

A Note on Analysts' Earnings Forecast Errors Distribution*

Daniel A. Cohen

and

Thomas Z. Lys**

Kellogg School of Management
Northwestern University
Evanston, Illinois 60208

Current Draft: November 20, 2003

Abstract

Abarbanell and Lehavy provide evidence that analysts' forecast errors are not normally distributed exhibiting a high occurrence of extreme negative forecast errors (left-tail asymmetry) and a high occurrence of small positive forecast errors (middle asymmetry). This is important for researchers who rely on techniques that are sensitive to the distributional assumptions of analysts' forecast errors. Many of the conclusions drawn by Abarbanell and Lehavy, however, are based on visual impressions (as opposed to formal empirical tests) or based on methods that are very sensitive to the empirical methods used (e.g., whether the serial correlation of forecast errors is caused by the left-tail asymmetry).

JEL classification: G10, G29, M41

Keywords: Analysts' forecasts, analysts' bias, analysts' under/overreaction to information; analysts' loss function; discretionary accruals.

* We would like to thank Aiysha Dey, Thomas Fields, Robert Korajczyk, David Larcker, Margaret Neale, Georgios Skoulakis, Shyam Sunder, Linda Vincent, Ross Watts (the editor), and Jerry Zimmerman for helpful comments on a previous draft. All remaining errors are our own responsibility.

** Corresponding Author (847) 491-2673, tlys@kellogg.northwestern.edu

1. Introduction

The purpose of Abarbanell and Lehavy's (henceforth, A&L) paper is to provide insights into the vast and, at times, conflicting research on the properties of analysts' earnings forecasts. The main point of their analysis is that many of the seemingly anomalous and conflicting results originate in researchers' failure to incorporate the appropriate distributional properties of analysts' forecast errors. Because those observations fall into two distinct regions of the distribution of forecast errors – the center and the tails – A&L refer to them as the middle and the tail asymmetries.

A&L show that scaled forecast errors are not normally distributed. This is an important observation as many researchers rely on statistical techniques that are sensitive to the distributional properties of the underlying variable. While the non-normal distribution of forecast errors is the central point of A&L's analysis, they note in their footnote 7 that there are no theoretical reasons to expect forecast errors to be normally distributed.

We interpret A&L's message as (1) that many conclusions drawn by researchers may be quite sensitive to the use of analyses that either explicitly or implicitly rely on specific distributional assumptions, and (2) that the observations in the two asymmetries may unduly affect the conclusions reached in the literature. In the spirit of A&L, recent research on analysts' loss function (Basu and Markov, 2003; Gu and Wu, 2003; and Liu, 2003) investigates the determinants of the distribution of analysts' forecast errors. That research shows that at least some of the previous results, such as analyst forecast bias and forecast inefficiencies, are sensitive to the distributional assumptions, thus supporting A&L's first point.

The conclusions reached by A&L with regards to the second point are more controversial, mainly because of their use of unconventional data analysis techniques. Based on their empirical analysis, A&L conclude that the inefficiency documented in the literature is caused by the ‘tail asymmetry,’ a small subset of extreme negative forecast errors. This implies, that analysts’ earnings forecasts are not nearly as inefficient as previous researchers have suggested. While this finding questions a substantial body of research, in the remainder of this discussion we raise points that may influence the reader’s confidence in this result. For example, our tests question A&L’s conclusion that the serial correlation in consecutive analysts’ forecast errors is caused by the tail asymmetry. Further, when we investigate whether these asymmetries are driven by the distribution of earnings or the distribution of analysts’ forecasts, we find that both asymmetries are present in the two components of forecast errors, albeit in attenuated form.

This discussion proceeds as follows: Section 2 discusses the analysis of the distributional properties of analyst forecasts. Section 3 focuses on analysts’ bias whereas analysts’ inefficiency is discussed in Section 4. Section 5 discusses the link between analysts’ forecast errors and unexpected accruals. Conclusions and suggestions for future research are presented in Section 6.

2. Distributional Properties of Analyst Forecast Errors

A&L’s analysis relies on price deflated consensus forecast errors ($CFE_{j,q}$), defined as:

$$CFE_{j,q} = \frac{ZAEPS_{j,q} - ZCF_{j,q}}{P_{j,q-1}} \times 100 \quad (1)$$

where for firm j and quarter q , $ZAEPS_{j,q}$ represents the actual earnings per share as reported by Zacks Investment Research (ZIR), $ZCF_{j,q}$ represents the ZIR consensus forecast outstanding just prior to the disclosure of the quarterly earnings, and $P_{j,q-1}$ represents the stock price at the beginning of quarter q .

Before discussing their results and interpretations, it is important to note that A&L subjected their sample to several screens that could affect the conclusions they reach. First, observations with data unavailable on COMPUSTAT are deleted. While this is unavoidable for some of the empirical evidence documented, it is not necessary for most of the analyses performed by A&L. In fairness, A&L indicate, that their results are not affected by this data selection. Second, the sample is winsorized based on values of $CFE_{j,q}$, the dependent variable in most of their analyses. This procedure selects observations for analysis based on their ex-post realization. Although this appears to be common in accounting research, it makes interpretation of the results difficult for two reasons. First, it may result in an understatement of the standard errors in subsequent analyses, thus overstating the levels of confidence when hypotheses are rejected. Second, it may induce associations that would not otherwise exist.¹ Both of these consequences are problematic when deriving the distributional properties of the dependent variable, especially when one of the main focuses of the analysis is to investigate the impact of extreme observations on prior results. Admittedly, some of the observations that are deleted in the winsorization are the result of data errors (e.g., Zacks enters a code of 999 for instances when the forecasts are not available). Another source of errors is the small denominator problem (e.g., when stock prices are near zero). For these data points, there

¹ In a different context, Kothari, Sabino, and Zach (2003) show that even small deletion of extreme observations can bias the results. Specifically, using data simulated to represent correct pricing (no market inefficiencies) they demonstrate that by deleting the top and bottom 1 to 2 percent of the observations, researchers are likely to wrongly reject market efficiency.

is no need to winsorize the dependent variable. Rather, it would be sufficient to delete observations with 999 codes, or very low stock prices or to consider other denominators, such as the standard error of forecast errors. Our own analysis (discussed below) shows that the tail asymmetry is most pronounced when deflating by price.

The distribution of analysts' forecast errors for the entire sample is summarized in A&L's Table 1 and plotted in their Figure 1. The results in Table 1, particularly the results reported in Panel A, confirm prior results: the mean of the forecast errors distribution is negative, the median is zero. In addition, Panel B of Table 1 shows that the absolute value of the forecast errors in the 5th and 10th percentiles is approximately twice as large as the forecast errors in the 90th and the 95th percentiles. This evidence is interpreted as indicative of the presence of the tail asymmetry. The existence of the middle asymmetry is documented in Panel C of Table 1 and Panel B of Figure 1. Both Panel C of Table 1 and Panel B of Figure 1 suggest that the frequency of small positive forecast errors is larger than the frequency of small negative forecast errors.

