TI 2013-121/III
Tinbergen Institute Discussion Paper

Forecasting Earnings Forecasts

Bert de Bruijn
Philip Hans Franses

Erasmus School of Economics, Erasmus University Rotterdam, Econometric and Tinbergen Institutes.

Forecasting Earnings Forecasts*

Bert de Bruijn[a,b,†]
Philip Hans Franses[a,b,‡]

[a] Econometric Institute, Erasmus School of Economics,
Erasmus University Rotterdam, the Netherlands
[b] Tinbergen Institute, the Netherlands

August 22, 2013

We analyze earnings forecasts retrieved from the I/B/E/S database concerning 596 firms for the sample 1995 to 2011, with a specific focus on whether these earnings forecasts can be predicted from available data. Our main result is that earnings forecasts can be predicted quite accurately using publicly available information. Second, we show that earnings forecasts that are less predictable are also less accurate. We also show that earnings forecasters who quote forecasts that are too extreme need to correct these as the earnings announcement approaches. Finally, we show that the unpredictable component of earnings forecasts can contain information which we can use to improve the forecasts.

Keywords: Earnings Forecasts; Earnings Announcements; Financial Markets; Financial Analysts.
JEL classifications: G17, G24, M41.

*Wharton Research Data Services (WRDS) was used in preparing this paper. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.

†Econometric Institute, Erasmus School of Economics, PO Box 1738, NL-3000 DR Rotterdam, The Netherlands; phone: +3110 4088902; email: debruijn@ese.eur.nl (corresponding author)
‡Econometric Institute, Erasmus School of Economics, PO Box 1738, NL-3000 DR Rotterdam, The Netherlands; phone: +3110 4081273; fax: +3110 4089162; email: franses@ese.eur.nl

1. Introduction

Earnings forecasts can provide useful information for investors. When investors in part rely on such forecasts, it is important to have more insights into how such earnings forecasts are created. A key research subject therefore concerns the drivers of the forecasts of earnings analysts. Such knowledge is relevant as the part that can be predicted from factors that are also observable to the end user of the forecast might not be the most interesting part of an earnings forecast. Indeed, it is the unpredictable component of the earnings forecast that amounts to the forecaster's true added value, based on latent expertise and domain-specific knowledge. As a consequence, in our perspective, the evaluation of the quality of earnings forecasts should mainly focus on that unpredictable part, as that is truly the added value of the professional forecaster.

There is much literature on the properties and accuracy of earnings forecasts, but there is no research that focuses on the prediction of such forecasts. Which variables are the most relevant drivers of earnings forecasts? Can we use the unpredictable part of the forecast to improve forecasts? In this paper we answer these questions using appropriate models. We apply these models to the earnings forecasts for a large number of firms which constitute the S&P500. Using this large sample of firms, we are confident to draw a few generalizing conclusions.

A key predictor of the earnings forecasts appears to be the average of all available earnings forecasts concerning the same forecast event. As an example, consider a forecaster who has produced his most recent forecast some period ago. If in the meantime information has been provided on the firm that has driven the forecasts of all (other) forecasters down, this forecaster will also on average produce a lower-valued forecast than before. A second predictor is the most recent difference between the individual forecaster's forecast and the average of the available contemporaneous forecasts. For example, a forecaster who previously was more optimistic about the earnings of a particular firm can be expected to persist in quoting above-average values. Other important conclusions that we draw from the data are that more unpredictable forecasts tend to be less accurate, and that the unpredictable component of the forecast can be used to improve the forecast. All in all, we document that earnings forecasts are quite predictable from data that are also available to the end user.

The outline of our paper is as follows. In Section 2 we develop several hypotheses to guide our empirical analysis, and we base these hypotheses on available studies, reviewed in Section 2. In Section 4 we discuss the data and in Section 5 we present our results. Section 6 concludes and provides various avenues for further research.

2. Literature review

Earnings forecasts have been the topic of interest for many researchers. For an extensive discussion of research on earnings forecasts in the period 1992-2007, see Ramnath et al. (2008). For earlier overviews we refer to Schipper (1991) and Brown (1993).

