THE JOURNAL OF BEHAVIORAL FINANCE, 9: 53–61, 2008

ISSN: 1542-7560 print / 1542-7579 online
DOI: 10.1080/15427560802093639

Biases in Individual Forecasts: Experimental Evidence

Lucy F. Ackert
Michael J. Coles College of Business, Kennesaw State University, Kennesaw, GA

Bryan K. Church
Georgia Tech, Atlanta, GA

Kirsten Ely
Sonoma State University, Rohnert Park, CA

Trueman [1994] provides a model of forecasting behavior in which analysts do not always make forecasts that are consistent with their private information. Using Trueman's model to provide theoretical direction, we conduct six experimental sessions to investigate individual forecasting behavior. In each session, four individuals predict earnings based on possibly divergent information. We manipulate forecast ability so that two individuals are strong analysts and two are weak. In three sessions, forecasts are released simultaneously. We find that forecasts do not always reflect private information. Both weak and strong analysts make forecasts that are inconsistent with private information, although the behavior is much more pronounced for weak analysts. In another three sessions, forecasts are released sequentially. We find that weak second analysts engage in herd behavior: that is, they mimic the reporting behavior of the first analyst. In contrast, strong second analysts are unaffected by the reporting behavior of the first analyst. The overall findings are consistent, in spirit, with the forecasting behavior suggested by Trueman [1994].

Keywords: Forecasts, Financial analysts, Private information, Herd behavior

Without question, capital markets are influenced by forecasts and, in particular, earnings forecasts. Managers make forecasts as part of budgeting and strategic planning, analysts turn out forecasts that provide benchmarks for performance, and individuals produce forecasts in considering investment strategies. Hence, forecast realizations (i.e., the relation between forecasts and actual outcomes) have economic consequences, be it implications for performance evaluation, reputation, or social welfare.

Research in economics and finance traditionally assumes that forecasts reflect individuals' private information in an unbiased manner. Yet extant evidence suggests that, under certain circumstances, individuals' forecasts may be biased. Trueman [1994] derives a model in which weak analysts (i.e., those with lower ability) release forecasts that do not fully reflect their private information. He shows that when all analysts (weak and strong) release forecasts simultaneously, the forecasts of weak analysts are less extreme than their private information warrants (i.e., the forecasts are too close to prior expectations). When analysts release forecasts sequentially, weak analysts disregard their private information and mimic the behavior of others (i.e., they follow the behavior of those perceived to have strong ability). In both cases (simultaneous and sequential forecast release), weak analysts suppress private information to favorably affect investors' assessment of their forecasting ability. Weak analysts want their clients to perceive them as strong.

While Trueman's work is cast in terms of analysts, it can be applied to all forecasters. An important implication is that market outcomes may be affected if individuals' forecasts do not fully reflect private information. Thus, researchers

Address correspondence to Lucy F. Ackert, Professor of Finance, Michael J. Coles College of Business, Kennesaw State University, Kennesaw, GA 30144. Tel: 770-423-6111. E-mail: lackert@kennesaw.edu

must gain insight into forecasting behavior to properly understand the functioning of financial markets. In this paper we empirically investigate whether individuals with different forecasting abilities produce forecasts that are consistent with the predictions of Trueman's model.

Our experimental examination of forecasting behavior is in a group setting that is analogous to one in which individuals produce forecasts of earnings per share based on possibly divergent private information. We examine forecasting behavior across two types of forecast abilities (weak and strong) and in two forecast settings (simultaneous and sequential release). We define forecast ability by the accuracy of the signal provided to the participant. Individuals receiving the more accurate signal are strong forecasters; those receiving the less accurate signal are weak forecasters. This operationalization is consistent with a variety of scenarios. Strong forecasters could be those with greater access to management, more industry expertise, or a greater ability to process information. Alternatively, differential forecasting ability could result from limited attention, where some investors attend more to the information required to forecast earnings than others (Hirshleifer and Teoh [2003]). For our purposes, the reason a forecaster might be strong or weak is unimportant because ability is exogenously determined by the accuracy of the signal provided to the participant.

In our first set of experimental sessions, weak and strong forecasters release forecasts simultaneously. We use the information given to participants to examine whether individuals issue forecasts that are consistent with their private information and whether this varies across weak and strong forecasters. In the second set of sessions, forecasts are released sequentially such that one individual publicly announces a reporting decision before other group members make their predictions. The ability of the forecaster who discloses first varies from period to period. This approach permits us to investigate whether the first analyst's reporting decision affects the forecast behavior of the remaining analysts.

