not normally distributed and the forecasts are not linear sums of signals. Let $f = f(y, c)$ be the analyst's forecast. We generalize our definitions of efficient weighting and misweighting as follows:

Definition 2. *Suppose public and private signals about earnings are generated according to (1) and (2) where ε_c and ε_y follow some general continuous distributions. We say an analyst efficiently weights information if his forecast f minimizes forecast errors. We say an analyst overweights (underweights) private information if his forecast deviates too much (too little) from the consensus in the same direction as his private signal's deviation from the consensus.*

The next Proposition extends Proposition 1 and specifies two criteria for minimizing forecast errors to accommodate the more general information structure (the proof is in Appendix A).

Proposition 3. *Suppose public and private signals about earnings are generated according to (1) and (2) where ε_c and ε_y follow some continuous distributions. (i) Suppose $E(\varepsilon_y) = 0$. If an analyst weights information to minimize $MSE = E(f - z)^2$, then β_0 in (5) is zero. Further, $\beta_0 > 0$ if the analyst overweights private information, and $\beta_0 < 0$ if he underweights. (ii) Suppose $Med(\varepsilon_y) = 0$, where Med stands for the median function. If an analyst weights information to minimize mean absolute forecast error $MAE = E|f - z|$, then π in (6) is $\frac{1}{2}$. Further, $\pi > \frac{1}{2}$ if the analyst overweights private information, and $\pi < \frac{1}{2}$ if he underweights.*

Proposition 3 shows that the consistency of the regression test and the probability test does not require that the underlying variables be normally distributed. Assuming $E(\varepsilon_y) = 0$ or $\text{Med}(\varepsilon_y) = 0$ is without loss of generality. As long as analysts minimize mean-squared forecast errors (MSE) and their signals are correct on average (i.e., they have zero-mean disturbance), then the regression-based test is consistent for testing misweighting. If analysts minimize mean absolute forecast errors (MAE) and their signals are correct in probability (i.e., they have zero median disturbance), then the probability test is consistent for detecting misweighting. Evidence that both MSE and MAE are reasonable criteria for forecast accuracy is provided by Abarbanell and Lehavy (2003), Cohen and Lys (2003), Gu and Wu (2003), and Basu and Markov (2004).

4.2 Sensitivity analyses

This section reports several sensitivity checks. First, firms may have incentives to manage earnings in order to beat or meet analysts' forecasts [see, e.g., Degeorge, Patel, and Zeckhauser (1999) and Abarbanell and Lehavy (2003)]. If so, the earnings number that analysts attempt to forecast may not be strictly exogenous. While modeling the earning management game is beyond the scope of the article [see, e.g., Liu and Yao (2002) and Liu (2003) for this topic], we assess the effect of earnings management by reestimating our main regressions excluding forecasts issued less than 30 days from the earnings release dates. To the extent that management is more likely to manage earnings in response to forecasts made right before the earnings announcement, reported earnings will be more exogenous to individual forecasts made farther away from the earnings announcement. The results (not reported) are very similar to those reported in Table 2.

Second, we obtain similar results when scaling both *FE* and *Dev* by stock price five days before the forecast date. Both scaled and unscaled variables identify similar optimistic weighting patterns (see also Figure 1). In the fixed-effects regressions with scaled variables, the $\hat{\beta}$ estimates for the positive and negative *Dev* subsamples are 0.95 ($t = 22.06$) and -0.59 ($t = 15.19$), respectively.²⁹

Lastly, we check the sensitivity of our results to the consensus measure (a weighted average of outstanding forecasts) in two ways. First, we note that the probability tests are more robust to measurement errors because such errors affect these tests only when the errors are large

²⁹ The only case where the scaled and unscaled variables produce significantly different results is the regression of *FE* on *Dev* alone for the whole sample: the estimated coefficient on *Dev* is -0.17 ($t = -7.09$), which contrasts with the average overweighting results in Section 3.2. The difference suggests that analysts covering firms with high stock prices overweight more. It is consistent with the results in Tables 2 and 3 that analysts overweight more for stocks with larger trading volumes. In our sample, the correlation between stock price and trading volume is 0.38.

enough to change the sign of Dev . In such cases, an observation that actually has $\text{sign}(FE^*) = \text{sign}(Dev^*)$ (in favor of overweighting) could be misclassified as $\text{sign}(FE) \neq \text{sign}(Dev)$ (in favor of underweighting), or vice versa. When this occurs, the measurement error in Dev becomes a misclassification error in the dependent variable in estimating (8). The impact of this misclassification error can be assessed using the Hausman, Abrevaya, and Scott-Morton (1998) method that allows the probability of misclassification in binary dependent variables to be estimated simultaneously with the γ_m coefficients.³⁰ The resulting coefficient estimates are qualitatively similar to those listed in Table 2 column 3, and are not reported.