One stream of earnings forecasts research has focused on relationships between forecast performance and forecaster characteristics. Performance can be measured by forecast accuracy and forecast impact on stock market fluctuations. The characteristics of these performance measurements have been related to timeliness (Cooper et al., 2001; Kim et al., 2011), the number of firms that the analyst follows (Kim et al., 2011; Bolliger, 2004), the firm-specific experience of the analyst (Bolliger, 2004), age (Bolliger, 2004), the size of the firm being followed and of the firm at which the analyst works (Kim et al., 2011; Bolliger, 2004), and whether the analyst works individually or in a team (Brown and Hugon, 2009).

Another stream of research concerns the value of an earnings forecast and how it is related to what other analysts do. In particular, herding behavior is considered, which occurs when forecasters produce forecasts that converge towards the average of those of the other forecasters. There has been an effort to categorize earnings forecasters into two groups, corresponding to leaders and followers or to innovators and herders (Jegadeesh and Kim, 2010; Clement and Tse, 2005). This is interesting as different types of forecasters might consult different amounts of information which in turn can be useful for investors to incorporate into their investment decisions. A leading or innovating forecaster might on average be more useful to follow than a herding forecaster. This does not directly imply that leading forecasts are also more accurate, as accuracy and the type of forecast are not necessarily related. In fact, it has been documented that aggregation of leading forecasts is a fruitful tactic to produce accurate forecasts (Kim et al., 2011).

Recently, Clement et al. (2011) have studied the effect of stock returns and other analysts' forecasts on what analysts do. In contrast to Jegadeesh and Kim (2010) and Clement and Tse (2005), Clement et al. (2011) do not consider categorizing the forecasters into different groups. Instead, they consider how the first forecast revision after a forecast announcement is affected by how the stock market and other analysts have reacted to that forecast announcement. Landsman et al. (2012) also look at how earnings announcements affect the stock market, where these authors focus on how mandatory IFRS adoption has influenced this effect. Sheng and Thevenot (2012) propose a new earnings forecast uncertainty measure, which they use to demonstrate that forecasters focus more on the information in the earnings announcement if there is high uncertainty in the available set of earnings forecasts.

In sum, earnings forecasts have been studied concerning their performance and a few of their potential drivers. In this paper we extend the knowledge base by considering many more drivers of earnings forecasts, while we pay specific attention to the value of the unpredictable component of earnings forecasts.

3. Hypotheses

To guide our empirical analysis, we put forward several useful questions and hypotheses. We start with a general question, relate this question to previous research, and at the same time we introduce relevant notation and definitions.

Consider, for one analyst and one year, a series of earnings forecasts. This series is a single time series, with irregularly spaced observations. Time series usually exhibit serial correlation with lagged observations, which allows to forecast a future observation of the series using previous observations. This suggests to use the previous forecast of the analyst as a predictor for the current forecast. In practice this is not straightforward, as oftentimes analysts only produce a few forecasts within a year and also the firm-specific data may change over time.

Ideally, we would want to have a daily-observed series of forecasts, but this is not the case in practice. Instead, we may use the forecasts of other individual analysts, who might have produced a forecast in the recent days. This would incorporate recent information as long as the total group of analysts is active. The herding literature (Jegadeesh and Kim, 2010; Clement and Tse, 2005) suggests that there is reason to believe that individual analysts can purposely follow the aggregate forecast, also because their forecasts are driven by a common factor, as the analysts share most of the information on the state of the firm.

Forecasters can be influenced by optimism (Easterwood and Nutt, 1999). This can result in an analyst giving a higher earnings forecast than other analysts. If this optimism is persistent over time, it can be used to forecast a next forecast. For example, an analyst who previously forecasted the earnings at a level higher than other analysts did at that time, can be expected to persist with above-average values. One reason for this might be a deliberate strategic approach to get attention from certain firms (Laster et al., 1999).

We expect that by relying on both predictors, which are both well documented as relevant for earnings forecasts, it is possible to achieve a good forecasting accuracy when forecasting earnings forecasts. Because of this, we arrive at the following hypothesis.

Hypothesis 1 *Individual earnings forecasts can be forecasted well by using publically available information. This information concerns (1) the average of the forecasts of all analysts and (2) the difference between the previous forecast of the analyst and the average of the available forecasts at that time.*

To test this hypothesis, we will regress the earnings forecast on the two explanatory variables amongst several other possible candidates. We can use the regression results across different firms to evaluate how much the effect of the different variables varies across firms.