The results of the first set of experimental sessions indicate that individuals' forecasts are not always unbiased reflections of their private information. Weak forecasters often release forecasts that are less extreme than their private information warrants. Strong forecasters disregard their private information on occasion, though much less frequently than weak forecasters. The results of the second set of experimental sessions indicate that weak forecasters regularly release forecasts that are similar to those previously announced by others. The first forecaster's reporting decision conveys private information to weak forecasters, which affects their reporting behavior. Strong forecasters, on the other hand, are not affected by the first forecaster's reporting decision.

The remainder of this paper is organized as follows. First we provide a framework for the study and develop the research hypotheses. Next we describe the experimental method and then present the experimental results. The final section provides concluding remarks.

FRAMEWORK

At the microeconomic level, individual behavior has received significant attention because it provides insight into the judgments and decisions of managers, analysts, and regulators. At the macroeconomic level, modern asset-pricing theory assumes that prices reflect available information so that individual psychology has no influence on stock prices (e.g., Hirshleifer [2001]). Traditional asset-pricing theory, however, cannot explain market anomalies such as Black Monday or the bursting of the dot-com bubble. Recent theories, developed to allow for market anomalies, suggest that individual behavior can influence market outcomes. As a result, individual behavior has become important at the macro level, too.

Trueman [1994] provides a model of individual behavior in which it can be rational to release forecasts that do not reflect private information in an unbiased manner. The model is also consistent with several other anomalies found in the capital-markets literature, including non-linearities in the relation between stock returns and earnings surprises and consistently optimistic forecasts. If Trueman's model is a valid representation of the way people behave, it can provide insight into what has previously been considered anomalous market behavior. Our experiment is designed to determine if Trueman's model provides a reasonable description of individual behavior in a forecasting task.

In Trueman's model, an analyst's forecast depends on private information (a signal) and ability (weak or strong). Earnings are signed (positive or negative) and take one of two levels (high or low). Extreme earnings are uncommon, such that high earnings are less likely than low earnings, regardless of the sign. The analyst knows the prior distribution of earnings realizations. Further, the analyst is able to gauge the sign of earnings with certainty, but not the level.

The analyst receives a signal of the earnings level that indicates high or low earnings and then makes a forecast. The analyst knows his or her own ability and, more specifically, the likelihood that the earnings level can be forecasted correctly. A strong analyst has a higher likelihood of making a correct forecast than a weak analyst because the strong analyst's signal is more accurate. Others cannot observe an analyst's ability but make inferences by comparing the accuracy of the analyst's forecast to that of other forecasters. Accordingly, analysts release forecasts to maximize the probability that others think the analyst's ability is strong.

Trueman's model uses Bayes' rule to predict analyst behavior under two settings: simultaneous and sequential release of forecasts. In the simultaneous setting, analysts announce forecasts concurrently. Analysts use knowledge of their own ability and private information to determine a posterior probability of high or low earnings and then choose a forecast (i.e., high or low earnings). Because weak analysts have less accurate private information than strong analysts, they may be less inclined to forecast high (extreme) earnings. According to Trueman, a weak analyst gains more, on average, than a strong analyst by reporting low earnings when the information signal is actually high.[1] If a forecast based on private information increases the probability of revealing that the weak analyst has lower ability, an analyst has the incentive to release less extreme forecasts than private information warrants. This leads us to our first set of hypotheses.

> H1a: Weak analysts sometimes report forecasts that do not reflect their private information.
>
> H1b: Strong analysts report forecasts that reflect their private information.

In the case of sequential forecast release, a weak analyst who receives a high signal, again, sometimes reports a forecast that is less extreme than suggested by private information. A weak analyst who reports after another analyst has incentives to mimic the previously released forecast if this forecast is made by a strong analyst. The first forecast allows a weak analyst to update the probability of each earnings level. Thus, a weak second analyst is prone to herding behavior when private information indicates a high earnings level. If private information indicates that earnings are low, the weak second analyst's forecast announcement is not affected by the first analyst. In addition, the forecasts of both the first analyst to report and strong second analysts do not change relative to the simultaneous release scenario. The first analyst to forecast faces the same set of circumstances as in the simultaneous release case and, therefore, has no reason to behave differently. Strong second analysts in Trueman's model receive no additional useful information from the first forecast because, even if the first analyst to release is strong, all strong analysts receive the same signal giving them no reason to revise their forecast. This leads to our second set of hypotheses.

> H2a: Weak second analysts sometimes release forecasts that are affected by previously announced forecasts.
>
> H2b: Strong second analysts release forecasts that are unaffected by previously announced forecasts.

RESEARCH METHOD

Participants

We conduct 6 experimental sessions, each administered to individuals in groups of 4, so that a total of 24 individuals take part in the study. Three sessions include the simultaneous release of forecasts and three include the sequential release of forecasts. Participants are recruited from third and fourth year undergraduate and fifth year post-baccalaureate students in business and economics at a medium-sized Canadian university. Students earn $14.50–25.00, with an average of $18.80, for participating approximately 75 minutes.