The conclusions related to the presence of the middle and tail asymmetries reached by A&L, however, are based on visual inspections of the table and figures. As such, it is difficult to infer the magnitude and statistical significance of those visual findings. Moreover, the inconsistency in the scales of the figures relative to the underlying distributions contributes to the readers' difficulty of interpretation. For example, while Figure 1 Panel A plots the cumulative distribution for the entire sample of forecast errors, Panel B (as does Figure 3) only plots the distribution truncated symmetrically around zero, which has been rescaled to represent 100 percent of the distribution within that interval. However, truncation and rescaling increases each bar in Panel B of Figure 1 by

a factor of 1.09.² A similar approach is used in Panels A, B, and C of Figure 3. In this case, the distributions are reported only for the -0.5 to +0.5 intervals, effectively rescaling the figures by a factor of 1.2. These rescalings can exaggerate the visual impression of the importance of the middle asymmetry.

To provide some additional evidence on the prevalence of these two asymmetries, we compute the frequency of the analysts' forecast errors distribution in symmetric intervals around the mean.³ We choose the width of those intervals to be 0.5 standard deviations of the analysts' forecast errors distribution. Following the instructions in A&L we collected a sample of 29,285 analysts' forecast errors (versus their 33,580).⁴ This sample size is limited because of matching with COMPUSTAT. However, all of the subsequent results are also performed using the entire sample of 55,386 analysts' forecast errors. None of our conclusions differ between the full and the reduced sample. We report the results for the reduced sample to make them comparable to the analysis of A&L.

We compute analysts' forecast errors following the same procedure as A&L by using the ZIR actual earnings and deflate by price, multiplying by 100. The mean, median, and standard deviation of the analysts' forecast errors sample is -0.0875, 0.7429, and 0.00523, respectively (compared to A&L's -0.126, 0.995, and 0.000, respectively). For this sample, we compiled the frequency of the forecast error distribution in symmetric intervals around the mean. In addition, we decompose the analysts' forecast errors into

² Panel C of Table 1 of A&L reports that 9 percent of the observations fall outside the [-1.0, +1.0] interval, thus the remaining observations are rescaled by a factor of 1.09 to add up to 100 percent.

³ This discussion then, raises the question on how one would document the presence of any unusual patterns, such as middle and tail asymmetries, for a given distribution. To derive such a test, one would have to first define what the distribution should look like absent any unusual patterns. The derivation of what the distribution should be and what should be considered as the tail of the analysts' earnings forecast errors distribution is beyond the scope of our discussion.

⁴ We use COMPUSTAT data for the period 1987-1999 given the evidence in Collins and Hribar (2002). Using this sample period we use SFAS No. 95 (FASB 1987) statement of cash from operations data to estimate accruals, rather than a balance sheet approach.

actual earnings and consensus forecasts (both price deflated). For comparison purposes, we included the expected frequencies and their 95 percent confidence interval for 29,285 independent draws from a normal distribution. Following A&L, we use the normal distribution as a benchmark.⁵ The results are reported in Table 1.

INSERT TABLE 1 AROUND HERE

Consistent with A&L, our results indicate, that relative to a normal distribution, the forecast errors distribution is characterized by a significant over-representation of very small negative forecast errors (the left-tail asymmetry), a significant under-representation of small forecast errors below the mean, and a significant over-representation of small forecast errors above the mean (jointly the middle asymmetry). Also, consistent with A&L, we find that the observations falling into the left-tail asymmetry are indeed very negative, and much more frequent than those suggested by the normal distribution. However, they comprise only 2.107 percent of the sample.⁶ Whether such a small frequency is likely to affect the conclusions reached in subsequent tests depends on the statistical procedures used. At any rate, it seems advisable for researchers to conduct diagnostics including the impact of influential observations and to examine the sensitivity of their conclusions depending on whether parametric or non-parametric tests are used. We will return to this issue in Section 4.

⁵ We recognize that there is no economic theory suggesting how analysts' forecast errors are distributed. We use the normal distribution as a benchmark only. Further, note that the confidence intervals computed in Table 1 assume independent draws. However, the confidence intervals need to be widened because analysts' earnings forecast errors are not likely to be independent.

⁶ Of course, this assessment depends on the width of the intervals used (0.5 standard deviations in our case). However, even if one were to include the next interval, the frequency of occurrence of the left-tail asymmetry would be not greater than 2.6 percent.

The middle asymmetry is much more pronounced than the left-tail asymmetry. We find 12.313 percent of the observations in the interval of 0.5 standard deviations below the mean (that is 6.8 percentage points less than expected from a normal distribution). However, there are 67.041 percent of the observations in the interval of 0.5 standard deviations above the mean (48.0 percentage points more than expected from a normal distribution). In summary, as concluded by A&L, the left-tail asymmetry is characterized by a small number (i.e., 2 to 3 percent) of very negative forecast errors; the middle asymmetry is characterized by a very large number of very small (0.05 percent of stock price) forecast errors. The economic implications of those two asymmetries depend on the specific economic context and on the research method used. For example, while very frequent, the very small magnitude analysts' forecast errors comprising middle asymmetry are not likely to be significant in a pricing context.

An interesting question raised is whether the “unusual” shape of the analysts' forecast errors distribution results from the distribution of earnings or analysts' earnings forecasts. When decomposing forecast errors into price-deflated earnings and price-deflated forecasts (the two components of analysts' forecast errors, see equation (1) above) we find that both the middle and tail asymmetries are still present, but attenuated. The results reported in Table 2 indicate that the percentage of analysts' forecast errors falling in the interval to the left of three standard deviation below the mean is 15 times larger than that of a normal distribution ($2.107/0.135$). In contrast, the ratio of observations in the tail (relative to a normal distribution) is 7.35 ($0.992/0.135$) for actual earnings and 6.53 ($0.882/0.135$) for analysts' forecasts. We find a similar attenuation for the middle asymmetry: the ratio of analysts' forecast errors falling in the interval one half standard deviations above the mean to the number of observations falling one half standard

deviations below the mean is 5.22 (67.041/12.313). In contrast, this ratio is 1.57 (48.266/30.646) for actual earnings and 1.31 (45.442/34.591) for analysts' forecasts. Thus, a challenge for future research is to explore why the pairing between actual earnings and forecasts increases the middle and the left-tail asymmetries.⁷ A promising line of research would be to analyze the economic setting in which analysts and managers operate; that is, what are analysts' and managers' objective functions (see Liu, 2003).

INSERT TABLE 2 AROUND HERE

The final element affecting the shape of the distribution of analysts' forecast errors is the choice of the deflator. The inverse of price has a right skewed distribution. While deflating by price is often suggested in the literature (Christie, 1987), the earnings-price ratio is, among other factors, a function of a company's assets-in-place relative to growth opportunities: Earnings typically result from assets-in-place, while price includes the value of assets-in-place plus growth opportunities. Thus, firms with a larger fraction of assets-in-place will have higher earnings-price ratios. Therefore, deflating by price is likely to affect the cross-sectional distribution of analysts' earnings forecast errors, depending on the composition of the sample.⁸ Preliminary tests using the deflators considered by A&L indicate the frequency in the lowest cell of analysts' forecast errors

⁷ In untabulated results, A&L reject normality on the basis of skewness and kurtosis (footnote 7), while acknowledging that there are no theoretical reasons to expect normality in the first place. Nevertheless, the interesting question is whether the non-normality is present in the components used to compute price deflated analysts' forecast errors. Burgstahler and Dichev (1997), Dechow, Richardson, and Tuna (2003) Beaver, McNichols, and Nelson (2003) document distributional properties of reported earnings – and the consensus that emerges is that the earnings distribution has a discontinuity around zero and is not symmetric.