Our second question concerns the relation between how accurate one can predict a particular earnings forecast and the forecast error associated with this earnings forecast. For example, one could argue that forecasts which cannot be predicted well apparently contain new analyst-specific information and this information might improve the forecast to have a smaller forecast error. This assumes that the analyst who makes the unpredictable forecast is well capable of correctly interpreting such new information for its impact in the future. On the other hand, it could also be argued that forecasts which are fairly unpredictable perform worse, as analysts may misinterpretate the information or have other reasons to forecast an unexpected forecast. The second option is supported by Kim et al. (2011), who document that aggregation of leading forecasts is a fruitful strategy to produce accurate forecasts. Lamont (2002) shows that this also holds for macroeconomic forecasts, as he finds that bolder forecasts turn out to be less accurate. Following these two studies, we hypothesize:

Hypothesis 2 *The forecast error is larger in size when earnings forecasts deviate more from their predictable component.*

We test this hypothesis by regressing the log squared forecast error on several transformations of the residual of the forecasting equation used for Hypothesis 1.

Our third question focuses on the variation in earnings forecasts and how that is correlated with the realization of the earnings. For example, in case of a high correlation between forecasts and realizations, it could well be that the forecasts fluctuate more than the realizations do. In this case, the forecast performance would improve if the forecasts were adjusted towards an average value. The reverse could also occur. The analysts produce forecasts that are too close to the average, even though publicly-available information suggests that more extreme forecasts would be relevant. If analysts do not only care about forecasting accuracy, but also about how much attention their forecasts get, they have an incentive to overreact (Laster et al., 1999). In sum, we hypothesize:

Hypothesis 3 *Earnings forecasts tend to be too extreme and need to be corrected towards average values to improve accuracy.*

To test the above hypothesis, we consider the following equation:

$$Actual = \beta_0 + \beta_1 Forecast + rest\, term \quad (1)$$

If $\beta_0 = 0$ and $\beta_1 = 1$, the forecasts are on average right. If $\beta_1 > 1$, then the analysts are too timid with their forecasts, while they quote too extreme values if $\beta < 1$. We will estimate this equation for different forecast horizons.

Our fourth research question focuses on whether we can seperate the forecasts into two parts and evaluate these parts concerning their forecasting contribution. The two parts are constructed using our forecast of the forecast. One part is the predictable component of the forecast, and the other part is what is left, that is, the unexplained part or residual of the forecast. This could lead to several interesting situations. For example, it could be that both parts are equally important in their contribution to the forecasting quality. It could however also be that one contains more information than the other. As aggregating forecasts can lead to superior forecasts (Kim et al., 2011), we expect the predictable part of the forecast to be most important. This does not necessarily mean that the unpredictable component contains no additional information. In fact, we hypothesize:

Hypothesis 4 *The unpredictable component of the earnings forecast can be used to improve forecast performance relative to using only the predictable component of the forecast. This improvement will be largest for the longest forecast horizon.*

We test this hypothesis by estimating an equation similar to (1):

$$Actual = \beta_0 + \beta_1 FOF + \beta_2 ROF + rest\ term \tag{2}$$

in which FOF stands for Forecast of Forecast and ROF denotes the Residual of Forecast, the two parts that together constitute the original earnings forecast. The regression results can be used to compute the contribution of FOF and ROF to the total forecasting power. Again, we estimate this equation for different forecast horizons of the earnings forecasts.

4. Data and sample selection

Data has been collected from WRDS[1], using the I/B/E/S database for the analyst forecasts and the CRSP data for the stock prices and returns.