TABLE 1
Earnings Levels and Associated Probabilities.

Earnings level	Sign	Magnitude	Probability
1	Negative	High	0.2
2	Negative	Low	0.3
3	Positive	Low	0.3
4	Positive	High	0.2

Procedures

Each experimental session consists of 48 periods. Each period group members forecast the earnings level of a firm. Participants are informed that earnings may be either positive or negative and either high or low. They also are informed of the distribution of earnings levels. The earnings levels and associated probabilities are shown in Table 1.

At the beginning of each period, the experimenter places an index card labeled "S" for signal in front of each group member. The participant's signal for the period is recorded on the reverse side of the card. Although the signal always accurately reflects the sign of earnings, it does not always accurately reflect the magnitude. To manipulate forecast ability, we vary signal accuracy between participants. Strong analysts receive a signal that is on average accurate 80% of the time, while weak analysts receive a signal that is on average accurate 55% of the time. Participants' signal accuracy remains constant over the entirety of the experimental session. Participants know their own signal accuracy, but not that of others: They know that signal accuracy can vary across group members. In each session, two participants are strong analysts and two are weak.

After each group member receives a signal card, participants are permitted to look at their signal. Each group member subsequently makes a forecast of earnings. In sessions 1-3, participants make forecasts simultaneously. Participants are given 30 seconds to forecast the earnings level for the period. Participants do so by placing another card facedown in front of them. At the beginning of each session, all group members are given a set of four forecast cards, one for each earnings level, labeled "F"for forecast. After all group members make their predictions, the experimenter publicly announces each member's forecast. The observed earnings level for the period is then announced. If a participant's forecast matches the observed earnings level, the participant accumulates 50 points. Otherwise, a participant does not accumulate any points.

In sessions 4–6, participants make forecasts sequentially. After each group member receives a signal card, participants are permitted to look at their signals. One group member (the first analyst) is given 30 seconds to place a forecast card in front of him/her, representing that group member's prediction for the period. The role of the first analyst alternates among the participants so that each is first every fourth period. The experimenter publicly announces the prediction of the first analyst and the remaining three group members are given 30 seconds to make a forecast. After the remaining group members make their predictions, the experimenter announces each of their forecasts. Firm earnings are then revealed and participants accumulate their points. As before, participants accumulate 50 points if their forecast matches the observed earnings level and zero otherwise.

At the end of each session, points are converted to dollars at a specified rate and each participant is paid in cash. During this time, participants complete a post-experiment questionnaire designed to collect general information about the participants and how they view the experimental session.

RESULTS

Do Individuals' Forecasts Fully Reflect Their Private Information?

The first set of hypotheses predicts that weak analysts sometimes report forecasts that are not consistent with their private information (H1a), whereas strong analysts report forecasts that are consistent with their private information (H1b). To assess the hypotheses, we focus on the data collected in the simultaneous forecast setting (sessions 1–3). Below, we examine the forecasting behavior of weak analysts and then the forecasting behavior of strong analysts.

Weak Analysts.

Table 2 presents the observed frequencies of forecasts reported by weak analysts conditioned on the signal received. We present the frequencies overall for sessions 1–3 (periods 1–48 in Panel A) and then splitting the data into halves (periods 1–24 in Panel B and periods 25–48 in Panel C). Looking at periods 1–48, weak analysts report forecasts that are not consistent with their private information 36.5% of the time (105 of 288, which represents the sum of the off-diagonal cells in Panel A of Table 2).

The data also suggest that weak analysts report forecasts that are not consistent with their private information more often when the earnings signal is high (level 1 or 4) as opposed to low (level 2 or 3), which is in line with Trueman [1994]. The forecast does not match the signal 48.3% of the time with a high signal (58 of 120) versus 28.0% with a low signal (47 of 168). A chi-square test indicates that the difference is statistically significant ($\chi^2 = 12.52$, $p < 0.001$). Splitting the data into halves (i.e., periods 1–24 and 25–48), the frequencies are similar.[2] The findings are consistent with H1a.

TABLE 2
Weak Analysts' Forecast by Signal Frequencies.

	Signal				
Forecast	1	2	3	4	Total
		Panel A: Periods 1–48			
1	18	24	—	—	42
2	36	66	—	—	102
3	—	—	55	22	77
4	—	—	23	44	67
Total	54	90	78	66	288
		Panel B: Periods 1–24			
1	11	8	—	—	19
2	25	22	—	—	47
3	—	—	26	15	41
4	—	—	10	27	37
Total	36	30	36	42	144
		Panel C: Periods 25–48			
1	7	16	—	—	23
2	11	44	—	—	55
3	—	—	29	7	36
4	—	—	13	17	30
Total	18	60	42	24	144

Notes: The cell entries indicate the frequency that signal by forecast combinations are observed in the simultaneous forecast setting for weak analysts.