⁸ For example, the accrual anomaly literature concludes that the choice of price versus sales as a deflator significantly changes the conclusions (e.g., Zach, 2003). Relatedly, Dechow et al. (2003) shows that the distribution of earnings is materially affected by the choice of the deflator.

declines to 1.588 percent for undeflated forecast errors, 0.866 percent for forecast errors deflated by the absolute value of the consensus forecast, and to 0.545 percent for forecast errors deflated by the consensus forecast.⁹ Consistent with A&L’s footnote 13, these results suggest that the presence of the tail asymmetry does not depend on the use of one of those deflators. However, the magnitude of the tail asymmetry is sensitive to the specific deflation, with price deflation providing the “highest” occurrence.

3. Analyst Forecast Error Bias

A&L emphasize that given the researcher’s choice of summary statistics, evidence can be found for analysts’ optimism, pessimism, or unbiasedness. The conclusion is derived from their analysis in Table 1 and Figure 1. Consistent with the literature, A&L document a negative mean forecast error – an indication that, on average, analysts are optimistic. They point out that the median forecast error is zero, which they interpret as “suggesting unbiased forecasts” (page 5). Finally, they suggest that one could interpret the fact that 48 percent of the forecast errors are positive compared to the fact that 40 percent are negative (Panel A of Table 1) as evidence of analyst pessimism.

This interpretation relies on the definition of bias. Typically, bias is defined relative to an expectations model. Thus, analysts’ forecasts would be unbiased if the expected value of the forecast errors were zero. Assuming a quadratic loss function, the mean represents the best estimate of expectations. Under these conditions, one would have to conclude that analysts’ forecasts are optimistic because the mean forecast error is negative. One alternative is to examine the median. The median forecast error is a measure of expectations of analysts’ forecasts only if analysts’ objectives are to forecast

⁹ In footnote 13, A&L report that they analyzed undeflated forecast errors and errors deflated by forecasts.

the median of the earnings distribution. A sufficient condition for this to occur is that analysts' loss function is linear, making forecasting the median the value-maximizing forecast. Indeed, this is the motivation of a growing literature that tests hypotheses about analysts' loss function (e.g., Gu and Wu 2003; Basu and Markov, 2003; Liu, 2003). It is surprising that A&L dismiss this line of research by saying that this research strategy "will not provide definitive answers to the question of whether analysts' forecasts are biased and inefficient." (page 2).

Finally, while strictly true that the percentage of positive forecast errors exceeds the percentage of negative forecast errors, A&L's evidence shows that a large fraction of the positive forecast errors are very small and are not likely to be of economic significance. In contrast, the negative forecast errors are large (hence, the negative average forecast error). Therefore, the main economic impact (as opposed to a frequency count without regards to the magnitudes) is one of pessimism.

4. Efficiency of Analysts' Forecasts

Prior research concludes that analysts' forecasts are inefficient. For example, Brown, Foster and Noreen (1985), Klein (1990), Lys and Sohn (1990), Elgers and Murray (1992), Abarbanell and Bernard (1992), and Chan, Jegadeesh, and Lakonishok (1996) document that analysts' forecast errors are correlated with prior stock returns. Similarly, Brown and Rozeff (1979), Mendenhall (1990), Abarbanell and Bernard (1992), Jacob and Lys (1992), Ali, Klein, and Rosenfeld (1993), Shane and Brous (2001), and Alford and Berger (1999) document that analysts' forecast errors are serially correlated. Finally, Jacob and Lys (1992) show that large absolute values of serial correlation in forecast

errors are associated with larger forecast errors – ignoring pertinent information leads to less precise forecasts.

A&L first focus on whether analysts over- or under-react to prior information (prior abnormal returns and prior earnings changes). They document that analysts under-react to both types of information. This result is robust with respect to parametric and non-parametric tests of the association between forecast errors and prior bad news. However, A&L detect an under-reaction to prior good news only when using non parametric tests. This result is consistent with their main conclusion that forecast errors are not normally distributed and, hence, the use of parametric tests that explicitly or implicitly rely on the normal distribution of the dependent variable may not be appropriate. However, in the end, their results are consistent with the tenor of prior literature: analysts under react to both prior good news and prior bad news and are, thus, inefficient.

A&L critique the literature documenting a serial correlation in analysts’ forecast errors because the test procedures are unduly influenced by the left-tail asymmetry. Specifically, A&L document in their Table 5 that the Pearson and Spearman correlations between consecutive forecast errors for the overall sample of 0.15 and 0.22 (significant at the 1 percent level) respectively are entirely driven by “observations in the extreme left tail associated with the tail asymmetry.” A&L derive their results by sorting the data on current forecast errors and computing the correlation between consecutive forecasts within each decile.¹⁰ However, sorting on the dependent variable is likely to lead to erroneous inferences.

A&L obtain these results by creating a matrix of forecast errors pairs; that is, each row of the matrix consists of $CFE_{j,q}$ and $CFE_{j,q-1}$, where j represents the firm and q

¹⁰ In previous drafts, A&L sorted on the preceding forecast error. Because the correlation of x and y is identical to the correlation of y and x , the discussion that follows applies equally to those previous results.

represents the quarter. Next, they sort this matrix by the current forecast (the first column) and divide it into deciles. Then, they compute the correlation between the two columns for the overall forecast errors and by decile ranking of current forecast errors.

These results are reported in their Table 5. Once the observations are sorted, only the decile with the most negative forecast errors has both Pearson and Spearman correlations that are positive and highly significant (0.17 and 0.19, respectively). Of the other 9 deciles, only 3 deciles have both the Pearson and Spearman correlations significant at conventional levels. Of the remaining six deciles, none have a significant Pearson correlation (and some are even negative), while all the Spearman correlations are significantly positive at conventional levels.

Interpreting these results is difficult for two reasons. First, the Spearman correlation in the unsorted sample is 0.22 and statistically significant at the 1 percent level. However, one typically would compute the Spearman (rank correlation) when there is concern that the distribution of the data makes the Pearson correlation undesirable or misleading (e.g., because of large influential observations). Note that converting forecast errors to ranks eliminates both the magnitude and the frequency aspect of the left-tail asymmetry. Thus, neither the magnitude nor the frequency aspects of the left-tail asymmetry can affect the Spearman correlation. Second, sorting on the dependent variable is likely to bias the tests, rendering interpretation of the results difficult.

