The Rational Modeling Hypothesis to Explain Analyst Underreaction to Earnings News*

Philip G. Berger
Booth School of Business, University of Chicago, 5807 S. Woodlawn Ave., Chicago, IL 60637

and

Zachary R. Kaplan
Booth School of Business, University of Chicago, 5807 S. Woodlawn Ave., Chicago, IL 60637

June, 2014

Please do not circulate or cite without permission. Comments welcome.

* We appreciate a helpful discussion with Jonathan Rogers, and helpful comments from Judson Caskey, Miles Gietzmann, Urooj Khan, Chad Larson, Volken Muslu, Peter Pope, and participants of workshops at ChicagoBooth, Cass Business School, the University of Houston's Bauer Accounting Symposium and Hong Kong Polytechnic University. We thank Gus DeFranco and Yibin Zhou for sharing their *Institutional Investor* all-star analyst data with us. We both gratefully acknowledge financial support from the University of Chicago Booth School of Business.

Abstract

Analysts publish earnings forecasts with serially correlated errors. We assess rational versus cognitive limitation explanations for analysts' underreaction to earnings news. Institutional investor voting for all-star analyst selections reveals whether these investors prefer analysts to issue forecasts with less serially correlated errors. Consistent with it being potentially rational for analysts to underreact to earnings news, institutional investors appear indifferent to serially correlated errors, although votes do indicate a preference for lower total error in earnings forecasts. Further evidence that analysts seem to be behaving rationally when they underreact to earnings news is that, despite consistently underreacting to their most recent forecast error, analyst reaction does a good job of capturing the extent to which earnings news is transitory versus permanent. We further probe the plausibility of our hypothesis by investigating potential explanations for why investors demand analyst forecasts with serially correlated errors.

I. Introduction

Analysts issue forecasts with serially correlated errors, but why they do so is unresolved. Abarbanell and Bernard (1992) suggest analysts misunderstand the time-series properties of earnings ("irrational modeling"). In contrast, Smith Ready et al. (2006) claim investors prefer forecast errors that have the same sign as the most recent earnings revision and that catering to this preference generates serial correlation in forecast errors ("rational modeling"). We further develop and test the rational modeling explanation and assess it versus the irrational modeling explanation.

To test whether institutional investors prefer forecasts without serially correlated errors, we examine whether analysts selected as institutional investor all-stars better incorporate last quarter's earnings surprise into their forecasts of this quarter's earnings. Performance on surveys of institutional investors has a substantial impact on sell-side analysts' compensation (Groysberg et al. 2011; Brown et al. 2013), so analysts have an incentive to adapt their forecasting strategy to the preferences of investors. An absence of an incentive to incorporate last quarter's forecast error into this quarter's forecast would make issuing forecasts with serially correlated errors rational for analysts, although the reason for investors' indifference to such errors would remain unclear. Alternatively, if an incentive to reduce forecast error autocorrelation exists and some analysts fail to respond to it, analyst underreaction to earnings news may be due to cognitive limitations.

Consistent with analysts being rational in producing serially correlated forecast errors, we do not find that *Institutional Investor (II)* all-stars incorporate more of last quarter's forecast error into this period's beginning-of-quarter earnings forecast (we actually find the all-stars underreact even more than non-all-stars to last quarter's earnings surprise). We find *II* all-stars nevertheless have lower forecast errors, inconsistent with the alternative explanation that investor indifference to forecast error leads to serial correlation in forecast error.

Our tests support the rational modeling explanation, but they do not address why investors prefer (or tolerate) underreaction. We attempt to identify the source of this investor preference by comparing the underreaction in revisions around earnings announcements to underreaction observed at other times. Last quarter's earnings announcement reveals error in the analyst's prior forecast whenever actual earnings deviate from the analyst's forecast. Although analysts choose to revise their forecasts in response to new information throughout the quarter, a distinction between within-quarter revisions and revisions which occur at the earnings announcement is that at least a portion of the revision at the earnings announcement is a response to the analyst's own error. Responding to his own error may involve different costs and benefits for the analyst than responding to other information. Any difference in the extent of analyst underreaction around earnings announcements versus other times may therefore reveal something about analysts' incentives and the investor preferences that drive these incentives. We find analyst revisions underreact more around earnings announcements and that this is true of both *II* all-stars and other analysts.

In section II, we explain in more detail the process analysts use to issue their reports (Weyns et al. 2007) to provide a better understanding of the theories that could explain why analyst underreaction is more intensive at earnings announcements. One potential explanation is based on the observation that analysts produce multiple linked outputs, such as qualitative analysis, financial models and earnings forecasts. Institutional investors report they value many of the other outputs more highly than the earnings forecast (Bagnoli et al. 2005). Publishing forecasts which fully adjust for the information in last quarter's earnings announcement, while the analyst has not fully incorporated the implications of last quarter's forecast error into his qualitative analysis or financial model, would create a research product without internal consistency. Given evidence that individuals find internally inconsistent arguments unappealing, adjusting the forecast (which can be done instantly at little cost) without adjusting the other inputs (which can have larger adjustment costs) may be suboptimal.

In section II, we elaborate on the considerations which may make the publication of forecasts with an identifiable source of error rational. In section III, we provide empirical tests related to the underreaction of all-star and non-all-star analysts.

In section IV, we further assess the rational modeling explanation by examining how analysts respond to inter-temporal variation in the persistence of last quarter's forecast error. We find that analyst forecasts vary strongly with variation in the persistence of last quarter's forecast error, even though analysts consistently underreact to last quarter's error. Explanations for the existence of serial correlation thus need to be consistent with analyst forecasts identifying inter-temporal variation in the persistence of earnings surprises even while consistently under-estimating their magnitude. While rational explanations can readily allow for analyst sophistication in understanding of intertemporal movements in the persistence of earnings surprises, it is unclear how cognitive limitations or other forms of irrationality could explain this aspect of analyst behavior.

In section V, we provide empirical tests examining a variant of the irrational modeling explanation in which the effect of additional analyst experience is to reduce the serial correlation in analyst forecast errors through a learning process. Mikhail et al. (2003; henceforth MWW) provide evidence that experience improves an analyst's understanding of the firm's earnings process and that the analyst underreacts less as he gains experience. Evidence that analysts learn from experience suggests misunderstanding of the time-series process of earnings causes initial serial correlation in forecast error, and that analysts react to the increased availability of information about the time-series by reacting more fully to earnings news. A concern with the

MWW evidence is, however, that they use a time-based identification strategy. A potential problem with time-based identification is that many things change over time, making it difficult to isolate the effect of time on the variable of interest. We present empirical results that suggest MWW's findings are driven by changes in the characteristics of the firms in their sample over time rather than by analysts learning how to reduce their underreaction to earnings news.

In (planned) section VI, to better understand the cause of the serial correlation in forecast error, we examine how the serial correlation in analyst forecast error responds to a reduction in the limits to arbitrage that eliminates post-earnings-announcement drift (PEAD). If serial correlation in analyst forecasts exists because analysts learn from price and price underreacts to earnings information, the elimination of post-earnings announcement drift would eliminate serial correlation in forecast errors. Conversely, if serial correlation exists because analysts gradually update their forecasts in response to earnings information (independent of price) serial correlation will persist even after a reduction to arbitrage costs that eliminates PEAD. We find continued serial correlation in analyst forecasts even after a reduction in the limits to arbitrage that is associated with stoppage of PEAD.