Next, we examine the reporting behavior of individual forecasters. We compute the percentage of time that the forecast does not match the signal for each forecaster over periods 1–48. The percentages are 52.1, 52.1 41.7, 35.4, 20.8, and 16.7.[3] Accordingly, four weak analysts release forecasts that are inconsistent with their private information at least one-third of the time.

We also investigate whether the frequency that weak analysts' forecasts reflect their private information is in line with the updated Bayesian probability. As mentioned earlier, Trueman relies on Bayesian updating as a basis for *why* weak analysts sometimes report forecasts that do not match their signal. We assess whether the percentage of time that weak analysts' forecasts do not match their signal differs from the updated probability that the analyst forecasts correctly if the forecast does not match the signal. Formally, we test the following

$$P(F \neq S) = P(E \neq S|S)$$

where F is the forecast, S is the signal, and E is the earnings realization. We perform binomial tests to compare the observed frequencies with the Bayesian posterior, conditioned on the signal being low or high. In conducting the tests, we use the normal approximation to the binomial distribution (Mendenhall, Scheaffer, and Wackerly [1981]). We perform tests using data from periods 1–48 as well as periods 1–24 and 25–48. The results are reported in Table 3. In five of six tests, we are unable to reject the null that the probabilities do not differ at the 5% level. In the one instance where we are

TABLE 3
Weak Analysts' Forecasting Behavior and Bayesian Posteriors.

Signal (S)	Periods	P(F≠S)	P(E≠S|S)	z-statistic
S = 2, 3 (Low)	1–48	0.280	0.353	1.985**
	1–24	0.273		1.360
	25–48	0.284		1.458
S = 1, 4 (High)	1–48	0.483	0.551	1.498
	1–24	0.513		0.675
	25–48	0.429		1.590

Notes: For notation, S is the signal, F is the forecast, and E is the earnings realization. The column headed P(F ≠ S) is the observed frequency that the forecast does not match the signal. The column headed P(E ≠ S|S) is the probability that the earnings realization does not match the signal, conditioned on the signal. It is the Bayesian posterior, which incorporates the analyst's ability (i.e., the weak analyst's signal is accurate 55% of the time). For each row, we conduct a binomial test using the normal approximation to determine whether P(F ≠ S) = P(E ≠ S | S). Two asterisks denote that the null hypothesis can be rejected at the 5% level (two-tailed test).

able to reject the null, weak analysts release forecasts that match their signal more often than predicted by the Bayesian posterior.

Strong Analysts.

Table 4 presents the observed frequencies of forecasts reported by strong analysts conditioned on the signal received. We present the frequencies for periods 1–48 (in Panel A) and separately for periods 1–24 and 25–48 (in Panels B and C, respectively). Over periods 1–48, strong analysts report forecasts that reflect their private information 81.2% of the time (234 of 288). They appear more likely to report consistent forecasts when the earnings signal is low rather than high. The forecast matches the signal 87.5% of the time with a low signal (147 of 168) versus 72.5% of the time with a high signal (87 of 120). The difference is statistically significant ($\chi^2 = 11.34$, $p = 0.001$). The result is not entirely consistent with H1b, which suggests that strong analysts release forecasts that reflect their private information, regardless of whether the signal is low or high earnings.

Splitting the data into halves, we gain additional insight. Strong analysts release a forecast that matches their signal 75.7% of the time over the first half of the sessions (109 of 144) versus 86.8% over the second half (125 of 244). The difference is statistically significant ($\chi^2 = 5.84$, $p = 0.016$). Further examination of the data indicates that the difference is driven by forecasting behavior when the signal is high. Strong analysts release forecasts that are consistent with a high signal 64.1% of the time over periods 1–24 (50 of 78) versus 88.1% over periods 25–48 (37 of 42). The difference is statistically significant ($\chi^2 = 7.88$, $p = 0.005$). Thus, as time progresses, the forecast reflects the signal with similar frequency regardless of whether the signal is low or high. The results over periods 25–48 are more in line with H1b, though on occasion strong analysts still report forecasts that do not match their private information.

TABLE 4
Strong Analysts' Forecast by Signal Frequencies.