We also construct a new consensus measure, $\hat{c}$, as a proxy for the best earnings predictor using all available public information at the time of the forecast. Briefly, we first project all realized earnings on variables that represent past public information on every sample forecast date t , including (i) firms' realized earnings four and eight quarters ago, (ii) the two most recent forecasts issued prior to the current forecast, (iii) the number of days from t to the end of forecasted quarter, and (iv) a dummy for whether the forecast is for firms' fourth fiscal quarter. We then use the estimated coefficients to extrapolate the current consensus (details available upon request). The resulting consensus measures $\hat{c}$ are less biased, but also less accurate, in terms of both MSE and mean absolute error (MAE), than the weighted averages of outstanding forecasts. This result is consistent with Brown, Hagerman, Griffin, and Zmijewski's (1985) finding that analysts' consensus forecasts are in general more accurate than predictions based on statistical models. We repeat our tests using this measure, assuming analysts and investors have access to all the public information as we do. Results (not reported) show significant optimistic weighting (overweight when $Dev > 0$ and underweight when $Dev < 0$), similar to the results reported in Section 2.

5. Conclusion

Much of the existing literature on security analysts studies properties of their observed forecasts such as accuracy and bias, and draw inference about behavior and incentive motivating such properties. We argue that these realized forecast properties do not provide unambiguous inferences about analysts' forecast behaviors because other factors, such as analysts' private information and randomness in the actual earnings

³⁰ Let Dev^* and Dev be the true and measured deviations from consensus. Then the misclassification probability is defined as $\theta = \Pr(\text{sign}(Dev) \neq \text{sign}(Dev^*))$. The maximum likelihood function is: $\mathcal{L}(\theta, \gamma) = n^{-1} \sum_{i=1}^n y_i \ln(\frac{1}{2}\theta + (1-\theta)\Phi(X_i\gamma)) + (1-y_i) \ln(1 - \frac{1}{2}\theta - (1-\theta)\Phi(X_i\gamma))$. We thank Jason Abrevaya for providing codes for the estimation routine.

realizations, also affect realized forecast properties. We address these issues by providing evidence on analysts' ex ante forecast behaviors in terms of how they weight private and public information when they forecast earnings. We develop two methods to extract information about analysts' weighting behavior and document two findings. First, on average, analysts overweight private information; and second, analysts weight information optimistically in that between private and public information, they overweight the relatively favorable one. We further explore the potential sources of analysts' misweighting behaviors and find that the degree of misweighting is positively related to the benefits, and negatively related to the costs, of misweighting. We interpret these findings as analysts' incentives playing a larger role in misweighting than their behavioral bias.

In addition to providing new evidence about analysts' forecast behaviors, the findings in this article also have implications for how the stock market forms expectations about firm earnings, given that analysts' forecasts are an important input into the market's expectation. Further, to the extent that analysts represent sophisticated investors, evidence on their weighting behaviors sheds light on how information is processed and transmitted in financial markets in general. Certain stock return anomalies such as price momentum, slow diffusion of bad news, or post earnings announcement drifts are consistent with investors' inefficient use of information. To account for these anomalies, researchers have developed models which make assumptions about how investors process information and respond to changes in their environment. Evidence presented here can potentially aid researchers in assessing the validity of these models.