Our finding indicates that forecast error autocorrelation is not caused by PEAD, which also implies that the association between PEAD and forecast error autocorrelation (see, e.g., Shane and Brous (2001)) may be due to common factors affecting both analyst and equity market underreaction to earnings news. While currently beyond the scope of our paper, attempting to identify these common factors could be of considerable interest given that PEAD can be viewed as reflecting serially correlated errors of the marginal investor whereas analyst underreaction can reasonably be viewed as reflecting error autocorrelation of the average investor.

An intriguing implication arises if we assume analysts' forecast behavior reflects the demands of representative investors and also make the stronger assumption that serial correlation in analysts' earnings forecast errors exists because the average investor demands only slow incorporation of earnings news into his valuation estimate. Under these (admittedly strong) assumptions, the continuation of serial correlation even in a post-PEAD world suggests the average investor may continue to make serially correlated errors valuing the firm even if the marginal investor does not. Additional thoughts on our findings are offered in Section VII, which concludes.

II. Theory and Hypotheses

An analyst typically forecasts earnings using a model that expresses earnings as a function of multiple inputs. The analyst typically publishes the model in his report or makes the model available to (some) clients. For instance, an airline analyst may forecast earnings as a function of fuel and labor costs, price per seat mile and seat miles. These inputs can be publicly available, such as average crude oil prices, or forecasts themselves such as price per seat mile. The role of the model is to structure the analysts' thoughts about the earnings process of the firm. Buy-side clients can then take the basic structure and adjust it to represent their views.

The approach the analyst takes to making adjustments to his model and earnings forecast will depend on his incentives. Groysberg et al. (2011) find buy-side client votes on analyst research are used to allocate soft commissions across investment banks and across analysts within a bank. Thus, an analyst seeking to maximize his compensation wants his forecasting method to create a more favorable opinion of his research.

Three pieces of evidence suggest earnings forecast accuracy may not be among the most important attributes to sell-side analysts or their clients. First, *Institutional Investor* asks respondents to the All-America Research Team survey to rank specified attributes in order of importance in assessing the worth of an equity analyst and his/her firm. Bagnoli et al. (2008)

examine the results from these surveys published in the October issue of *Institutional Investor* magazine for the years 1998 - 2003 and report (in their Table 1) the results, which show that "Earnings Estimates" rank anywhere from 12th of 15 attributes (for 2002 and 2003) to 5th of 10 attributes (for 2000). In contrast, "Industry Knowledge" ranks as the top attribute every year during 1998 - 2003 and "Written reports" ranks above "Earnings Estimates" in all of these years. Thus, buy-side users usually place earnings forecast accuracy toward the bottom of the attributes they value, whereas attributes related to qualitative insights are ranked higher.[1] Second, Groysberg et al. (2011) find that earnings forecast accuracy is not correlated with compensation after controlling for institutional investor status. Third, the view that predictive accuracy of any component (earnings, revenue, price, etc.) of an analyst's report is the primary goal of the report is challenged by recent papers. For example, Louth et al. (2010) argue a major role played by sell-side analysts is to describe and predict the impact of possible "jump events" that would tend to sharply raise or lower the price of followed firms. They use Morgan Stanley's analyst database to provide empirical support for their argument and in doing so cite the Weyns et al. (2007) explanation of the database, which emphasizes that analysts should use probabilistic thinking about ranges of uncertainty because "single point estimates only obscure valuable insights and give a false impression of precision."

In light of the preceding evidence about analyst incentives and the preferences of buy-side clients, consider the way an analyst will respond to an earnings surprise. If the analyst incorrectly predicts earnings, the analyst knows that because earnings are persistent his forecast of next quarter's earnings can be expected to contain error. In addition, the model of the followed firm, his written narrative about the firm, his assessment of upside and downside risks

[1] We confirmed that the same pattern of ranking of attributes still persists in the *Institutional Investor* survey by viewing the 2012 ranking at the *Institutional Investor* website and found that "Earnings Estimates" was ranked ninth among 12 attributes, whereas "Industry Knowledge" was ranked first and "Written Reports" sixth.

or "jump events" for the firm, and other aspects of his multidimensional report output need updating. Ideally, the analyst learns the cause of his mistake from the conference call, earnings announcement, or financial statements, and can update his earnings forecast and other outputs with this new information.

In a less ideal situation, the cause of the error may still be unclear to the analyst at the time of the earnings announcement. The analyst knows his forecast contains error, knows the error will affect earnings next quarter, but does not know why. What aspect of the company's operations did he fail to consider properly when he published last quarter's earnings forecast?

Even with uncertainty as to its cause, it would still be simple to publish a forecast which minimizes error after an earnings surprise. The analyst would multiply the earnings surprise by the average persistence of earnings and add this product to his previous estimate. A cost of doing so would be to sacrifice the internal consistency between the earnings forecast and other research outputs. To achieve internal consistency between the model and the earnings forecast, the analyst would have to adjust line items in the model until the earnings forecast equals the output from the model. If the analyst decides after further thought that he changed the wrong line items, the analyst will have to adjust the items he changed initially back to their previous values, when he adjusts the line items actually affected.[2]

Maintaining internal consistency between the model, the earnings forecast and the written report would require the analyst to explain these adjustments to clients. Explaining these machinations imposes processing costs on clients, which may exceed the benefit to the client of a fully updated model.[3] Explaining the analyst's own uncertainty about the cause of the earnings surprise could also undercut the analyst's position as an expert.

[2] Analysts' models frequently contain inputs which are not disclosed by firms in their earnings announcements.
[3] Analysts have incentives to internalize their clients' costs, because of the importance of institutional investor votes in the determination of their compensation (Groysberg et al. 2011).

To summarize, if the clients of analysts value internal consistency between the analyst's multiple research outputs, analysts may not want to fully incorporate the information from last quarter's earnings surprise as quickly as possible. Conversely, if clients view the earnings forecast in isolation, eliminating serial correlation would seem to offer the benefit of lower forecast error at a negligible adjustment cost.

As there are trade-offs for analysts with multiple output reports to fully incorporate last quarter's earnings forecast, it as an empirical question whether clients prefer analysts to fully incorporate last quarter's forecast error. To measure the preference of clients we use the analyst's status as an institutional investor all-star.

H1: All-star analysts will differ from non-all-stars in the extent to which they incorporate information from last quarter's earnings surprise.

Shane and Brous (2001) demonstrate that analysts underreact to news at other times during the quarter and not only to earnings news. However, they do not examine whether the underreaction in revisions around the earnings announcement differs from the underreaction observed at other times.

We argue there are two significant differences between revisions issued around earnings announcements and those issued at other times, which likely drive any difference in underreaction. First, analysts nearly always issue revisions around the earnings announcement, but more selectively issue revisions later in the quarter (Ivkovic and Jegadeesh 2004). Analysts may find it more difficult to fully respond to news around the earnings announcement if they are issuing reports "automatically" at this time. Conversely, given that revisions issued later in the quarter appear to be a self-selected choice by the analyst, analysts who wish to try to respond relatively completely to news may be better able to do so when they self-select the timing of their report. Second, revisions around earnings announcements respond to the earnings announcement, which reveals error in the analyst's previous estimate. Reactions to one's own error may differ from reactions to news. Given the differences between revisions issued around earnings announcements versus later in the quarter, we test whether analyst underreaction differs across revisions made at these different times.

H2: Analysts will differ from in the extent to which their forecast revisions underreact to information around earnings announcements, compared to other times.

III. Analysts' forecasting method

3.1 Summary

We begin investigating the way analysts react to information by reviewing the prior literature and describing how its empirical evidence is consistent with rational modeling. For certain empirical facts, we reference the prior literature, but when we want to address a consideration raised by a subsequent literature we present new analysis based on prior work. We also test H1 and H2. Overall, the findings in this section are consistent with analysts reacting more slowly to earnings-announcement window news than non-earnings-announcement window news.