	Signal				
Forecast	1	2	3	4	Total
		Panel A: Periods 1–48			
1	43	11	—	—	54
2	11	79	—	1	91
3	—	1	68	21	90
4	—	—	9	44	53
Total	54	91	77	66	288
		Panel B: Periods 1–24			
1	21	4	—	—	25
2	9	32	—	—	41
3	—	—	27	19	46
4	—	—	3	29	32
Total	30	36	30	48	144
		Panel C: Periods 25–48			
1	22	7	—	—	29
2	2	47	—	1	50
3	—	1	41	2	44
4	—	—	6	15	21
Total	24	55	47	18	144

Notes: The cell entries indicate the frequency that signal by forecast combinations are observed in the simultaneous forecast setting for strong analysts. The cell entries indicate two mistakes in making forecast. In one period, a participant reported a forecast of 3 after receiving a signal of 2. In another period, a participant reported a forecast of 2 after receiving a signal of 4. In both cases, the participant inadvertently placed the wrong forecast card in front of him/her. The mistakes occurred in different sessions and, as such, involved different participants.

We also assess the forecasting behavior of individual forecasters. As before, we compute the percentage of time that the forecast is the same as the signal for each forecaster over periods 1–48. The percentages are 58.3, 68.8, 79.2, 89.6, 95.8, and 97.9.[4] Hence, three strong analysts seldom release forecasts that differ from their private information (five times or less over 48 periods). The other three strong analysts release inconsistent forecasts more frequently but, overall, they release forecasts that match their signal a majority of the time.

Next, we compare the frequency that strong analysts' forecasts reflect their private information with the Bayesian posterior. We conduct binomial tests using data from periods 1–48, 1–24, and 25–48, conditioned on the signal being low or high. The results are shown in Table 5. In four of six tests, we are unable to reject the null that the probabilities do not differ at the 10 percent level. In one instance, we are able to reject the null at the 5% level and in another instance at the 10% level. In both cases, the rejections occur when the signal is high: the observed frequency is below the Bayesian posterior in periods 1–24 and above it in periods 25–48.

TABLE 5
Strong Analysts' Forecasting Behavior and Bayesian Posteriors

Signal (S)	Periods	P(F = S)	P(E = S| S)	z-statistic
S = 2, 3 (Low)	1–48	0.875	0.857	0.666
	1–24	0.894		0.859
	25–48	0.863		0.173
S = 1, 4 (High)	1–48	0.725	0.727	−0.049
	1-24	0.641		−1.705*
	25-48	0.881		2.240**

Notes: For notation, S is the signal, F is the forecast, and E is the earnings realization. The column headed P(F = S) is the observed frequency that the forecast matches the signal. The column headed P(E = S|S) is the probability that the earnings realization matches the signal, conditioned on the signal. It is the Bayesian posterior, which incorporates the analyst's ability (i.e., the weak analyst's signal is accurate 55% of the time). For each row, we conduct a binomial test using the normal approximation to determine whether P(F = S) = P(E = S|S). Two asterisks denote that the null hypothesis can be rejected at the 5% level and one asterisk denotes at the 10% level (two-tailed test).

Do Individuals' Follow the Reporting Behavior of Others?

The second set of hypotheses indicates that weak second analysts sometimes follow the reporting behavior of the first analyst (H2a), whereas strong second analysts are unaffected by the reporting behavior of the first analyst (H2b). To assess the hypotheses, we focus on the data collected in the sequential forecast setting (sessions 4–6). Once again, we examine the forecasting behavior of weak analysts and then strong analysts.

Weak Analysts.

We tabulate the frequencies that weak second analysts' forecasts ($F_{2,w}$) are the same and different from the first analyst's forecast (F_1). With sequential reporting of forecasts, additional notation is introduced to differentiate the forecast order as well as the analyst's ability. The subscript indicates the order of forecast release and analyst ability where w (s) indicates weak (strong) ability, when necessary. If no superscript appears, ability needs no differentiation. Here we are interested in the number of times that the forecasts are the same ($F_{2,w} = F_1$). However, a complication arises with the data. When the forecasts are the same, it is difficult to determine whether the second analyst is mimicking the behavior of the first analyst or reacting to private information. For example, when the second analyst's signal matches the first analyst's forecast, a second analyst who releases an identical forecast may be reacting to private information rather than the behavior of others. Therefore, we partition the frequency counts by whether the second analyst's signal ($S_{2,w}$) matches the first analyst's forecast (F_1). We can say with greater confidence that the second analyst follows the reporting behavior of the first analyst when $F_{2,w} = F_1 | S_{2,w} \neq F_1$. The summary data are presented in Table 6 for periods 1–48 (Panel A) and separately for periods 1–24 and 25–48 (Panels B and C, respectively).