To elaborate on the second point, we perform the following experiment. We first generate two vectors x and ε of random variables each representing 35,000 observations drawn from a standard normal distribution. This corresponds approximately to the number of forecast errors in A&L's sample. Next, we define a new variable y_q as:

$$y_q = 0.15 \times x_q + \varepsilon_q \quad (2)$$

One can interpret equation (2) as the relation between the current forecast error (y_q) and the previous forecast error (x_q). It follows from (2) that since x and ε are normally distributed so is y . In other words, the distribution of our ‘forecast errors’ follow a normal distribution by construction.

We now compute the Pearson and Spearman correlations between y and x , both for the overall sample and for the ten deciles sorted on y . The results for 100 repetitions of the above simulation are reported in Panel A of Table 3 which has been formatted to resemble Table 5 in A&L.

INSERT TABLE 3 PANEL A and PANEL B AROUND HERE

As indicated in the last row of Panel A, the overall Pearson and Spearman correlations are 0.148 and 0.140, consistent with the seeded correlation of 0.15 in equation (2). However, once the data are sorted on the dependent variable y and grouped into deciles, the correlations within each decile become very small in magnitude. For the lowest and highest deciles, we find average Pearson (Spearman) correlations of 0.060 and 0.062 (0.055 and 0.056), and average p-values of 0.009 and 0.008 (0.013 and 0.013). However, for the middle eight deciles the correlations range from 0.009 to 0.020 and the average p-values range from 0.331 to 0.451. In addition, we report in Panel B of Table 3 the empirical 5th and 95th percentiles of estimated correlations using 100 simulations of 35,000 observations each. As indicated in Panel B, a zero correlation is within the 90 percent range for all but one of middle deciles (2 through 9; row 7 column 3 is the sole exception). However, for those middle deciles, the 95th percentiles ranges from 0.034 to

0.050, making it likely that for single draws, at least some of those middle correlations will be significantly larger than zero.

What this experiment demonstrates, is that even absent any asymmetries, one would conclude that the extreme deciles contribute to the overall correlation. In summary, Table 3 suggests that sorting on the variable of interest has biased the estimated correlation coefficients within each decile toward zero.¹¹ More importantly, this procedure rejects the null hypothesis that all observations equally contribute to the overall serial correlation, even when we know this hypothesis to be true by construction. It is important to note that the purpose of our simulation is not to recreate the results of A&L's Table 5;¹² rather, the purpose of this simulation is to question the reliability of the results created by the same procedure that A&L use. Thus, what our simulation suggests is that the question whether the tail asymmetry drives the overall correlation cannot be answered using their sorting procedure.¹³

Providing evidence whether extreme observations are responsible for the overall serial correlation is very difficult because such tests necessarily pre-select the data based on ex-post realizations. We attempt to answer this question by first identifying influential observations with commonly used econometric procedures. While we are not aware of such procedures to identify influential (or suspicious) observations for correlations, such procedures are commonly used in a regression context. For example, observations whose studentized residual have an absolute value greater than two are considered "suspicious."

¹¹ The way to visualize this effect is to imagine a regression with a slope of 0.15. Now, partition the data horizontally into ten subgroups. Within each group, the number of observations that lie on each side of the regression line is the same. This is true for all subgroups, except the lowest and highest groups. For those two extreme groups, there are somewhat more observations above the original regression line than below that line. As a result, for each subgroup, the estimated regression will have a zero slope, while the two extreme groups will have a positive slope.

¹² To recreate their Table 5 we would have to choose a skewed, rather than a symmetric distribution.

¹³ For a related discussion of the effect introduced by sorting see Berk (2000).

Note that as long as the data are stationary, the probability limit of the slope coefficient in an auto-regression (regressing forecast errors on preceding forecast errors) equals the Pearson correlation among consecutive forecast errors.

To establish a base line of the differences between estimates using regression vs. correlation techniques, we first computed the lag one serial correlation using Pearson, Spearman and auto-regression using the sample described in Table 1. For the entire sample, those results are 0.242, 0.235, and 0.265, all statistically significant at the one percent level. Then, we performed an analysis of influential observations using studentized residuals. We marked all observations with an absolute studentized residual in excess of 2.0 as suspicious. The procedure netted 1,322 (4.57 percent) influential observations. This is slightly more than the 4.224 percent of the observation falling outside the two standard deviations of the sample identified in Table 1 ($2.107+0.492+0.775+0.444+0.355+0.051=4.224$). Consistent with A&L's conjecture, 904 (68.4 percent) of those influential observations are in the lowest decile, while 350 (26.5 percent) are in the most positive decile. The remaining 68 (5.1 percent) influential observations are scattered across the middle eight deciles (with 7, 2, 8, 2, 4, 5 observations for deciles two through nine, respectively). Note, that given the nature of the data (forecast errors and lag forecast errors) it is not surprising that most of the observations with an absolute value of the studentized range in excess of 2.0 fall in the extreme deciles.

To investigate whether those influential observations unduly affect the correlation results, we re-estimate the serial correlation using Pearson, Spearman, and regression, after excluding those 1,322 observations. (It is important to note, however, that we are not advocating excluding observations. Rather, we only perform this analysis to

investigate the sensitivity of the overall correlations to the inclusion/exclusion of those suspicious observations.) The resulting correlations are 0.307, 0.248, and 0.195, respectively, all statistically significant at the one percent level. Note that, as suggested by our discussion above, the Spearman (rank) correlation did not change substantially by excluding the influential observations.

Based on these analyses, we find no evidence to suggest that influential observations drive the overall serial correlation between consecutive analysts' earnings forecasts. Indeed we hypothesize that the serial correlation of forecast errors is a real phenomenon that both characterizes analysts' behaviors and has real economic consequences. To demonstrate this, we perform the following analysis using the sample described in our Table 1. We first rank the firms into deciles using the firm-specific serial correlation and plot the associated analysts' forecast errors. The results are plotted in Figure 1.

INSERT FIGURE 1 AROUND HERE

For comparison purposes, we also plot A&L's forecast errors as reported in column four of Table 5. In contrast to A&L's S-shaped relation (denoted as FE A&L in Figure 1), we obtain an inverse U-shaped relation between ranking of serial correlation and average forecast errors (denoted as AVG_FE in Figure 1).

To provide an economic interpretation of the results using a more formal test and, we compute the average absolute forecast error for each company. We then estimate a quadratic equation

$$\overline{FE}_j = a + b \times Cor_j + c \times Cor_j^2 + \varepsilon_j \quad (3)$$

where for firm j , \overline{FE} is the average analysts' forecast error, and Cor is the serial correlation in analysts' forecast errors. The analysis is performed on the entire sample of 2,740 companies. To facilitate visual representation, we also form fifty groups based on the serial correlations (plotting the entire sample results in a very dense cloud of points that does not allow visualization).¹⁴

INSERT TABLE 4 AROUND HERE

The results reported in Table 4 indicate that both b and c are significant at the 1 percent level. We plot the actual group means and the fitted values in Figure 2. Figure 2 shows a U-shaped relation between the average absolute forecast errors and the serial correlation decile. For a similar approach see Jacob and Lys (1992). That paper demonstrates that the serial correlation of forecast errors is indeed a quadratic (U-shaped) function with the minimum forecast error occurring near zero correlation (as one would expect).¹⁵ This relation has a clear economic interpretation – analysts who ignore the serial correlation have larger forecast errors. In other words, the U-shaped relation is an indication of inefficiency by analysts' whose forecasts are serially correlated. (For a more detailed analysis of this relation, see Jacob Lys, 1992.)