3.2 Serial correlation in forecast error

To investigate the serial correlation in forecast error we estimate model (1):

$$Fe_{i,t} = \propto + \beta_1 Fe_{i,t-1} + \epsilon_{i,t} \qquad (1)$$

We present estimates of model (1) in Table I. We define forecast error as actual earnings minus forecasted earnings. For all estimates, we treat each analyst-firm-quarter as an observation and cluster standard errors at the firm level. We treat each analyst as an individual observation because we expect the strategy an analyst follows will be related to his own forecast error. In addition, recent evidence on the herding behavior of analysts suggests the average analyst anti-herds, consistent with analysts not intensively using information from the forecasts of competing analysts (Chen and Jiang 2006; Bernhardt et al. 2006).[4]

In Panel A column (1), we present analysis using OLS (Abarbanell and Bernard 1992). This column demonstrates that, if analysts attempt to minimize their mean-squared forecast errors, they do not sufficiently adjust their forecasts in response to last quarter's earnings. In column (2) we present results using median regression, which assumes the analyst attempts to minimize the absolute deviation in his forecast (Gu and Wu 2003). The coefficient of interest, β_1, is still highly significant, inconsistent with the significant coefficient in column (1) being attributable to the particular analyst loss function implicit in the OLS estimation. Finally, column (3) presents results from regressing the sign of this quarter's earnings surprise on the sign of last quarter's earnings surprise. The column (3) results demonstrate that beating the analyst's forecast last quarter shifts the probability upward of beating the forecast again this quarter. Collectively, these three columns of results suggest that last quarter's forecast error shifts the distribution of this quarter's error and that this finding is robust to varying the implicit loss function of analysts via variation in the regression approach used to estimate the relation.

Column (4) presents results including only firms with a positive earnings surprise last quarter and column (5) presents results including only firms with a negative earnings surprise. The positive coefficient in each regression suggests that the serial correlation does not relate to analysts intensively incorporating positive or negative news (Easterwood and Nutt 1999), and is more consistent with a general underreaction to last quarter's earnings surprise. Although columns (4) – (5) present results using OLS, in untabulated analysis we estimate model (1) using median regression and confirm that the results are not driven by the choice of (implicit) loss

[4] We note that the evidence that analysts do not herd in their earnings estimates is consistent with the incentives that underlie our rational modeling hypothesis for analyst underreaction to earnings news. Although the analyst can improve forecast accuracy in a number of ways (including by incorporating information into his forecast from the consensus), investors do not demand this type of forecast revision from analysts because it does not present original insights.

function. We conclude that serial correlation is distinct from the optimistic-pessimistic bias documented in the prior literature (Ke and Yu 2006; Libby et al. 2008; Richardson et al. 2004), as it seems to affect forecasts with both negative and positive earnings surprises last quarter.

In panel B columns (1) – (3), we examine how the serial correlation in analyst forecasts decays over time by presenting this quarter's forecast error regressed on the forecast error from two, three and eight quarters ago. In column (1), we find that the estimated serial correlation declines by 19% at lags of two quarters, compared to one-quarter, suggesting that serial correlation is distinct from the forecast consistency Hilary and Hsu (2013) report. In column (2), we demonstrate that the serial correlation declines another 8% at three lags, but we note that the decay in the serial correlation is far less than we would expect if forecast error followed an AR(1). Finally, column (3) demonstrates that over long lags forecast errors are essentially uncorrelated.

Another aspect of serial correlation which has been noted in prior literature is that analysts incorporate more of last period's forecast error into this period's forecast as the next quarter approaches (Smith Raedy et al. 2006). In columns (4) and (5) of panel B, we present the results of estimating model (1) using only the first revision of the quarter and again using only the final revision. The serial correlation declines by 36% during the quarter. This suggests non-earnings announcement information revealed during the quarter plays a role in correcting analysts' initial underreaction to last quarter's earnings surprise. Non-earnings announcement information refers to information analysts obtain from sources other than earnings announcements. Examples include monthly sales figures, conference calls, investor meetings, and private discussions with management. As detailed by Lang and Lundholm (1993, 1996), these additional sources of corporate disclosure are valuable parts of the information set used by analysts to arrive at their earnings forecasts.

We conclude from the results in this section that any rational theory of the serial correlation in analysts' forecast errors must demonstrate why it is optimal for analysts to gradually react to last quarter's earnings information over the course of the current quarter.

3.3 Revealed preference tests

If investors elect analysts who issue forecasts with certain properties more frequently to be *Institutional Investor* all-stars, then investors reveal a preference for those forecast methods.[5] To test whether investors' prefer forecasts which minimize the error that last quarter's earnings surprise explains, we estimate model (1A), which is model (1) with a dummy set equal to one for *Institutional Investor* all-stars and to zero otherwise ("II"), as well as the interaction of this dummy variable with last quarter's forecast error.

$$Fe_{i,t} = \propto + \beta_1 Fe_{i,t-1} + \beta_2 Fe_{i,t-1} * II + \beta_3 II + \epsilon_{i,t} \qquad (1A)$$

We estimate model (1A) using only analyst forecasts issued within three days of the earnings announcement, to isolate how the analyst responds to the earnings announcement. We find a positive coefficient ($\beta_2 = 0.046, t = 1.14$), suggesting that all-stars incorporate less of last quarter's forecast error into this quarter's forecast of earnings. When we estimate model (1A) using the final forecast of the quarter, in untabulated analysis we find a slightly negative coefficient estimate, suggesting all-stars erase the initial difference and by the end of the quarter better incorporate last quarter's forecast error into this quarter's forecast than non-all-stars ($\beta_2 = -0.007, t = 0.24$).

Inconsistent with the alternative explanation that investor indifference to forecast error leads to serial correlation in forecast error, we also confirm prior findings that all-star analysts have lower forecast errors (citations needed).

3.4 Analyst response to news

[5] We obtain *Institutional Investor* all-star status for all analysts whose reports appear on Investext from 2002 – 2010.

Prior literature has demonstrated that analysts underreact to a variety of information (Lys and Sohn 1990; Abarbanell 1991). Smith Raedy et al. (2006) assert that the general tendency of analysts to underreact to news may be attributable to an incentive for analysts to issue revisions with the same sign as the forecast news disclosed in the revision. Although it is not completely clear why investors would demand analyst underreaction, a theory of analyst underreaction based on investor demand is potentially appealing given the pervasiveness of analyst underreaction and the importance of investor opinion in determining analyst compensation.

To investigate analysts' reaction to news further, we estimate model (2):

$$Fe_{i,t} = \propto + \beta_1 Rev_{i,t} + \epsilon_{i,t} \qquad (2)$$

The dependent variable measures the forecast error before the analyst issues the revision. If analysts fully incorporate information into their revised forecast of earnings the coefficient on the revision will be one. If analysts underreact (overreact) to information, the coefficient estimate on Rev will be larger (smaller) than one, with the deviation from one increasing in the degree of underreaction (overreaction).

In Table II, column (1), we estimate model (2) using the final forecast revision of the quarter. Consistent with analysts underreacting to information on average, we find that the estimate of β_1 is 1.26, suggesting analysts would minimize post-revision forecast error if they increased the magnitude of all revisions by 26%. To compare the reaction to earnings announcement window news to the reaction to non-earnings announcement window news, in column (2) we estimate model (2) using only observations where the analyst revises his forecast of earnings within three days of the earnings announcement. For these observations, the coefficient of interest is 1.5, significantly larger than the coefficient in column (1). Finally, in column (3) we estimate model (2) using the final revision of the quarter for all revisions where the analyst had already issued a revision after the prior quarter's earnings announcement. We assert that these non-earnings announcement window revisions will more likely relate to non-earnings announcement news rather than earnings announcement news. Comparing the coefficients in columns (2) and (3), we find the deviation from one is nearly 2.5 times larger in column (2) than in column (3). We conclude that analysts underreact substantially more to information when the information relates more to earnings announcement news than non-earnings announcement news.