Appendix A

Proof of Proposition 1. It is equivalent to show that the probability $\Pr(\text{sign}(f - z) = -\text{sign}(f - c))$ is $\frac{1}{2}$ when $k = h$, and is greater (smaller) than $\frac{1}{2}$ if $k < (>)h$. With symmetric distributions, $\Pr(\text{sign}(f - z) = -\text{sign}(f - c))$ amounts to

$$\Pr(c < f < z) + \Pr(z < f < c) = 2\Pr(c < f < z).$$

First, we prove that $\Pr(c < f < z) = 1/4$ when $k = h$. Notice that

$$\Pr(c < f < z) = \Pr(\varepsilon_c < k\varepsilon_y + (1 - k)\varepsilon_c < 0) = \Pr\left(\varepsilon_c < \varepsilon_y < -\frac{1 - k}{k}\varepsilon_c\right). \quad (\text{A1})$$

Since only the relative precision of ε_c and ε_y matters, we normalize ε_c to be a standard normal variable. Let ϕ and Φ stand for the probability density and cumulative probability functions of the standard normal distribution. Then, when $k = h$, (A1) becomes

$$\begin{aligned}
 \Pr\left(\varepsilon_c < \varepsilon_y < -\frac{1-h}{h}\varepsilon_c\right) &= \Pr\left(\varepsilon_c < \varepsilon_y < -\sigma_y^2\varepsilon_c\right) \\
 &= \int_{-\infty}^0 \left[1 - \Phi\left(\frac{\sigma_y^2\varepsilon_c}{\sigma_y}\right) - \Phi\left(\frac{\varepsilon_c}{\sigma_y}\right)\right] \phi(\varepsilon_c) d\varepsilon_c. \\
 &= \frac{1}{2} - \int_{-\infty}^0 \left[\Phi(\sigma_y\varepsilon_c) + \Phi\left(\frac{\varepsilon_c}{\sigma_y}\right)\right] \phi(\varepsilon_c) d\varepsilon_c.
 \end{aligned} \tag{A2}$$

Integration by parts and change of variables yield:

$$\begin{aligned}
 \int_{-\infty}^0 \Phi(\sigma_y\varepsilon_c) \phi(\varepsilon_c) d\varepsilon_c &= [\Phi(\sigma_y\varepsilon_c)\Phi(\varepsilon_c)]_{-\infty}^0 - \sigma_y \int_{-\infty}^0 \phi(\sigma_y\varepsilon_c) \Phi(\varepsilon_c) d\varepsilon_c \\
 &= \frac{1}{4} - \int_{-\infty}^0 \phi(t) \Phi\left(\frac{t}{\sigma_y}\right) dt,
 \end{aligned} \tag{A3}$$

and

$$\int_{-\infty}^0 \Phi\left(\frac{\varepsilon_c}{\sigma_y}\right) \phi(\varepsilon_c) d\varepsilon_c = \frac{1}{4} - \int_{-\infty}^0 \phi(t) \Phi(\sigma_y t) dt. \tag{A4}$$

Substituting (A3) and (A4) into (A2) yields

$$\begin{aligned}
 &\frac{1}{2} - \int_{-\infty}^0 \left[\Phi(\sigma_y\varepsilon_c) + \Phi\left(\frac{\varepsilon_c}{\sigma_y}\right)\right] \phi(\varepsilon_c) d\varepsilon_c \\
 &= \frac{1}{2} - \left\{ \frac{1}{4} - \int_{-\infty}^0 \phi(t) \Phi\left(\frac{t}{\sigma_y}\right) dt + \frac{1}{4} - \int_{-\infty}^0 \phi(t) \Phi(\sigma_y t) dt \right\} \\
 &= \int_{-\infty}^0 \phi(t) \Phi\left(\frac{t}{\sigma_y}\right) dt + \int_{-\infty}^0 \phi(t) \Phi(\sigma_y t) dt \\
 &= \int_{-\infty}^0 \left[\Phi(\sigma_y\varepsilon_c) + \Phi\left(\frac{\varepsilon_c}{\sigma_y}\right)\right] \phi(\varepsilon_c) d\varepsilon_c.
 \end{aligned}$$

Therefore,

$$\int_{-\infty}^0 \left[\Phi(\sigma_y\varepsilon_c) + \Phi\left(\frac{\varepsilon_c}{\sigma_y}\right)\right] \phi(\varepsilon_c) d\varepsilon_c = 1/4.$$

Now suppose $k > h$ (overweighting). Then $\frac{1-k}{k} < \frac{1-h}{h}$. Accordingly

$$\begin{aligned}
 \Pr\left(\varepsilon_c < \varepsilon_y < -\frac{1-k}{k}\varepsilon_c\right) &= \Pr\left(\varepsilon_c < \varepsilon_y < -\frac{1-k}{k}\varepsilon_c \mid \varepsilon_c < 0\right) \Pr(\varepsilon_c < 0) \\
 &< \Pr\left(\varepsilon_c < \varepsilon_y < -\frac{1-h}{h}\varepsilon_c \mid \varepsilon_c < 0\right) \Pr(\varepsilon_c < 0) = 1/4.
 \end{aligned}$$

Similar argument shows that $\Pr(\varepsilon_c < \varepsilon_y < -\frac{1-k}{k} \varepsilon_c) > 1/4$ when $k < h$ (underweighting), which completes the proof.