INSERT FIGURE 2 AROUND HERE

In summary, in this section, we demonstrate that A&L's sorting procedure can lead to biased results. We then show that using standard econometric techniques to identify

¹⁴ The results are not sensitive to whether we form 20, 50, and 100 groups.

¹⁵ In our Figure 2, the minimum forecast error occurs at correlation = 0.19.

influential observation does not lead one to conclude that the serial correlation in analysts' forecast errors is related to the tail asymmetry. Finally, we show, that as one would expect, the serial correlation of analysts' forecast errors results in less accurate forecasts.

5. Biased Forecasts or Biased Earnings?

In trying to establish an empirical link between the two documented asymmetries and discretionary accrual choices, A&L focus on the association between unexpected accruals and the distribution of forecast errors. Regarding the negative tail asymmetry, A&L suggest that: "If analysts' forecasts do not account for the fact that some firms will recognize accruals placing them in the extreme negative tails of the distribution of unexpected accruals, then there will be a correspondence between the negative tail of this distribution and the extreme negative tail of the forecast error distribution" (page 25). However, A&L do not provide a model or an economic basis for establishing the link between discretionary managerial choices and analysts' forecast errors.

Note that if the left-tail asymmetry were caused by managers' accrual choices, it should be no less pronounced in the earnings distribution than in the forecast errors distribution. One indication that other forces are likely to be at work is provided in our Table 2. In particular, our results indicate that the left-tail asymmetry is twice as frequent in the distribution of analysts' forecast errors as compared to the distribution of actual (deflated) earnings (2.107 percent versus 0.992 percent). Thus, this evidence suggests that other factors may influence the left-tail asymmetry and, as a consequence, affect the conclusion of whether "extreme negative unexpected accruals included in

reported earnings go hand in hand with observations in the cross-section that generate the tail asymmetry” (A&L, page 5).

Is the forecast error distribution caused by analysts’ issuing their forecasts in reaction to managers’ choices? Economic intuition would suggest that if analysts’ objectives were to issue accurate earnings forecasts, then analysts would anticipate and incorporate future managerial discretionary accrual behavior in their forecasts. Liu (2003) studies the properties of analysts’ earnings forecasts by hypothesizing that analysts are aware of earnings management practices of firms, and incorporate such discretionary behavior into their forecasts. She finds that analysts’ earnings forecasts are systematically below the median of the earnings distribution (assuming that analysts’ objective is to minimize mean absolute forecast errors) for firms with high accounting reserves, negative forecasted earnings, and negatively skewed unmanaged earnings. Her results suggest that analysts choose to forecast earnings below the median earnings to avoid large optimistic forecast errors caused by downward earnings management by firms.

A&L’s conclusion is reached by their visual inspection of Figure 6. Therefore, while suggestive, their evidence is not based on formal tests that control for other factors that might affect forecast errors and unexpected accruals. For example, when we add back the unexpected accruals (computed following the method outlined in A&L) we find that approximately 35 percent of the forecast errors in the tail asymmetry would be positive. The question then becomes, why did those managers take such extreme discretionary accruals, given that they could have beaten the forecast if they had not done so? One possible reason is that those accruals are not, in fact, discretionary – possibly resulting from the reversal of previous discretionary accruals or a strategic write-off to manage subsequent earnings. If this were the case, then analysts are indeed very unsophisticated

not to anticipate those reversals and or earnings management. Alternatively, those discretionary accruals may not be related to earnings management but result from other corporate activities, for example restructuring charges following mergers or the like. Again, depending on the circumstances, sophisticated analysts may have been able to anticipate those accruals as well. Finally, our own replication of the analysis underlying A&L's Figure 6 suggests that negative forecast errors are in general associated with negative accruals, while positive forecast errors, with positive accruals.

In summary, the relation between unexpected accruals and analysts' forecast errors is an interesting question which requires more formal analysis. At this stage, the reader is left with an unanswered question – is the distribution of unexpected accruals due to managerial discretion, the estimation technique undertaken by A&L, or due to a third factor that is omitted from the analysis? Addressing these questions and discussing the strategic interactions between management and financial analysts is necessary to conclude whether unexpected accruals are related to extreme analysts' forecast errors.

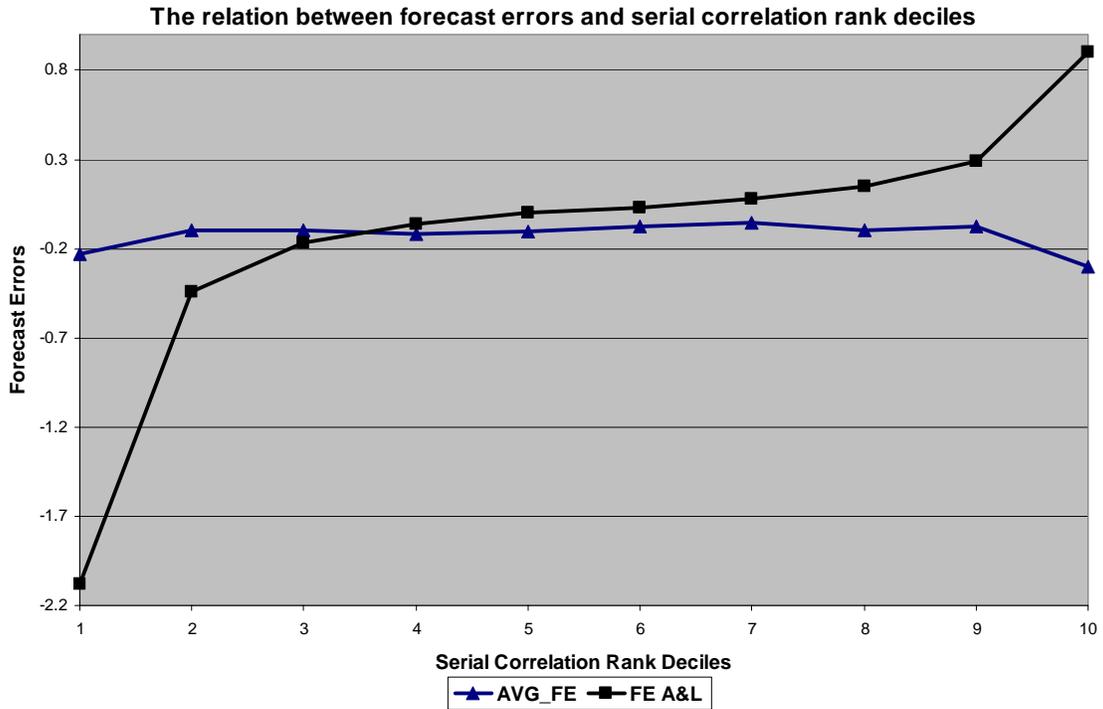
6. Conclusions

As stated in the introduction, A&L's major contribution is that analysts' forecast errors are not normally distributed. This is an important insight because research often relies on techniques that are sensitive to the distributional assumptions of analysts' forecast errors. Second, A&L document that the forecast error distribution is characterized by two asymmetries, the left-tail asymmetry and the middle asymmetry.