[1] http://wrds-web.wharton.upenn.edu/wrds/

Concerning the earnings forecasts, we have collected data for all firms which have been part of the S&P500 during the period 1995 to 2011. This amounts to 658 firms due to mergers, name changes and entry and exit of firms. We focus on the within-year yearly earnings forecasts, that is, the forecasts that are produced to forecast the earnings of the current year. The structure of the data is characterized by Figure 1. This figure shows a cross for the moment an analyst makes a forecast available, which is not at the same moment or with the same frequency for all analysts. Next, this figure shows that there are variables which we measure at the highest frequency. As an example, the returns are shown, which we measure daily. Finally, this figure shows vertical lines depicting the moment of the earnings announcement, at which point the realization occurs of the variable that is to be forecasted by the analysts. We only use within-year earnings forecasts, which means that we only include forecasts that are forecasting the variable announced at the next upcoming yearly earnings announcement.

We have linked the earnings data to the stock data where possible. As this link could not be established for all firms, our initial sample is cut down to 596 firms. Some descriptives of the remaining sample are shown in Table 1. This table shows that the number of forecasters per firm is asymmetricly distributed. This is also the case for the number of forecasts per firm, and the number of forecasts per forecaster per firm. This asymmetry shows that there relatively few firms with high earnings forecast activity, and many firms with low earnings forecast activity. One can expect this to be linked to the size of the firm, that is, that larger firms also receive more attention from earnings analysts.

Table 1: Descriptive statistics of the 596 firms with earnings and stock market data, for the entire sample period.

Number of:	Forecasters per firm	Forecasts per firm	Forecasts per forecaster per firm
Average	100	1177	11.32
Median	94	1003	10.54
Minimum	11	84	3.43
Maximum	310	4890	27.14
Standard Deviation	52.35	807.39	3.92

5. Empirical Results

Hypothesis 1 states that individual earnings forecasts can be forecasted using (1) the average of the available forecasts and (2) the difference between the previous forecast of the analyst and the average forecast at that time. We will test this hypothesis by regressing the earnings forecast on several explanatory variables, and we expect the regression coefficients to be positive and significant for both these variables. These two variables are depicted in the top panel of Table 2, along with other variables that we include in the regression (bottom panel), to be discussed below. We will describe the regression by using the notation

$$y_{i,j,t} = X_{i,j,t}\beta_j + \varepsilon_{i,j,t}, \tag{3}$$

with subscript i denoting the individual forecaster, j the firm for which the earnings are forecasted and t the day on which the forecast is produced. The parameter coefficients are denoted by β_j, which is a vector consisting of $\beta_{j,k}$ for $k = 1, .., K$, one parameter for each variable in $X_{i,j,t}$. We will let the vector of parameter coefficients differ per firm, but not per individual nor for different time periods. Also, the error variance $\sigma^2_{\varepsilon,j}$ differs per firm.

Additional to the two earlier-mentioned variables, we also include the first difference in the average of the active forecasts. Forecasters tend to herd (Jegadeesh and Kim, 2010; Clement and Tse, 2005), but not every forecaster will respond during the same day, so that leads us to suspect that some forecasters will respond one day later. We expect these herders to follow the trend and move in the same direction as the change in the previous day, so we expect the associated parameter to be positive.

Next, we also include the previous forecast, on top of already including the difference between the previous forecast and the average forecast at that time. Some forecasters might not so much be influenced by what other forecasters do. Therefore, we do not want their relative forecast (compared to the average forecast), but the forecast itself as an additional predictor.

Finally, we also include some information about the stock market. If the stock market in general, or the market for the firm-specific stocks, is healthy, forecasters might be more positive on the future than if the situation is unhealthy. This also holds in the short-term case, which is why we expect the forecasts to be higher if the daily returns have been higher. This implies that we expect all associated signs to be positive.

For estimating this regression, we will start with the standard Ordinary Least Squares (OLS). There might be some firms for which the results will differ much from the other firms due to outliers, especially if the number of forecasts for such a firm is not high. Extreme cases will be left out of the sample, for which we use the criterion that none of the regression estimates should be more than four times the standard deviation away from the mean of that parameter. Also, firms with less than 50 data points in the regression are left out. If we would include these firms (with estimates based on a low number of data points, or with very outlying estimates), we would add noise to our results.