Although the above analysis demonstrates analysts underreact more to information around earnings announcements, it does not imply institutional investors prefer greater underreaction. To examine whether institutional investors prefer greater underreaction in revisions around earnings announcements, in untabulated analysis we estimate model (2) separately for all-stars and non-all-stars, and compare the coefficients. We find that all-star analysts underreact slightly more for revisions issued within three days of an earnings announcement (difference = 5.4%, t=1.58). We find that all-stars underreact slightly less (difference = 3.9%, t= 1.00) for late in quarter revisions, when the analyst has already responded within three days of the earnings announcement. We conclude that investors prefer greater underreaction in revisions issued within three days of an earnings announcement, because the difference in underreaction among all-stars is even larger than the underreaction among non-all-stars.

3.5 Decomposing forecast error

To analyze the relative importance of incorporating last quarter's earnings surprise into this quarter's forecast, compared to incorporating other information, we decompose the analyst's beginning of quarter forecast error into a component related to last quarter's earnings surprise ("lagged quarter forecast error") and a component orthogonal to it ("start of quarter forecast error"). Lagged quarter forecast error is last quarter's forecast error multiplied by the persistence of earnings, or the predicted value from estimates of model (1), where the analyst's beginning of quarter forecast error has been substituted for the end of quarter forecast error. Start of quarter forecast error is beginning of quarter forecast error minus lagged quarter's forecast error, or the residual from the regression described above. It represents forecast error which a very simple econometric analysis cannot identify.

The purpose of this decomposition is two-fold: (1) to better understand the relative contribution to total error of error which can be identified from last quarter's earnings surprise and (2) to better understand all-star and non-all-star analysts' abilities to incorporate the two sources of error into forecasts by the end of the quarter.

We find, in untabulated tests, that the R-squared of regressing the beginning of quarter forecast error on the end of quarter forecast error is 15 percent. Thus, fully incorporating last quarter's forecast error using a naïve strategy where the analyst assumes all earnings surprises have average persistence would reduce beginning of quarter forecast error 15 percent. As the two components of forecast error must equal initial forecast error, the magnitude of start of quarter forecast error equals 85 percent. Overall, these results suggest that start of quarter forecast error is a much larger component of total error than lagged quarter forecast error is, so if the analyst faces costs in incorporating lagged quarter forecast error, he may rationally choose not to do so.

Next, we conduct regression analysis to learn which component of earnings analysts better incorporate into their forecasts. If analysts have a greater ability to address lagged quarter forecast error, than final forecast error should be lower holding all else constant when lagged quarter forecast error is higher. To test this proposition, we take the absolute value of the end of quarter forecast error and regress it on the absolute value of initial forecast error and the absolute value of last quarter's forecast error (model 3):

$$Abs(Final\ Fe)_{i,t} = \propto + \beta_1 Abs(Initial\ Fe)_{i,t} + \beta_2 Abs(Fe)_{i,t-1} + \epsilon_{i,t} \tag{3}$$

We interpret β_1 as the percentage of start of quarter forecast error the analyst incorporates into his forecast. We interpret β_2 as the incremental ability of analysts to identify lagged quarter forecast error relative to start of quarter forecast error. It is important to use absolute rather than signed values in model (3), as signed values would imply reductions of forecast error when forecast error isn't really being reduced. It is also worth noting that regression estimation of model (3) raises some concerns. Specifically, if earnings are heteroskedastic, volatility from the prior quarter is associated with current quarter volatility. In turn, this implies that a high forecast error from the prior quarter can be associated with a high current quarter forecast error even if the analyst does learn from his prior quarter error.

We present estimates of model (3) in Table III. In column (1), we find unconditionally that analysts incorporate 53% of beginning of quarter forecast error into their forecasts by the end of the quarter. In column (2), we include the absolute value of last quarter's forecast error and find that when last quarter's forecast error is higher, final forecast error is also higher. We interpret the positive association between last quarter's forecast error and final forecast error as analysts responding somewhat less to the lagged quarter forecast error component.

To test institutional investors' demands for forecast methods we interact the independent variables in model (3) above with *Institutional Investor* all-star status. If investors elect analysts who issue forecasts with certain properties more frequently to be *Institutional Investor* all-stars, then investors reveal a preference for those forecast methods. This creates an incentive for analysts to adopt such forecast methods.

In column (4), we fully interact model (3) with *Institutional Investor* all-star status. We find that on average all-stars incorporate more of their start of quarter forecast error than non-all-stars ($\beta = -0.066, t = 3.77$). However, we find the opposite results for last quarter's forecast error. When last quarter's forecast error is higher, all-stars have a significantly higher forecast error than non-all-stars ($\beta = 0.088, t = 2.76$).

We conclude that *Institutional Investor* all-star votes suggest institutional investors prefer forecast methods that identify start of quarter forecast error, but that variation in ability to incorporate last quarter's forecast error does not seem to have a great impact on voting. Perhaps this explains why analysts on average incorporate only about half of last quarter's forecast error even though incorporating the other half appears to be so simple.

IV. Responding to news

4.1 Method

In this section, we examine how forecasts of earnings and actual earnings respond to specific news events. Testing the properties of reported and forecasted earnings requires a model of the way in which past earnings and forecasts map into future earnings and forecasts. Previous research (Ball and Bartov, 1996; Markov and Tamayo 2006) assumes quarterly earnings and expectations of quarterly earnings follow an auto-regressive process in fourth differences with a drift.

$$Q_t = \delta + Q_{t-4} + \varphi(Q_{t-1} - Q_{t-5}) + \epsilon_t \qquad (4A)$$

$$Fore(Q_t) = \delta + Q_{t-4} + \varphi(Q_{t-1} - Q_{t-5}) + \epsilon_t \qquad (4B)$$

In equations (4A) and (4B), δ and φ are the true drift and auto-regressive parameters. A potential problem with this model is that analysts forecast a portion of the seasonal change in earnings $(Q_{t-1} - Q_{t-5})$. If analysts' expectations differ systematically over the previously forecasted component of earnings and the surprise component of earnings, failing to decompose the change in earnings into a forecasted and surprise component may affect inferences. Therefore, we decompose $(Q_{t-1} - Q_{t-5})$ into a component related to previously forecasted earnings change $(Fore(Q_{t-1}) - Q_{t-5})$ and earnings surprise $(Q_{t-1} - Fore(Q_{t-1}))$.

$$Q_t = \delta_a + Q_{t-4} + \varphi_{1A}(Q_{t-1} - Fore(Q_{t-1})) + \varphi_{2A}(Fore(Q_{t-1}) - Q_{t-5}) + \epsilon_t \qquad (5A)$$

$$Fore(Q_t) = \delta_b + Q_{t-4} + \varphi_{1B}(Q_{t-1} - Fore(Q_{t-1})) + \varphi_{2B}(Fore(Q_{t-1}) - Q_{t-5}) + \epsilon_t \qquad (5B)$$

We use the same variables to estimate analysts' expectations of earnings and the actual earnings process, with any differences between the actual model and the expectations model resulting in error. Table IV contains estimates of model (5), for both actual earnings and expectations of earnings.[6]