Beyond this insight, A&L provide some detailed and provocative results. However, our analysis raises the question of the reliability of these results. In the end, the interplay

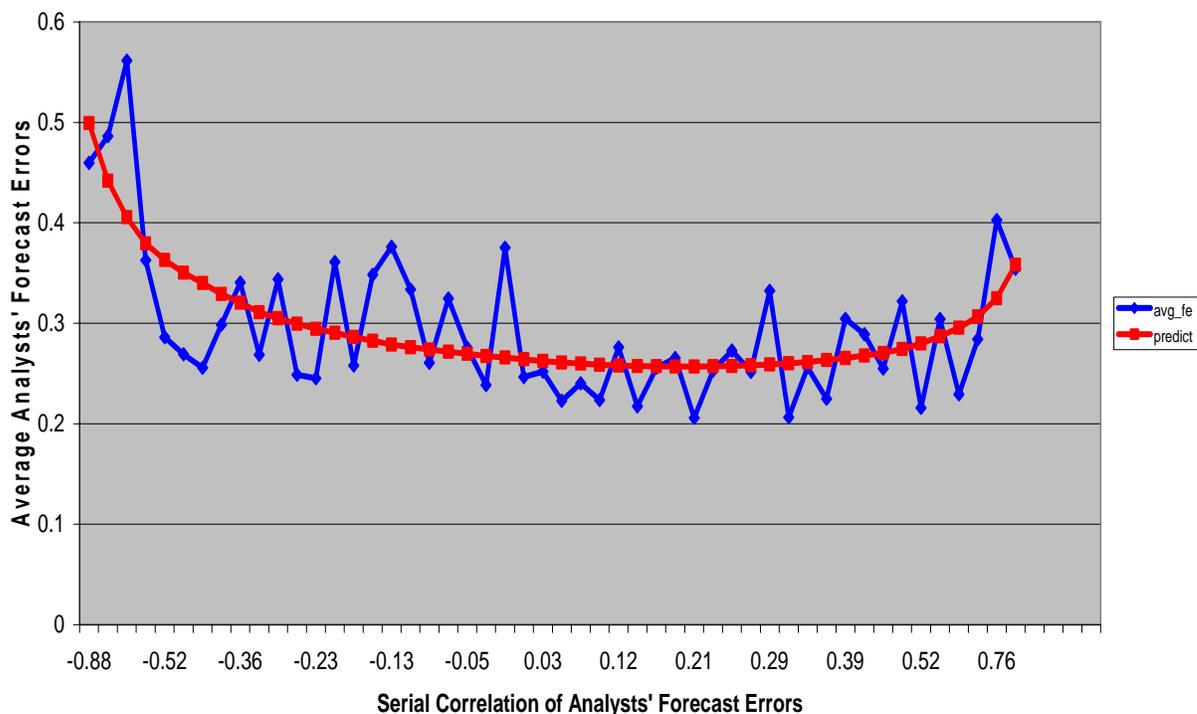
between A&L's results and our discussion provide some clear avenues for future research.

Figure 1



This figure compares the relation between the serial correlation in consecutive analysts' forecast errors and the magnitude of those forecast errors. FE A&L reports the relation as documented in Table 5 of Abarbanell and Lehavy, obtained by sorting on the dependent variable (current forecast errors). In contrast we first compute the firm specific serial correlation. We then sort those correlations into deciles and compute for each decile the average forecast error (AVG_FE). In contrast to the sharp S shape of the relation between serial correlation decile and average forecast errors, we obtain a very "mild" inverted U-shaped relation.

Figure 2
The Relation between Forecast Errors and Serial Correlation Ranks



This figure compares the relation between the serial correlation in consecutive analysts' forecast errors and the average absolute forecast errors. We first compute the firm specific serial correlation. We then sort those correlations into 50 groups and compute for each group the average absolute forecast error. The fitted quadratic function is $\overline{FE}_j = a + b \times Cor_j + c \times Cor_j^2 + \varepsilon_j$. The regression coefficients are reported in Table 4. The figure documents that analysts that have either a low (negative) or high (positive) correlation of their forecast errors are less accurate in forecasting earnings.

Table 1
Distribution of Analysts' Forecast Errors, Actual Earnings and
Analysts' Consensus Forecasts 1987-1999
Sample Size = 29,285

Mean Median Standard Deviation	-0.0875 0.0052 0.7429	Forecast Errors		Normal Distribution		
		Frequency	Mean	Expected frequency	5 th and 95 th percentiles	
		$FE \leq \overline{FE} - 3.0\sigma$	2.107 ⁺	-4.016	0.135	[0.093, 0.177]
		$\overline{FE} - 3.0\sigma < FE \leq \overline{FE} - 2.5\sigma$	0.492	-2.127	0.486	[0.406, 0.566]
		$\overline{FE} - 2.5\sigma < FE \leq \overline{FE} - 2.0\sigma$	0.775 ⁻	-1.746	1.654	[1.507, 1.801]
		$\overline{FE} - 2.0\sigma < FE \leq \overline{FE} - 1.5\sigma$	1.014 ⁻	-1.372	4.406	[4.169, 4.642]
		$\overline{FE} - 1.5\sigma < FE \leq \overline{FE} - 1.0\sigma$	1.800 ⁻	-0.993	9.185	[8.852, 9.518]
		$\overline{FE} - 1.0\sigma < FE \leq \overline{FE} - 0.5\sigma$	3.722 ⁻	-0.617	14.988	[14.577, 15.400]
		$\overline{FE} - 0.5\sigma < FE \leq \overline{FE}$	12.313 ⁻	-0.221	19.146	[18.693, 19.600]
		$\overline{FE} < FE \leq \overline{FE} + 0.5\sigma$	67.041 ⁺	0.046	19.146	[18.693, 19.600]
		$\overline{FE} + 0.5\sigma < FE \leq \overline{FE} + 1.0\sigma$	6.905 ⁻	0.421	14.988	[14.577, 15.400]
		$\overline{FE} + 1.0\sigma < FE \leq \overline{FE} + 1.5\sigma$	2.059 ⁻	0.183	9.185	[8.852, 9.518]
		$\overline{FE} + 1.5\sigma < FE \leq \overline{FE} + 2.0\sigma$	0.922 ⁻	1.204	4.406	[4.169, 4.642]
		$\overline{FE} + 2.0\sigma < FE \leq \overline{FE} + 2.5\sigma$	0.444 ⁻	1.564	1.654	[1.507, 1.801]
		$\overline{FE} + 2.5\sigma < FE \leq \overline{FE} + 3.0\sigma$	0.355 ⁻	1.943	0.486	[0.406, 0.566]
		$\overline{FE} + 3.0\sigma < FE$	0.051 ⁻	2.195	0.135	[0.093, 0.177]

The sample was collected using the same criteria as Abarbanell and Lehavy (this issue). Forecast errors (FE) are defined as the actual earnings per share as reported by Zacks Investment Research (ZIR), minus the ZIR consensus forecast outstanding just prior to the disclosure of the quarterly earnings. Actual earnings are as reported by ZIR. Forecast errors, actual earnings and consensus forecasts are deflated by the stock price at the beginning of the quarter and multiplied by 100 (see equation (1)). The last column reports the expected frequency and the 5th and 95th percentiles of the expected frequency for a normal distribution.