Proof of Proposition 2. Due to symmetry,

$$\begin{aligned}\Pr[\text{sign}(f - c) = \text{sign}(z - c)] \\ &= \Pr(f > c | z > c) = \Pr(y > c | z > c) \\ &= \Pr(\varepsilon_y > \varepsilon_c | \varepsilon_c < 0) = 1 - \Phi\left(\frac{\varepsilon_c}{\sigma_y} | \varepsilon_c < 0\right) \\ &= \Phi\left(\frac{|\varepsilon_c|}{\sigma_y}\right)\end{aligned}$$

which is strictly decreasing in σ_y , or increasing in h , but completely free from k for a given h .

Proof of Proposition 3. (i) Conditional on the signals, $E(z|y, c)$ minimizes MSE because

$$\text{MSE} = E(f - z)^2 = [f - E(z|y, c)]^2 + E([E(z|y, c) - z]^2).$$

Let $FE^* = E(z|y, c) - z$ and $Dev^* = E(z|y, c) - c$. Then under efficient weighting,

$$\begin{aligned}\beta_0 &\propto E(FE^* \cdot Dev^*) - E(FE^*)E(Dev^*) \\ &= E[E(FE^* | Dev^*) \cdot Dev^*] \\ &= E[E(E(z|y, c) - z | y, c) \cdot Dev^*] \\ &= 0.\end{aligned}$$

Let f be the analyst's actual forecast, $Dev = f - c$ and $FE = f - z$. Suppose the analyst overweights private information in his forecast. Assuming no perversion in forecasts, then $Dev > Dev^*$ when $Dev > 0$, and $Dev < Dev^*$ when $Dev < 0$. Consequently,

$$\begin{aligned}\beta_0 &\propto \text{cov}(FE, Dev) \\ &= \text{cov}[FE^* + (Dev - Dev^*), Dev] \\ &= \text{cov}[(Dev - Dev^*), Dev] \\ &= \sigma_{Dev}^2 - \rho_{Dev, Dev^*} \sigma_{Dev} \sigma_{Dev^*},\end{aligned}$$

where ρ and σ stand for the coefficient of correlation and standard deviation. Because $\rho_{Dev, Dev^*} \leq 1$ and $\sigma_{Dev^*} \leq \sigma_{Dev}$, $\beta_0 > 0$ when the analyst overweights. Symmetrically, $\beta_0 < 0$ when the analyst underweights.

(ii) Let f^* be the forecast that minimizes $\text{MAE} = E|f - z|$ conditional on the signals, then $\text{Med}(FE(f^*) | c, y) = 0$. Let $FE^* = f^* - z$ and $Dev^* = f^* - c$. Then under efficient weighting,

$$\begin{aligned}\pi &= \Pr(\text{sign}(FE^*) = \text{sign}(Dev^*)) \\ &= \Pr(FE^* > 0 | Dev^* > 0) \Pr(Dev^* > 0) + \Pr(FE^* < 0 | Dev^* < 0) \Pr(Dev^* < 0) \\ &= \frac{1}{2} [\Pr(Dev^* > 0) + \Pr(Dev^* < 0)] = \frac{1}{2}.\end{aligned}$$

Let f be the actual forecast, FE and Dev be f 's forecast error and deviation from consensus, respectively.

Suppose $Dev > 0$. If the analyst overweights private information ($Dev > Dev^*$), then

$$\begin{aligned}\Pr(\text{sign}(FE) = \text{sign}(Dev) | Dev > 0) \\ &= \Pr(FE > 0 | Dev > 0) \\ &= \Pr(FE^* + (Dev - Dev^*) > 0 | Dev > 0) \\ &> \Pr(FE^* > 0 | Dev > 0) = \frac{1}{2}.\end{aligned}$$

Similar arguments show that $\Pr(\text{sign}(FE) = \text{sign}(Dev) | Dev < 0) > \frac{1}{2}$. Therefore $\pi > \frac{1}{2}$ when the analyst overweights private information. By symmetry, $\pi < \frac{1}{2}$ when the analyst underweights.