The differences in the coefficient estimates between models (5A) and (5B) suggest that analysts considerably underestimate the persistence of the surprise component ($\varphi_{1A} - \varphi_{1B}$= 0.37), but slightly overestimate the persistence of the forecasted component ($\varphi_{2A} - \varphi_{2B}$= -0.03). These results strongly suggest that the forecasted and surprise components of seasonal earnings change do not have equal effects on next quarter's forecast of earnings. As a result, all subsequent analysis will deviate from the prior literature and estimate model (5) in testing how expectations of earnings differ from actual earnings.[7]

4.2 Inter-temporal variation in earnings persistence

To obtain additional evidence on analysts' ability to identify variation in the persistence of earnings, we examine whether analysts' forecasts incorporate more of last period's earnings change when earnings have more persistence. Figure one (two) plots estimates for each quarter from 1993 - 2009 of the estimated persistence of actual and forecasted earnings surprise (forecast

[6] All variables are winsorized at the first and ninety-ninth percentiles. The inferences are unchanged using data scaled by price, but in many instances the coefficient estimates are different using the two techniques. We present all results using unscaled data because scaling by price results in a few very small firms receiving large weights (having high expected values of variance). To the extent that not all firms receive the same weight in a regression equation, we prefer to assign larger weights to the largest firms in the economy, which make up a greater fraction of the economic activity. In untabulated analysis we find the coefficient estimates are similar using a GLS procedure to weight each observation by an expectation of its variance.

[7] From column (1) of Table IV it appears there may be a small systematic difference between the persistence of the forecasted component and the surprise component of earnings. This suggests either that there is a systematic difference between the earnings innovations analysts do and do not impound into earnings or that firms systematically manage earnings to exceed earnings expectations, and the managed earnings do not recur in the subsequent period. The difference in persistence between the surprise and forecasted components of earnings is not pursued further in this paper.

change), φ_{1A} (φ_{2A}) from model (5A) and φ_{1B} (φ_{2B}) from model (5B). As the figures show, the forecasted persistence moves with the actual persistence, for both earnings surprise and forecast change.

To test how closely the estimates of actual and forecasted persistence covary, we regress estimates of $\widehat{\varphi_{1B}}$ from equation 5B on estimates of $\widehat{\varphi_{1A}}$. The results of this regression are reported in Table V. The coefficient estimate on the actual persistence is 0.52, meaning that forecasts of earnings incorporate a little over half of the inter-temporal variation in the persistence of earnings. The intercept is near-zero, suggesting that all of the variation in the persistence of earnings causes variation in the persistence of forecasted earnings. If analysts followed a naive process in which they consistently adjusted next quarter's forecast by a constant fraction of last quarter's earnings news, forecasted earnings would capture none of the inter-temporal variation in earnings persistence.

The column (1) results suggest analysts integrate substantial information about the time-series variation in earnings persistence into their earnings forecasts. These results would not, however, represent sophistication on the part of analysts if analysts simply adjust their forecasts in response to observable properties of earnings. For instance, some periods contain a greater number of observations with negative earnings and negative earnings have less persistence. If analysts are aware of this, they may correctly forecast variation in the aggregate persistence of earnings without integrating information from sources other than the earnings number. We therefore address the possibility that variation in the persistence of earnings can be predicted by observable time-series variation in the distribution of earnings surprises.

In untabulated analysis, we pool observations across time periods and orthogonalize forecasted earnings and actual earnings with respect to a number of earnings variables (percentage change in revenue, a flag indicating revenue increased, and separate dummies indicating Q1 and/or Q5 was a loss year, as well as these four variables interacted with the two components of earnings) to control for time-series variation in the properties of actual earnings. We then regress the residual forecasted change in earnings and the residual actual change in earnings on earnings surprise and forecast change in each quarter. After eliminating the effect of observable differences from the time-series variation in actual and forecasted earnings persistence, we find almost no change in the coefficient estimate on the variable of interest. We conclude that analysts process non-earnings information in a sophisticated way in producing their earnings forecasts.

V. Do analysts learn over time?

Mikhail, Walther and Willis (2003), henceforth MWW, suggest analysts respond to increased knowledge of the time-series of earnings by incorporating more of last quarter's earnings surprise into this quarter's forecast of earnings. These results suggest that the serially correlated errors in analyst forecasts are undesirable, because analysts respond to the increased accessibility of information by decreasing the serial correlation in their forecast errors. The notion that analysts find serially correlated errors undesirable contradicts the theory advanced in our paper, that analysts rationally (and gradually) adapt their model in response to information disseminated at last quarter's earnings announcement.

To examine the implications of MWW for the theory tested in our paper, we examine their identification strategy. It relies on the passage of time to identify the effect of experience on forecast errors. In particular, MWW compares forecast errors in an earlier period to forecast errors in a later period and attributes any difference between time periods to experience. A threat to the internal validity of this identification strategy is that many firm characteristics change systematically over time and these characteristics may themselves cause analyst forecasts to be more or less autocorrelated. To address this potential threat to the internal validity, we investigate the same question with an identification strategy that we argue better isolates the effect of experience on forecast errors.

MWW hypothesize that more experienced analysts better learn a firm's earnings process and, as a result, issue forecasts with less serially correlated forecast errors. To test this hypothesis, the authors define experience as the number of prior forecasts issued by a unique analyst-firm combination and estimate model (6) below, in which *i* indexes the analyst-firm and *t* indexes time:

$$Fe_{i,t} = \propto + \rho_1 Fe_{i,t-1} + \rho_2 Exp + \rho_3 Exp * Fe_{i,t-1} + \epsilon_{i,t} \qquad (6)$$

The authors find a significantly negative estimate for ρ_3 and conclude from this that experience reduces the serial correlation in forecast error. While the finding is consistent with experience reducing forecast error, firms that have been followed by analysts for a long period of time are necessarily surviving firms. These firms' information environments may have evolved over time in a way that would affect the serial correlation in forecast error for the average analyst. In particular, surviving firms are larger and more profitable than the average firm. In untabulated analysis, we find that both of these characteristics are significantly negatively associated with the serial correlation in forecast error. As a result, it is unclear if experience causes the decrease in the serial correlation of forecast error, or if the changing firm characteristics affect the information environment in a way that causes all analysts (regardless of experience level) to issue forecasts with less serially correlated errors. To control for any possible change in firm characteristics, we match experienced analysts to less experienced analysts following the same firm, and compute the difference in their experience levels. Then we estimate the following regression, in which *i* indexes a firm followed by the matched pair of analysts and *t* indexes time:

$$Fe_{i,t} = \propto + \rho_1 Fe_{i,t-1} + \rho_2 Diff\_Exp + \rho_3 Diff\_Exp * Fe_{i,t-1} + \epsilon_{i,t} \qquad (7)$$

The coefficient of interest is ρ_3, which measures the effect of experience on the serial correlation of the forecast error (we set *Diff_Exp* to zero for the less experienced analyst).

MWW obtain their data from a different database and for an earlier time period. When we replicate their study using IBES data from 1992 - 2010, we find a significantly negative coefficient estimate, consistent with their findings.

When we implement the matching procedure that aims to control for changing firm characteristics, we obtain a small and insignificantly positive coefficient estimate (ρ_3=0.0007, t=0.25). This regression has considerable power because there are over 58,000 unique firm-quarters for which analysts with different levels of experience issue a forecast. Matching by firm eliminates the effect of changing firm characteristics on the coefficient of interest and isolates the effect of experience on the serial correlation in forecast error. We find that the difference between the coefficient estimates obtained estimating equations (6) and (7) is statistically significant at the 1% level, suggesting that the two designs are unlikely to be measuring the same effect. We argue the matched estimate provides a better measure of the effect of experience on the serial correlation of analyst forecast errors, because it controls for all unobserved firm characteristics. Our finding suggests the results MWW report differ from the matched sample results because, as the analysts gain experience, the firms they continue to follow change and the change in firm characteristics drives the statistically significant coefficient estimate on experience that they document.