⁺ indicates a frequency above the expected 95th percentile of the frequency interval.

⁻ indicates a frequency below the expected 5th percentile of the frequency interval.

Table 2				
Distribution of Analysts' Forecast Errors, Actual Earnings and Analysts' Consensus Forecasts 1987-1999				
Sample Size = 29,285	Forecast Errors	Actual Earnings	Analysts' Consensus Forecasts	Normal Distribution
Mean (\bar{X})	-0.0875	0.7764	0.8575	Expected frequency [5 th and 95 th percentiles]
Standard Deviation (σ_X)	0.7429	2.881	2.706	
$X \leq \bar{X} - 3.0\sigma_X$	2.107 ⁺	0.992 ⁺	0.882 ⁺	0.135 [0.093, 0.177]
$\bar{X} - 3.0\sigma_X < X \leq \bar{X} - 2.5\sigma_X$	0.492	0.398 ⁻	0.315 ⁻	0.486 [0.406, 0.566]
$\bar{X} - 2.5\sigma_X < X \leq \bar{X} - 2.0\sigma_X$	0.775 ⁻	0.698 ⁻	0.519 ⁻	1.654 [1.507, 1.801]
$\bar{X} - 2.0\sigma_X < X \leq \bar{X} - 1.5\sigma_X$	1.014 ⁻	1.179 ⁻	0.999 ⁻	4.406 [4.169, 4.642]
$\bar{X} - 1.5\sigma_X < X \leq \bar{X} - 1.0\sigma_X$	1.800 ⁻	2.157 ⁻	1.787 ⁻	9.185 [8.852, 9.518]
$\bar{X} - 1.0\sigma_X < X \leq \bar{X} - 0.5\sigma_X$	3.722 ⁻	4.124 ⁻	4.231 ⁻	14.988 [14.577, 15.400]
$\bar{X} - 0.5\sigma_X < X \leq \bar{X}$	12.313 ⁻	30.646 ⁺	34.591 ⁺	19.146 [18.693, 19.600]
$\bar{X} < X \leq \bar{X} + 0.5\sigma_X$	67.041 ⁺	48.266 ⁺	45.442 ⁺	19.146 [18.693, 19.600]
$\bar{X} + 0.5\sigma_X < X \leq \bar{X} + 1.0\sigma_X$	6.905 ⁻	8.992 ⁻	8.753 ⁻	14.988 [14.577, 15.400]
$\bar{X} + 1.0\sigma_X < X \leq \bar{X} + 1.5\sigma_X$	2.059 ⁻	1.646 ⁻	1.504 ⁻	9.185 [8.852, 9.518]
$\bar{X} + 1.5\sigma_X < X \leq \bar{X} + 2.0\sigma_X$	0.922 ⁻	0.456 ⁻	0.456 ⁻	4.406 [4.169, 4.642]
$\bar{X} + 2.0\sigma_X < X \leq \bar{X} + 2.5\sigma_X$	0.444 ⁻	0.183 ⁻	0.204 ⁻	1.654 [1.507, 1.801]
$\bar{X} + 2.5\sigma_X < X \leq \bar{X} + 3.0\sigma_X$	0.355 ⁻	0.111 ⁻	0.104 ⁻	0.486 [0.406, 0.566]
$\bar{X} + 3.0\sigma_X < X$	0.051 ⁻	0.152	0.214 ⁺	0.135 [0.093, 0.177]

The sample was collected using the same criteria as Abarbanell and Lehavy (this issue). X represents forecast errors in column 2, actual earnings in column 3, and consensus forecasts in column 4. Forecast errors are defined as the actual earnings per share as reported by Zacks Investment Research (ZIR), minus the ZIR consensus forecast outstanding just prior to the disclosure of the quarterly earnings. Actual earnings are as reported by ZIR. Forecast errors, actual earnings and consensus forecasts are deflated by the stock price at the beginning of the quarter and multiplied by 100 (see equation (1)). The last column reports the expected frequency and the 5th and 95th percentiles of the expected frequency for a normal distribution.

⁺ indicates a frequency above the expected 95th percentile of the frequency interval.

⁻ indicates a frequency below the expected 5th percentile of the frequency interval.

Table 3						
Panel A						
Correlations using Simulated Data						
Decile ranking of forecast errors	Average Pearson correlation (2a) and average p-values (2b) across 100 simulations		Average Spearman correlation (3a) and average p-values across 100 simulations		Current quarter forecast errors	Prior quarter forecast errors
	(2a)	(2b)	(3a)	(3b)	(4)	(5)
Lowest	0.060	0.009	0.055	0.013	-1.7748	-0.2586
2	0.020	0.331	0.018	0.170	-1.0559	-0.1509
3	0.012	0.451	0.011	0.216	-0.6850	-0.1002
4	0.010	0.446	0.013	0.209	-0.3906	-0.0572
5	0.013	0.405	0.010	0.229	-0.1273	-0.0171
6	0.009	0.448	0.012	0.216	0.1277	0.0195
7	0.013	0.394	0.009	0.231	0.3916	0.0535
8	0.016	0.326	0.012	0.214	0.6858	0.1024
9	0.019	0.331	0.019	0.161	1.0562	0.1555
Highest	0.062	0.008	0.056	0.013	1.7734	0.2554
Overall	0.148	0.000	0.140	0.000	0.0001	0.0002

This table reports the average Pearson and Spearman correlations estimated from 100 random samples. Each sample is constructed by drawing 35,000 random variables x and ε independently from a standard normal distribution. We then construct a new random variable y using $y_q = 0.15 \times x_q + \varepsilon_q$, thus inducing a correlation of 0.15 between y and x . We label y as the ‘current quarter forecast error’ and x as the ‘prior quarter forecast error.’ Each sample is first sorted on current forecast errors (y) and assigned to 10 deciles of 3,500 observations each. The first ten rows, from “Lowest” to “Highest” represent the average correlations for the ten decile sub-samples obtained after sorting the entire sample on y . The last row “Overall” reports the correlation before sorting the samples.