References

- Abarbanell, J., and R. Lehavy, 2003, "Biased Forecasts or Biased Earnings? The Role of Earnings Management in Explaining Apparent Optimism and Inefficiency in Analysts' Earnings Forecasts," *Journal of Accounting and Economics*, 36, 105–146.
- Abarbanell, J. S., and V. Bernard, 1992, "Analysts' Overreaction/Underreaction to Earnings Information as an Explanation for Anomalous Stock Price Behavior," *Journal of Finance*, 47, 1181–1207.
- Avery, C., and J. Chevalier, 1999, "Herding over the Career," *Economics Letters*, 63, 327–333.
- Basu, S., and S. Markov, 2004, "Loss Function Assumptions in Rational Expectations Tests on Financial Analysts' Earnings Forecasts," *Journal of Accounting and Economics*, 38, 171–203.
- Brown, L. D., R. Hagerman, P. Griffin, and M. E. Zmijewski, 1985, "Security Analyst Superiority Relative to Univariate Time-Series Models in Forecasting Quarterly Earnings," *Journal of Accounting and Economics*, 61–87.
- Chen, Q., J. Francis, and W. Jiang, 2005, "Investor Learning about Analyst Predictive Ability," *Journal of Accounting and Economics*, 39, 3–24.
- Clement, M., and S. Tse, 2004, "Financial Analyst Characteristics and Herding Behavior in Forecasting," *Journal of Finance*, 60, 307–341.
- Cohen, D., and T. Lys, 2003, "A Note on Analysts' Earnings Forecast Error Distribution," *Journal of Accounting and Economics*, 36, 147–164.
- Cooper, R. A., T. E. Day, and C. M. Lewis, 2001, "Following the Leader: A Study of Individual Analysts' Earnings Forecasts," *Journal of Financial Economics*, 61, 383–416.
- Davison, A., and D. Hinkley, 1997, *Bootstrap Methods and Their Application*, Cambridge University Press, Cambridge, UK.
- DeBondt, W. F. M., and R. H. Thaler, 1990, "Do Security Analysts Overreact?," *American Economic Review*, 80, 52–57.
- DeGeorge, F., J. Patel, and R. Zeckhauser, 1999, "Earnings Management to Exceed Thresholds," *Journal of Business*, 72, 1–33.
- Diether, K., C. Malloy, and A. Scherbina, 2001, "Differences of Opinion and the Cross-Section of Stock Returns," *Journal of Finance*, 57, 2113–2141.
- Dugar, A., and S. Nathan, 1995, "The Effect of Investment Banking Relationships on Financial Analysts' Earnings Forecasts and Investment Recommendations," *Contemporary Accounting Research*, 12, 131–160.
- Easterwood, J. C., and S. R. Nutt, 1999, "Inefficiency in Analysts' Earnings Forecasts: Systematic Misreaction or Systematic Optimism?," *Journal of Finance*, 54, 1777–1797.
- Ehrbeck, T., and R. Waldmann, 1996, "Why are Professional Forecasters Biased? Agency versus Behavioral Explanations," *Quarterly Journal of Economics*, 111, 21–41.