VI. Impact of elimination of PEAD on forecast error serial correlation

The work on this section is in process and not yet ready for the NYU Summer Camp draft of the paper.

VII. Conclusion

Financial analysts use models to help predict earnings. Financial models generate earnings predictions by assuming a relation between a series of known inputs and next period's earnings. The economic relation between these inputs and earnings is uncertain and likely modeled with error. Thus, a rational Bayesian seeking to minimize forecast error would adjust his posterior expectation of the relation between the inputs and the outputs each time he observes an earnings realization which differs from the forecast.

We assert that analysts instead adjust their models gradually as their thinking about the economics of the company evolves. We assert gradual adjustment may be incentive compatible for the analyst, as institutional investors do not seem to reward analysts who more fully respond to last quarter's surprise with *Institutional Investor* all-star status. We also demonstrate that analyst revisions capture variation in the persistence of last quarter's earnings surprise, consistent with analysts rationally modeling rather than naively responding to the earnings surprise.

We note that rational modeling may also explain why managers' forecast errors are serially correlated (Gong et al. 2011), as managers often produce forecasts using models as well. Anecdotal evidence suggests that accountants and board members occasionally audit the forecasts managers use, suggesting managers may have difficulty deviating from the forecasts their models produce.

Overall, we hypothesize that serially correlated forecast errors may arise out of a rational system where analysts choose a number of inputs to include in a model to forecast earnings.

References

[1] J.S. Abarbanell. Do analysts earnings forecasts incorporate information in prior stock price changes? *Journal of Accounting and Economics*, 14(2):147-165, 1991.

[2] J.S. Abarbanell and V.L. Bernard. Tests of analysts overreaction-underreaction to earnings information as an explanation for anomalous stock price behavior. *The Journal of Finance*, 47(3):1181-1207, 1992.

[3] Mark Bagnoli, Susan G Watts, and Yong Zhang. Reg-fd and the competitiveness of all-star analysts. *Journal of Accounting and Public Policy*, 27(4):295_316, 2008.

[4] R. Ball and E. Bartov. How naive is the stock market's use of earnings information? *Journal of Accounting and Economics*, 21(3):319-337, 1996.

[5] D. Bernhardt, M. Campello, and E. Kutsoati. Who herds? *Journal of Financial Economics*, 80(3):657-675, 2006.

[6] Q. Chen and W. Jiang. Analysts weighting of private and public information. *Review of Financial Studies*, 19(1):319_355, 2006.

[7] J.C. Easterwood and S.R. Nutt. Inefficiency in analysts earnings forecasts: systematic misreaction or systematic optimism? *The Journal of Finance*, 54(5):1777-1797, 1999.

[8] G. Gong, L.Y. Li, and J.J. Wang. Serial correlation in management earnings forecast errors. *Journal of Accounting Research*, 49(3):677-720, 2011.

[9] B. Groysberg, P.M. Healy, and D.A. Maber. What drives sell-side analyst compensation at high-status investment banks? *Journal of Accounting Research*, 49(4):969-1000, 2011.

[10] Zhaoyang Gu and Joanna Shuang Wu. Earnings skewness and analyst forecast bias. *Journal of Accounting and Economics*, 35(1):5-29, 2003.

[11] Gilles Hilary and Charles Hsu. Analyst forecast consistency. *The Journal of Finance*, 68(1):271-297, 2013.

[12] Amy P. Hutton, Lian Fen Lee, and Susan Z. Shu. Do managers always know better? The relative accuracy of management and analyst forecasts. *Journal of Accounting Research*, 50(5):1217-1244, 2012.

[13] Zoran Ivkovic and Narasimhan Jegadeesh. The timing and value of and recommendation revisions. *Journal of Financial Economics*, 73(3):433-463, 2004.

[14] Bin Ke and Yong Yu. The effect of issuing biased earnings forecasts on analysts access to management and survival. *Journal of Accounting Research*, 44(5):965-999, 2006.

[15] Mark H. Lang and Russell J. Lundholm. Cross-sectional determinants of analyst ratings of corporate disclosures. *Journal of Accounting Research*, 31(Autumn):246-271, 1993.

[16] Mark H. Lang and Russell J. Lundholm. Corporate disclosure policy and analyst behavior. *The Accounting Review*, 71 (October):467-492, 1996.

[17] Robert Libby, James E. Hunton, Hun-Tong Tan, and Nicholas Seybert. Relationship incentives and the optimistic-pessimistic pattern in analysts forecasts. *Journal of Accounting Research*, 46(1):173-198, 2008.

[18] R.J. Louth, P. Joos, S.E. Satchell and G. Weyns. Understanding analyst forecasts. *The European Journal of Finance*, 16(2):97-118, 2010.

[19] Thomas Lys and Sungkyu Sohn. The association between revisions of financial analysts earnings forecasts and security-price changes. *Journal of Accounting and Economics*, 13(4):341-363, 1990.

[20] Stan Markov and Ane Tamayo. Predictability in financial analyst forecast errors: Learning or irrationality? *Journal of Accounting Research*, 44(4):725-761, 2006.

[21] Michael B. Mikhail, Beverly R.Walther, and Richard H. Willis. The effect of experience on security analyst underreaction. *Journal of Accounting and Economics*, 35(1):101-116, 2003.

[22] Jana Smith Raedy, Philip Shane, and Yanhua Yang. Horizon-dependent underreaction in financial analysts earnings forecasts. *Contemporary Accounting Research*, 23(1):291-322, 2006.

[23] Scott Richardson, Siew Hong Teoh, and Peter D. Wysocki. The walk-down to beatable analyst forecasts: The role of equity issuance and insider trading incentives. *Contemporary Accounting Research*, 21(4):885-924, 2004.

[24] Philip Shane and Peter Brous. Investor and (Value Line) analyst underreaction to information about future earnings: the corrective role of non-earnings-surprise information. *Journal of Accounting Research*, 39(2):387-404, 2001.

[25] G. Weyns, J.L. Perez and V. Jenkins. Risk-reward views: unlocking the full potential of fundamental analysis. *Morgan Stanley Global Research Paper*, December, 2007.

Appendix A: Variable Definitions

This appendix describes each variable used in our study. All data are from I/B/E/S for the years 1983 – 2010.