Table 3
Panel B
Distribution of Correlation And Associated p-Values of Simulated Data

		Estimated correlations across 100 random samples			average p-value of correlations across 100 random samples
		Average correlation	5 th percentile	95 th percentile	
Lowest	Pearson	0.060	0.036	0.085	0.009
	Spearman	0.055	0.026	0.083	0.013
2	Pearson	0.020	-0.005	0.050	0.331
	Spearman	0.018	-0.013	0.046	0.170
3	Pearson	0.012	-0.011	0.038	0.451
	Spearman	0.011	-0.018	0.044	0.216
4	Pearson	0.010	-0.018	0.036	0.445
	Spearman	0.013	-0.015	0.040	0.209
5	Pearson	0.013	-0.016	0.041	0.405
	Spearman	0.010	-0.017	0.034	0.229
6	Pearson	0.009	-0.018	0.037	0.448
	Spearman	0.012	-0.019	0.037	0.216
7	Pearson	0.013	-0.015	0.037	0.395
	Spearman	0.009	0.019	0.037	0.231
8	Pearson	0.016	-0.014	0.038	0.326
	Spearman	0.012	-0.013	0.041	0.214
9	Pearson	0.019	-0.013	0.048	0.331
	Spearman	0.019	-0.070	0.044	0.161
Highest	Pearson	0.062	0.033	0.092	0.008
	Spearman	0.056	0.031	0.084	0.013
Overall	Pearson	0.148	0.132	0.152	0.000
	Spearman	0.140	0.131	0.149	0.000

This table reports the 5th and 95th percentiles of the Pearson and Spearman correlations and the associated average p-values for testing the null hypothesis of zero serial correlation estimated from 100 random samples. Each sample is constructed by drawing 35,000 random variables x and ε independently from a standard normal distribution. We then construct a new random variable y using $y_q = 0.15 \times x_q + \varepsilon_q$, thus inducing a correlation of 0.15 between y and x . We label y as the ‘current quarter forecast error’ and x as the ‘prior quarter forecast error.’ Each sample is first sorted on current forecast errors (y) and assigned to 10 deciles of 3,500 observations each. The first ten rows, from “Lowest” to “Highest” represent the average correlations for the ten decile sub-samples obtained after sorting the entire sample on y . The last row “Overall” reports the correlation before sorting the samples.

Table 4
The Association Between Analysts' Forecast Errors and the Serial Correlation of Consecutive Analysts' Forecast Errors

$$\overline{FE}_j = a + b \times Cor_j + c \times Cor_j^2 + \varepsilon_j$$

	Individual Company Level	50 Groups based on Correlation
N	2,740	50
Intercept	0.3208	0.2647
t-statistic	26.92	27.29
Cor	-0.0875	-0.0808
t-statistic	-3.69	-4.21
Cor ²	0.2325	0.2104
t-statistic	4.88	5.42
R-squared	0.0129	0.4892
F	17.94	22.51

This table reports the association between analysts' earnings forecast errors and the serial correlation between consecutive forecast errors. \overline{FE}_j is the average consensus analysts' forecast error. Cor_j represents the serial correlation in consecutive consensus forecasts. Column 2 reports the regression results for the entire sample. Column 3 reports the regression results after grouping the sample into 50 groups based on Cor. The actual observations and the fitted values of the results in Column 3 are plotted in Figure 2.

References

- Abarbanell, J. and R. Lehavy, 2003, "Biased Forecasts or Biased Earnings? The Role of Reported Earnings in Explaining Apparent Bias and Over/Underreaction in Analysts' Earnings Forecasts," *Journal of Accounting and Economics*, this issue
- Abarbanell, J., and V. Bernard, 1992, "Tests of analysts' overreaction/underreaction to earnings information as an explanation for anomalous stock price behavior," *Journal of Finance* 47, 1181–1207.
- Alford, A., and P. Berger, 1999, "A simultaneous equations analysis of forecast accuracy, analyst following, and trading volume," *Journal of Accounting, Auditing & Finance*, 14, 219–240.
- Ali, A., A. Klein, and J. Rosenfeld, 1992, "Analysts' use of information about permanent and transitory earnings components in forecasting annual EPS," *Accounting Review*, 67, 183–198.
- Basu, S. and S. Markov, 2003, "Loss function assumptions in rational expectations tests on financial analysts' earnings forecasts," working paper Goizueta Business School, Emory University.
- Beaver, W. H., M. F. McNichols, and K. K. Nelson, 2003, "An alternative interpretation of the discontinuity in earnings distributions," Working paper, Graduate School of Business, Stanford University.
- Berk, J. B., 2000, "Sorting Out Sorts," *The Journal of Finance*, Vol. LV, No. 1, 407-428.
- Brown L., G. Foster, and E. Noreen, 1985, "Security analysts' multi-year earnings forecasts and the capital markets," *Studies in Accounting Research*, no. 21, Sarasota, Florida, American Accounting Association.
- Brown, L., and M. Rozeff, 1979, "University time series models of quarterly accounting earnings per share: a proposed model," *Journal of Accounting Research* 17, 179–189.
- Burgstahler, D., Dichev, I., 1997, "Earnings management to avoid earnings decreases and losses," *Journal of Accounting and Economics* 24, 99-126.
- Chan, L., N. Jegadeesh, and J. Lakonishok, 1996, "Momentum strategies," *Journal of Finance* 51, 1681–1713.
- Christie, A. A., 1987, "On cross-sectional analysis in accounting research," *Journal of Accounting and Economics* 9, 231-258.
- Collins, D. W., and P. Hribar, 2002, "Errors in estimating accruals: Implications for empirical research," *Journal of Accounting Research* 40, 105-134.

- Dechow, P. M., Richardson, S. A., Tuna, I. A., 2003, "Why are earnings kinky? An examination of the earnings management explanation," *Review of Accounting Studies*, 8, 355-384.
- Elgers, P., and D. Murray, 1992, "The relative and complementary performance of analyst and security-price-based measures of expected earnings," *Journal of Accounting and Economics* 15, 303–316.
- Financial Accounting Standard Board (FASB), 1987, "*Statement of Cash Flows*," Statement of Financial Accounting Standards No. 95, Norwalk, CT: FASB
- Gu, Z and S. Wu, 2003, "Earnings skewness and analysts' forecasts bias," *Journal of Accounting and Economics* 35, 5-29.
- Jacob, J. and T. Lys, 1992, "Determinants and implications of the serial-correlation in analysts' earnings forecast errors," working paper, Kellogg School of Management, Northwestern University.
- Klein, A., 1990, "A direct test of the cognitive bias theory of share price reversals," *Journal of Accounting and Economics* 13, 155–166.
- Kothari, S.P., J. Sabino, and T. Zach, 2003, "Implications of survival and data restrictions on tests of market efficiency," working paper, Massachusetts Institute of Technology.
- Liu, X., 2003, "Analysts' response to earnings management," working paper, Kellogg School of Management, 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 (December), 341-364.
- Marais, L. 1984, "An Application of the Bootstrap Method to the Analysis of Squared, Standardized Market Model Prediction Errors," *Journal of Accounting Research*, Vol. 22, Studies on Current Econometric Issues in Accounting Research, pp. 34-54.
- Mendenhall, R., 1991, "Evidence on the possible underweighting of earnings-related information," *Journal of Accounting Research*, 29 (1):170–179.
- Shane, P., and P. Brous, 2001, "Investor and (Value Line) analyst underreaction to information about future earnings: the corrective role of non-earnings- surprise information," *Journal of Accounting Research*, Sept. 2001, 387–415.
- Zach, T., 2003, "Inside the accrual anomaly," unpublished dissertation, The University of Rochester.