TABLE 6
Weak Second Analysts' Forecasts by Signal Match and First Analyst's Forecast

First Analyst's	$S_2 = F_1$		$S_2 \neq F_1$	
Forecast (A)	$F_2 = F_1$	$F_2 \neq F_1$	$F_2 = F_1$	$F_2 \neq F_1$
Panel A: Periods 1–48				
1	20	3	10	6
2	51	3	9	4
3	47	1	12	11
4	29	3	4	3
Total	147	10	35	24
Panel B: Periods 1–24				
1	12	3	3	4
2	15	3	8	4
3	22	0	7	3
4	21	2	0	1
Total	70	8	18	12
Panel C: Periods 25–48				
1	8	0	7	2
2	36	0	1	0
3	25	1	5	8
4	8	1	4	2
Total	77	2	17	12

Notes: For notation, F_1 is the first analyst's announced forecast, S_2 is the second analyst's signal, and F_2 is the second analyst's forecast.

The data indicate that when weak second analysts' signal matches the first analyst's forecast ($S_{2w} = F_1$), the second analyst's forecast is the same as that of the first analyst 93.6% of the time over periods 1–48 (147 of 157), which is not particularly surprising. When weak second analysts' signal does not match the first analyst's forecast ($S_{2w} \neq F_1$), the second analyst's forecast is the same as that of the first analyst 59.3% of the time over periods 1–48 (35 of 59). Though the frequency decreases when $S_{2w} \neq F_1$, weak second analysts still follow the lead of the first analyst a majority of the time. Splitting the data into halves, the frequencies are similar.

To formally test whether weak second analysts follow the reporting behavior of the first analyst, we focus on instances in which $S_{2,w} \neq F_1$. Under this condition, we compute the probability that the second analyst's forecast is the same as the first analyst's forecast or $P(F_{2,} = F_1 | S_{2,w} \neq F_1)$. To provide a benchmark for comparison, we compute the probability that weak analysts release a forecast that is not consistent with their private information, absent other information (i.e., absent a first forecast). In the sequential forecast setting, one analyst goes first each period. We isolate the reporting behavior of weak first analysts and determine the frequency that the forecast does not coincide with private information (i.e., $F_{1,w} \neq S_{1,w}$).

According to H2a, weak second analysts follow the reporting behavior of the first analyst. We test whether $P(F_{2,w} =$

TABLE 7
Do Weak Second Analysis' Follow the Forecast of the First Analyst?

Periods	$P(F_{2,w} = F_1 \mid S_{2,w} \neq F_1)$	$P(F_{1,w} \neq S_{1,w})$	z-statistic
1–48	0.593	0.208	7.290***
1–24	0.600	0.306	3.494***
25–48	0.586	.0.111	8.147***

Notes: For notation, $F_{2,w}$ is the weak second analyst's forecast, F_1 is the first analyst's forecast, $S_{2,w}$ is the weak second analyst's signal, F $F_{1,w}$ is the weak first analyst's forecast, and $S_{1,w}$ is the weak first analyst's signal. The column headed $P(F_{2,w} = F_1 \mid S_{2,w} \neq F_1)$ is the probability that the second analyst follows the first analyst's forecast and disregards private information. The column headed $P(F_{1,w} \neq S_{1,w})$ is the probability that weak weak analysts release a forecast that does not coincide with private information when weak analysts go first. For each row, we conduct a binomial test using the normal approximation to determine whether $P(F_{2,w} = F_1 \mid S_{2,w} \neq F_1) > P(F_{1,w} \neq S_{1,w})$, which is indicative of second analyst following the behavior of the first analyst.Three asterisks denote that the null hypothesis can be rejected at the 1% level (one-tailed test).

$F_1 \mid S_{2,w} \neq F_1) > P(F_{1,w} \neq S_{1,w})$, which is indicative of herd behavior. We perform binomial tests using the normal approximation to the binomial distribution (Mendenhall, Scheaffer, and Wackerly [1981]). The probability that weak analysts disregard private information absent other forecasts (i.e., $P(F_{1,w} \neq S_{1,w})$) serves as the baseline probability for comparison. We conduct tests using data from periods 1–48 and separately for periods 1–24 and 25–48. The results, reported in Table 7, indicate that in all cases the null can be rejected at the 1% level. The findings suggest that weak second analysts follow the reporting behavior of the first analyst and are supportive of H2a.

Strong Analysts

We tabulate the frequencies that strong second analysts' forecasts ($F_{2,s}$) are the same and different from the first analyst's forecast (F_1). As before, we partition the data by whether the second analysts' signal matches the first analysts' forecast. The summary data are presented in Table 8 for periods 1–48 (Panel A) and separately for periods 1–24 and 25–48 (Panels B and C, respectively).