(i = firm; t = time of the fiscal period end, t' = time of forecast (if different than end of fiscal period); j = analyst)

Variable	Description	Formula
Forecast Error (also called final forecast error)	Actual earnings minus the last forecast of earnings issued by the analyst. Forecast of earnings must be issued **after** last quarter's earnings announcement.	$= Actual_{i,t} - Forecast_{itj}$
Initial forecast Error (also called final forecast error)	Actual earnings minus the last forecast of earnings issued by the analyst. Forecast of earnings must be issued **before** last quarter's earnings announcement.	$= Actual_{i,t} - Forecast_{i,t,t'-1,j}$
Forecast Change	The difference between earnings five quarters ago and the analyst's forecast of last quarter's earnings	$= Forecast_{i,t-1,j} - Actual_{i,t-5}$
Actual Change	The difference between earnings four quarters ago and this quarter's actual earnings	$= Actual_{i,t} - Actual_{i,t-4}$
Lagged quarter forecast error	The portion of this quarter's forecast error implied by last quarter's forecast error. This is equal to the last quarter's forecast error multiplied by the coefficient estimate obtained from regressing initial forecast error on last quarter's forecast error.	
Start of quarter forecast error	The initial forecast error minus the lagged quarter forecast error.	= Initial FE - Lagged Quarter FE
Revision	Forecast of earnings minus the same analyst's previous forecast of earnings	
Last quarter's forecast error (earnings surprise)	Actual earnings for last quarter minus the last forecast of earnings issued by the analyst. Forecast of earnings must be issued **after** the earnings announcement from two quarters ago.	$= Actual_{i,t-1} - Forecast_{i,t-1,j}$
Institutional Investor All-star	A flag set equal to one if the analyst was voted as an *Institutional Investor* all-star in either the year the fiscal period ended or the year after the fiscal period ended. We have this variable populated for all Investext analysts between 2002 – 2010.	
Difference in revisions	The signed difference in the revisions for the firm with the larger revision in absolute value, and zero for the firm with the smaller revision in absolute value.	$= \Delta \mathrm{Re}\, v = rev_{i,t} - rev_{i,t}$

Table I: Effect of prior period's forecast error on this period's forecast error -- This table presents the results of regressing one quarter's forecast error on a previous quarter's forecast error. For all columns in Panel A and columns (1) - (2) in panel B, we include all observations where the analyst issued a forecast of this quarter's earnings announcement after last quarter's earnings announcement for the same firm for three consecutive quarters. For column (3) in panel B, we require the analyst issue a forecast of this quarter's earnings announcement after last quarter's earnings announcement, and does similarly for the earnings announcement eight quarters prior. For columns (4) and (5) of panel B, we require the analyst issue two forecasts after last quarter's earnings announcement and a forecast of last quarter's earnings.

Panel A: All columns present the results of regressing this quarter's final forecast error on last quarter's forecast error. Column (1) presents OLS, Column (2) presents median regression and column (3) takes the sign of both the independent and dependent variable before performing OLS. Columns (4) and (5) include only observations with prior quarter forecast errors which are positive and negative respectively ("Restriction on independent variable").

Panel B: In panel B, all results use OLS and include all observations, but differ as far as the time the analyst issued the revision. In columns (1), (2) and (3) we present results where the prior quarter's forecast error is taken from the quarter two, three and eight quarters prior, respectively. In column (4) we present results using the forecast error calculated using the first forecast an analyst issues during the quarter.

Panel A: variation in loss function and sign

	(1) Forecast Error$_{t=0}$	(2) Forecast Error$_{t=0}$	(3) Forecast Error$_{t=0}$	(4) Forecast Error$_{t=0}$	(5) Forecast Error$_{t=0}$
Intercept	0.001	0.008	0.227	0.009	0.003
	(0.10)	(360.63)	(49.73)	(4.54)	(0.79)
Lagged forecast error	0.351	0.250	0.210	0.250	0.383
	(16.61)	(87.85)	(55.48)	(6.86)	(12.55)
Restriction on independent variable	NONE	NONE	NONE	Positive	Negative
Type of regression	OLS	Median	Signed OLS	OLS	OLS
N	427,672	427,672	427,672	252,266	129,355
R-squared	0.093	--	0.044	0.026	0.114

Panel B: variation in time of forecast

	(1) Forecast Error$_{t=0}$	(2) Forecast Error$_{t=0}$	(3) Forecast Error$_{t=0}$	(4) Forecast Error$_{t=0}$	(5) Forecast Error$_{t=0}$
Intercept	0.000	0.000	0.002	-0.054	0.000
	(0.16)	(0.30)	(1.23)	(-17.46)	(-0.08)
Lagged forecast error	0.286	0.257	0.067	0.540	0.347
	(12.48)	(11.61)	(6.49)	(12.14)	(12.74)
Observation	Final	Final	Final	First	Final
Number of lags	2	3	8	1	1
N	427,672	427,672	346,439	242,239	242,239
R-squared	0.054	0.041	0.013	0.083	0.075

See Appendix A for variable definitions. Reported below the coefficients are t-statistics. All t-statistics, except the t-statistics reported in panel A, column (2) are clustered by firm. All variables are Winsorized at the first and ninety-ninth percentiles.

Table II: Correlation between revisions and forecast error -- This table presents the results of regressing the analyst's initial forecast error on a revision the analyst issues during the quarter. Column (1) reports results for the final revision the analyst issues during the quarter. Column (2) reports results for all analysts who issue an earnings forecast within three days of the prior quarter's earnings announcement. Column (3) reports results for the final revision of the quarter for all analysts who issued a revision after last quarter's earnings announcement prior to their final revision.

	(1) Initial FE	(2) Initial FE	(3) Initial FE
Intercept	0.008	-0.008	0.010
	(5.23)	(-3.78)	(5.87)
Revision	1.266	1.506	1.203
	(60.17)	(53.40)	(47.54)
Sample Selection	Final Forecast	EAD Window	2nd Revision
N	715,008	429,562	328,274
R-squared	0.402	0.286	0.454

See Appendix A for variable definitions. Reported below the coefficients are t-statistics. All t-statistics are clustered by firm. All variables are Winsorized at the first and ninety-ninth percentiles.

Table III: Relationship between initial forecast error and final forecast error -- This table presents results of regressing the final forecast error of the quarter on the initial forecast error, to estimate the percentage of forecast error analysts identify during the quarter. We include all observations where (1) the analyst issues a forecast of this quarter's earnings both before and after last quarter's earnings announcement and (2) the analyst issues a forecast of last quarter's earnings after the announcement of earnings two quarters ago. In columns (1) and (3), both the dependent variable and initial forecast error are first regressed on last quarter's forecast error before we take their absolute value and regress them on each other (model 3). In columns (2) and (4) we use raw values. In columns (3) and (4) all independant variables are interacted with all-star status and we only include observations for which we have all-star status available.

	(1) Abs(Final FE)	(2) Abs(Final FE)	(3) Abs(Final FE)	(4) Abs(Final FE)
Intercept	0.005	0.001	0.006	0.003
	(3.27)	(1.01)	(3.02)	(2.14)
Absolute Initial FE	0.543	0.512	0.551	0.515
	(45.62)	(45.18)	(33.67)	(32.55)
Absolute Last Quarter FE		0.001		0.152
		(1.01)		(8.29)
II Status	0.543	0.512	0.004	-0.001
	(45.62)	(45.18)	(1.54)	(-0.36)
ABS(Init_FE)*II			-0.053	-0.066
			(-3.01)	(-3.77)
ABS(Last_FE)*II				0.088
				(2.76)
N	593,629	593,629	204,059	204,059
R-squared	0.642	0.715	0.637	0.703

See Appendix A for variable definitions. Reported below the coefficients are t-statistics. All t-statistics are clustered by firm. All variables are Winsorized at the first and ninety-ninth percentiles.

Table IV: Effect of earnings surprise and forecast change components of current quarter earnings -- This table contains OLS regressions of actual change in earnings (column 1), forecasted change in earnings (column 2) and forecast error (column 3) on last quarter's forecasted change in earnings and forecast error. All observations are taken from the split-adjusted I/B/E/S detail file between 1993 and 2009.

We compute the forecast of earnings in both this quarter (dependent variable) and the prior quarter (independent variable) as the average forecast from all analysts who issue a forecast in each quarter. We only include forecasts issued after the prior quarter's earnings announcement.

Forecast Change is the average forecast of earnings minus actual earnings four quarters ago.

Actual Change is this quarter's earnings minus actual earnings four quarters ago.