- Ellison, G., and S. Ellison, 2000, "A Simple Framework for Nonparametric Specification Testing," *Journal of Econometrics*, 96, 1–23.
- Festinger, L., 1957, *Theory of Cognitive Dissonance*, Stanford University Press, Stanford.
- Francis, J., and D. Philbrick, 1993, "Analysts' Decisions as Products of a Multi-Task Environment," *Journal of Accounting Research*, 31, 216–230.
- Gervais, S., and T. Odean, 2001, "Learning to be Overconfident," *Review of Financial Studies*, 14, 1–27.
- Griffin, D., and A. Tversky, 1992, "The Weighing of Evidence and the Determinants of Overconfidence," *Cognitive Psychology*, 24, 411–435.
- Gu, Z., and J. S. Wu, 2003, "Earning Skewness and Analyst Forecast Bias," *Journal of Accounting and Economics*, 35, 5–29.
- Hastorf, A. H., D. Schneider, and J. Polefka, 1970, *Person Perception*, Addison-Wesley, Reading, MA.
- Hausman, J., J. Abrevaya, and F. Scott-Morton, 1998, "Misclassification of the Dependent Variable in a Discrete-Response Setting," *Journal of Econometrics*, 87, 239–269.
- Holmstrom, B., 1999, "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66, 169–182.
- Hong, H., and J. Kubik, 2003, "Analyzing the Analysts: Career Concerns and Biased Earnings Forecasts," *Journal of Finance*, 63, 313–351.
- Hong, H., J. Kubik, and A. Solomon, 2000, "Security Analysts' Career Concerns and Herding of Earnings Forecasts," *Rand Journal of Economics*, 31, 121–144.
- Hong, H., T. Lim, and J. C. Stein, 2000, "Bad News Travels Slowly: Size, Analysts Coverage, and Profitability of Momentum Strategies," *Journal of Finance*, 55, 265–295.
- Hong, H., and J. Stein, 1999, "A Unified Theory of Underreaction, Momentum Trading and Overreaction in Asset Markets," *Journal of Finance*, 54, 2143–2184.
- Irvine, P. J., 2004, "Analysts' Forecasts and Brokerage-Firm Trading," *Accounting Review*, 79, 125–149.
- Kahneman, D., P. Slovic, and A. Tversky, 1982, *Judgment under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge, UK.
- Lim, T., 2001, "Rationality and Analysts' Forecast Bias," *Journal of Finance*, 56, 369–385.
- Lin, H.-W., and M. F. McNichols, 1998, "Underwriting Relationships, Analysts' Earnings Forecasts and Investment Recommendations," *Journal of Accounting and Economics*, 25, 101–127.
- Liu, H. M., and T. Yao, 2002, "The Consensus-Beating Game," Boston College and University of Arizona Working Paper.
- Liu, X., 2003, "Analysts' Response to Earnings Management," Working paper, Northwestern University.
- Lys, T., and S. Sohn, 1990, "The Association between Revisions of Financial Analysts' Earnings Forecasts and Security-Price Changes," *Journal of Accounting and Economics*, 13, 341–363.
- McNichols, M., and P. C. O'Brien, 1997, "Self-Selection and Analyst Coverage," *Journal of Accounting Research, Supplement*, 35, 167–199.
- Michaely, R., and K. L. Womack, 1999, "Conflict of Interest and the Credibility of Underwriter Analyst Recommendations," *Review of Financial Studies, Special*, 12, 653–686.
- Mikhail, M. B., B. R. Walther, and R. H. Willis, 1999, "Does Forecast Accuracy Matter to Security Analysts?," *Accounting Review*, 74, 185–200.
- Park, C., and E. Stice, 2000, "Analyst Forecasting Ability and the Stock Price Reaction to Forecast Revisions," *Review of Accounting Studies*, 5, 259–272.

- Prendergast, C., and L. Stole, 1996, "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, 104, 1105–1134.
- Rajan, R., and H. Servaes, 1997, "Analyst Following of Initial Public Offerings," *Journal of Finance*, 52, 507–529.
- Richardson, S., S. H. Teoh, and P. Wysocki, 2001, "Tracking Analysts' Forecasts over the Annual Earnings Horizon: Are Analysts' Forecasts Optimistic or Pessimistic?," Working paper, University of Pennsylvania, Ohio State University and MIT.
- Stickel, S. E., 1992, "Reputation and Performance among Security Analysts," *Journal of Finance*, 48, 1811–1836.
- Welch, I., 2000, "Herding among Security Analysts," *Journal of Financial Economics*, 58, 369–396.
- Wooldridge, J., 2003, "Cluster-Sample Methods in Applied Econometrics," *American Economic Review*, 93.
- Zitzewitz, E., 2001a, "Measuring Herding and Exaggeration by Equity Analysts and Other Opinion Sellers," Stanford University GSB Working Paper.
- Zitzewitz, E., 2001b, "Opinion-Producing Agents: Career Concerns and Exaggeration," Stanford University GSB Working Paper.

Copyright of Review of Financial Studies is the property of Society for Financial Studies and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.