The data indicate that when strong second analysts' signal matches the first analyst's forecast ($S_{2s} = F_1$), the second analyst's forecast is the same as that of the first analyst 98.9% of the time over periods 1–48 (172 of 174), which again is not surprising. When strong second analysts' signal does not match the first analyst's forecast ($S_{2s} \neq F_1$), the second analyst's forecast is the same as that of the first analyst only 4.8% of the time over periods 1–48 (2 of 42). Inferences are unchanged looking at periods 1–24 and 25–48 separately. Hence, the data provide little evidence that strong second analysts engage in herd behavior, which is consistent with H2b.

TABLE 8
Weak Second Analysts' Forecasts by Signal Match and First Analyst's Forecast

First Analyst's Forecast (A)	$S_2 = F_1$		$S_2 \neq F_1$	
	$F_2 = F_1$	$F_2 \neq F_1$	$F_2 = F_1$	$F_2 \neq F_1$
Panel A: Periods 1–48				
1	29	0	0	9
2	58	0	1	13
3	51	2	1	16
4	34	0	0	2
Total	172	2	2	40
Panel B: Periods 1–24				
1	15	0	0	4
2	22	0	1	5
3	21	1	1	14
4	24	0	0	0
Total	82	1	2	23
Panel C: Periods 25–48				
1	14	0	0	5
2	36	0	0	8
3	30	1	0	2
4	10	0	0	2
Total	90	1	0	17

Notes: For notation, F_1 is the first analyst's announced forecast, S_2 is the second analyst's signal, and F_2 is the second analyst's forecast.

To formally test H2b, we assess whether $P(F_{2,s} = F_1 \mid S_{2,s} \neq F_1) = P(F_{1,s} \neq S_{1,s})$ for the strong second analysts. The term on the left hand side of the equality is the probability that second analysts follow the first analyst's reporting behavior, disregarding their private information. The term on the righthand side is the probability that analysts release forecast that do not coincide with their private information, absent other forecasters. As before, we perform binomial tests using the normal approximation of the binomial distribution. We take $P(F_{1,s} \neq S_{1,s})$ as the baseline for comparison and conduct

TABLE 9
Do Strong Second Analysts' Follow the Forecast of the First Analyst?

Periods	$P(F_{2,s} = F_1 \mid S_{2,s} \neq F_1)$	$P(F_{1,s} \neq S_{1,s})$	z-statistic
1–48	0.048	0.042	0.182
1–24	0.080	0.056	0.522
25–48	0.000	0.028	−0.700

Notes: For notation, $F_{2,s}$ is the strong second analyst's forecast, F_1 is the first analyst's forecast, $S_{2,s}$ is the strong second analyst's signal, $F_{1,s}$ is the strong first analyst's forecast, and $S_{1,s}$ is the strong first analyst's signal. The column headed $P(F_{2,s} = F_1 \mid S_{2,s} \neq F_1)$ is the probability that the second analyst follows the first analyst's forecast and disregards private information. The column headed $P(F_{1,s} \neq S_{1,s})$ is the probability that strong analysts release a forecast that does not coincide with private information when strong analysts go first. For each row, we conduct a binomial test using the normal approximation to determine whether $P(F_{2,s} = F_1 \mid S_{2,s} \neq F_1) > P(F_{1,s} \neq S_{1,s})$, which is indicative of second analyst following the behavior of the first analyst.

test for periods 1–48, 1–24, and 25–48. The results, shown in Table 9, indicate that in no case are we able to reject the null at conventional levels. The findings suggest that strong second analysts do not mimic the behavior of the first analyst and provide compelling evidence in support of H2b.

CONCLUSION

We conduct six experimental sessions to investigate individual forecasting behavior. In each session, four individuals predict the earnings of a hypothetical firm based on possibly divergent information. We manipulate forecast ability so that two individuals are strong analysts and two are weak analysts. In the first set of experimental sessions, all four group members release forecasts simultaneously. Our theoretical predictions suggest that weak analysts do not always release forecasts that are consistent with their private information. Strong analysts, on the other hand, are expected to release forecasts that coincide with their private information. In the second set of experimental sessions, forecasts are released sequentially. One individual publicly announces a forecast before the remaining group members make their predictions. The role of the first analyst varies from period to period. Our theoretical predictions suggest that weak second analysts follow the reporting behavior of the first analyst. By comparison, strong second analysts are not expected to engage in herd behavior.

The results from the simultaneous forecast setting indicate that individuals' forecasts are not always unbiased reflections of their private information. Both weak and strong analysts release forecasts that are not consistent with their private information, though the behavior is much more pronounced for weak analysts. Moreover, such behavior occurs much more frequently when private information indicates that the earnings level is extreme (i.e., individuals are inclined to report a low earnings level when their private information indicates a high earnings level). The results from the sequential forecast setting indicate that weak second analysts follow the reporting behavior of the first analyst. In contrast, strong second analysts are not affected by the reporting behavior of the first analyst. Across experimental sessions, the results are largely consistent with the theoretical predictions.