	(1) Actual Change	(2) Forecast Change	(3) Forecast Error
Intercept	-0.008	0.002	-0.008
	(-9.95)	(2.86)	(-10.85)
Earnings Surprise	0.523	0.146	0.370
	(31.72)	(9.23)	(20.61)
Forecast Change	0.578	0.607	0.000
	(48.78)	(56.43)	(-0.02)
N	171,338	171,338	171,338
R-squared	0.256	0.324	0.102

See Appendix A for variable definitions. Reported below the coefficients are t-statistics. All t-statistics are clustered by firm. All variables are Winsorized at the first and ninety-ninth percentiles.

Table V. Effect of Actual Persistence of Earnings on the Forecasted Persistence of Earnings -- This table contains OLS regressions of the estimated actual persistence of earnings on the estimated forecasted persistence of earnings for each quarter from 1993 - 2009. The dependant variable, the estimated actual persistence of earnings is β_1 in the following regression, estimated separately for each quarter: Actual Earnings Change = $\alpha+\beta_1$*Earnings_Surprise+β_2*Forecast_Change + ε. The independent variable, the estimated forecast persistence of earnings is β_1 in the following regression, estimated separately for each quarter: Forecasted Earnings Change = $\alpha+\beta_1$*Earnings_Surprise+β_2*Forecast_Change + ε. We define forecasted change as the forecast of last quarter's earnings minus actual earnings reported the same quarter the prior year ($Fore_{t-1}$- $Actual_{t-5}$). We compute the forecast of last quarter's earnings as the average of the final forecast on the I/B/E/S detail file. We exclude all forecasts issued before the previous quarter's earnings announcement from the consensus. We also exclude firms without earnings information on COMPUSTAT. We define earnings surprise as the difference between last quarter's actual earnings and last quarter's forecast ($Actual_{t-1}$- $Fore_{t-1}$). We define the actual change in earnings as the actual earnings reported this quarter minus the actual earnings reported the same quarter the prior year ($Actual_t$- $Actual_{t-4}$). We define the forecasted earnings change as the consensus forecast of earnings this quarter minus the actual earnings reported the same quarter the prior year ($Fore_t$- $Actual_{t-4}$).

	(1) Estimated Actual Persistence
Intercept	-0.110
	(-2.64)
Estimated Actual Persistence	0.524
	(7.28)
Reported Results	2nd Stage
N	67
R-squared	0.449

See Appendix A for variable definitions. Reported below the coefficients are t-statistics. In the first stage variables are Winsorized at the first and ninety-ninth percentiles.

Figure 1. Time-Series Variation in the Persistance of Actual and Forecasted Earnings Surprise

The blue series, the estimated actual persistence of earnings is β_1 in the following regression, estimated separately for each quarter: Earnings Change = $\alpha+\beta_1$*Earnings_Surprise+β_2*Forecast_Change + ε. The red series, the estimated forecast persistence of earnings is β_1 in the following regression, estimated separately for each quarter: Forecast Change = $\alpha+\beta_1$*Earnings_Surprise+β_2*Forecast_Change + ε.
All regressions were estimated using the last IBES consensus forecast prior to the earnings announcement date. Estimating the regression requires actual earnings data for Q_t, Q_{t-1} and Q_{t-4} and forecasted earnings data for Q_t and Q_{t-1}.

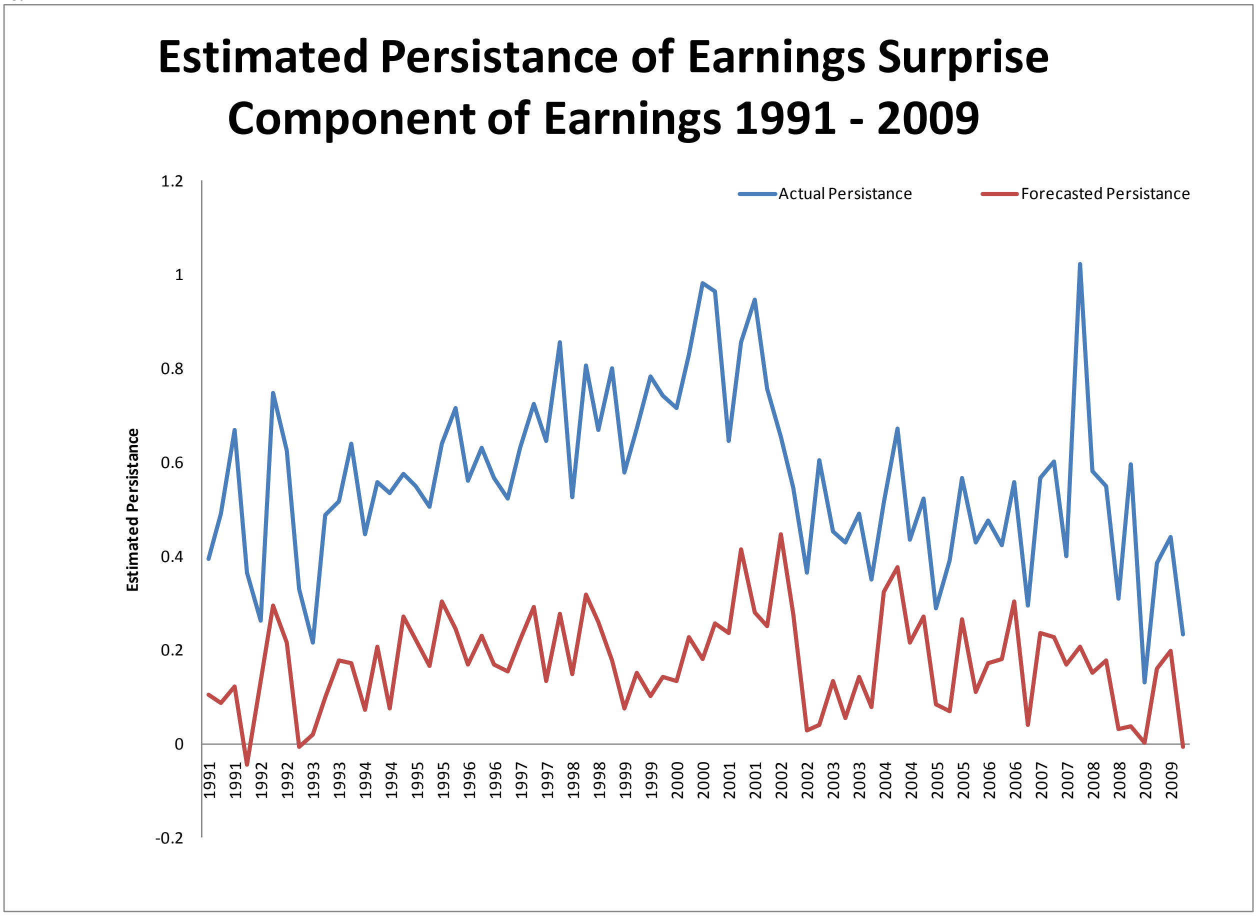

Figure 2. Time-Series Variation in the Persistance of Actual and Forecasted Forecast Change

The blue series, the estimated actual persistence of earnings is β_2 in the following regression, estimated separately for each quarter: Earnings Change = $\alpha+\beta_1$*Earnings_Surprise+β_2*Forecast_Change + ε. The red series, the estimated forecast persistence of earnings is β_2 in the following regression, estimated separately for each quarter: Forecast Change = $\alpha+\beta_1$*Earnings_Surprise+β_2*Forecast_Change + ε.

All regressions were estimated using the last IBES consensus forecast prior to the earnings announcement date. Estimating the regression requires actual earnings data for Q_t, Q_{t-1} and Q_{t-4} and forecasted earnings data for Q_t and Q_{t-1}.