The results from the simultaneous forecast setting have implications for research on the market's reaction to earnings announcements. Many archival-based studies use analysts' forecasts as a proxy for market expectations. This proxy is used to compute the earnings surprise associated with an earnings announcement (i.e., the difference between analysts' forecasts and reported earnings). Yet investors likely realize that analysts' forecasts are not always unbiased reflections of their private information, and investors' earnings expectations may differ from those of analysts (Trueman [1994]). As a consequence, earnings surprises may be measured incorrectly, and the observed reaction to earnings announcements may contain error (see also Lev [1989]).

The results from the simultaneous forecast setting also have implications for research on the biasedness of analysts' forecasts. Several studies show that, on average, analysts' forecasts are optimistic. One explanation for this finding is that analysts are able to generate trading activity by releasing optimistic forecasts, which indirectly impacts their compensation (i.e., they have incentives to release biased forecasts). In our experimental setting, however, participants are compensated based on their forecast accuracy. Yet, they still release biased forecasts.

The results from the sequential forecast setting provide insight into *why* individuals engage in herd behavior. Previous theoretical studies suggest that individuals follow the behavior of others because others' actions convey private information and in order to positively affect assessments of ability (e.g., Banerjee [1992]; Trueman [1994]). Our findings are consistent with this assertion. Other theoretical studies suggest that individuals have different incentives to engage in herd behavior. Scharfstein and Stein [1990, p 466] contend that managers avoid contrarian behavior when mimicking the behavior of other managers. Herd behavior results from attempts to manipulate the labor market's inferences regarding managerial ability (i.e., aptitude for decision making). Managers may be evaluated more favorably if they follow the decisions of others than if they behave in a contrarian manner. A similar sentiment was echoed by John Maynard Keynes who in *The General Theory* [1936, p 158] goes so far as to argue "that it is better for reputation to fail conventionally than to succeed unconventionally." Future research is necessary to investigate whether other incentives drive herd behavior.

ACKNOWLEDGEMENT

The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. The authors gratefully acknowledge the financial support of the Federal Reserve Bank of Atlanta and the Certified General Accountants of Canada and the research assistance of Karen Butler, Dan Li, Betty Soares, and Shawna White.

NOTES

1. Refer to observations 1 and 2 in Trueman [1994, p. 105]. According to proposition 2 (p. 106), the probability that a weak analyst reports low earnings when the information signal indicates high earnings is strictly positive as long as the analysts' ability is less than two times the probability of low earnings. This restriction holds in our experimental setup.
2. The only discernible difference is that with a high earnings signal, the forecast does not match the signal more often over periods 1–24 than periods 25–48 (51.3%

versus 42.9%). The difference, however, is not statistically significant ($\chi^2 = 0.78$, $p = 0.378$).

3. Splitting the data into halves, the percentages are similar for five of the six forecasters. One weak analyst releases inconsistent forecasts 41.7% of the time over periods 1–24, and 0% of the time over periods 25-48.
4. Splitting the data into halves, the percentages are similar for five of the six forecasters. One strong analyst releases consistent forecasts 45.8% of the time over periods 1–24, and 91.6% of the time over periods 25–48.

REFERENCES

Banerjee, A.V. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107, (1992), pp. 797–817.

Hirshleifer, D. "Investor Psychology and Asset Pricing." *The Journal of Finance*, 51, (2001), pp. 1533–1597.

Hirschleifer, D. and S.H. Teoh. "Limited Attention, Information Disclosure, and Financial Reporting." *Journal of Accounting and Economics*, 36(1–3), (2003), pp. 337–386.

Keynes, J.M. *The General Theory of Employment, Interest, and Money*. First Harbinger Edition, 1964. New York: Harcourt, Brace, Jovanovich, 1936.

Lev, B. "On the Usefulness of Earnings and Earnings Research: Lessons and Directions from Two Decades of Empirical Research." *Journal of Accounting Research*, 27, (1989), pp. 153–192.

Mendenhall, W., R.L. Scheaffer and D.D. Wackerly. *Mathematical Statistics with Applications*, Second Edition. Boston: Duxbury Press, 1981.

Scharfstein, D.S. and J.C. Stein. "Herd Behavior and Investment." *American Economic Review*, 80, (1990), pp. 465–479.

Trueman, B. "Analyst Forecasts and Herding Behavior." *Review of Financial Studies*, 7, (1994), pp. 97–124.