For the remaining firms we introduce a latent variable model for β_j. We can use this latent variable model to correct estimates that have been estimated with just over 50 data points and which are thus less accurate and more prone to outliers. These estimates can be adjusted towards the overall mean of that respective parameter, and we do that in such a way that estimates based on more than thousand observations are hardly affected. As necessary assumption for this model, we use

$$\beta_j \sim N(\beta^*, \Sigma_\beta) \tag{4}$$

which means that the latent parameter vector β_j (the estimated parameters for firm j) is related to the overall mean parameter vector β^*. For simplicity, we will assume the covariance matrix Σ_β to be diagonal. Then we employ the following steps:

1. The elements of β^* and Σ_β are estimated by taking the weighted average and weighted variance of all individual estimates.

2. We update each individual estimate by taking a weighted average:

$$\beta_{j,k}^{(u)} = w_{j,k}\beta_k^* + (1 - w_{j,k})\beta_{j,k} \tag{5}$$

$$w_{j,k} = \frac{\frac{1}{\sigma_{\beta,k}}}{\frac{1}{\sigma_{\beta,k}} + \frac{n_k}{\sigma_{\varepsilon,j}}} \tag{6}$$

   The weights are calculated using the inverses of the latent variable standard deviation and the standard error of the regression, as these determine how accurate both sources of information on the $\beta_{j,k}$ estimate are.

We will repeat (5) and (6) until convergence.

Table 2: The variables that enter the regression used to test Hypothesis 1. All regressors use one-day lagged information.

Variable	Description	Expected sign
Average Active Forecast	The average of all most recent forecasts of every forecaster, until the previous day	+
Delta Previous Forecast	The difference between the previous forecast of the forecaster, and the average active forecast at that time	+
Constant		
Δ Average Active Forecast	First difference in Average Active Forecast	+
Previous Forecast	The previous forecast of the individual forecaster	+
Stock Index Firm	The stock market index of the firm for which the earnings are forecasted	+
Stock Returns Firm	The stock market daily returns of the firm for which the earnings are forecasted	+
Stock Index S&P500	The stock market index of the S&P500 index	+
Stock Returns S&P500	The stock market daily returns of the S&P500 index	+

The results of this estimation process are summarized in Table 3 to 5. Of the 596 firms that have both earnings forecast and stock market data, 133 are left out due to less than 50 observations in the regression. Also, 43 firms have at least one estimate that is more than 4 standard deviations away from the mean of all other firms. This means that 420 firms are left after our filtering approach.

Tables 3 to 5 contain the results of applying above methodology. Table 3 shows the aggregated results on the parameter estimates of (3), both with and without applying the correction method based on (4). Similarly, Table 4 shows results for the standardized estimates, which are the estimates one finds after first standardizing all regressors. Table 5 contains results for the t-statistic and related statistics. As the corrected estimates are not t-distributed, we only report results in this table for the uncorrected estimates. While Table 4 can be used to evaluate economic significance of the estimates, Table 5 is the basis for a statistical evaluation.

We now discuss the results per table. The results for the uncorrected estimates are shown in the top half of Table 3. These results show that both main regressors have an effect in the expected direction, while only four of the remaining seven estimated parameters have on average the expected sign. The only variable for which zero is not included in the 95 % interval of parameters is Average Active Forecast.

In the bottom half of Table 3, the parameter estimates have been corrected using the estimates for all firms. This does not really affect the average or median estimate, but the spread is highly affected: the standard deviation is lower for every variable, and also the minimum and maximum estimate are closer to the average. Because of this, now both the Average Active Forecast and the Delta Previous Forecast have 95 % estimate intervals that do not include zero.

The top half of Table 4 shows the results for the standardized estimates. The two main regressors have a larger average standardized estimate. Looking at the contribution to the fit, both main regressors again perform well. The returns of the S&P500 is the only other relevant variable, but the direction of the parameter is not stable (third column of second panel) and thus the effect of the variable is not predictable. We find the contribution to the fit of the two main regressors to be 94.8%. Not shown is the average R^2, which equals 96.2%. Also, the two main

Table 3: Aggregated results on the estimates for testing Hypothesis 1, both raw and corrected.

Uncorrected estimates	Average	Minimum	Maximum	StDev	Median
Average Active Forecast	1.067	-0.076	1.763	0.274	1.090
Delta Previous Forecast	0.589	-0.735	1.622	0.303	0.608
Constant	-0.056	-1.248	0.887	0.198	-0.037
Δ Average Active Forecast	0.852	-4.100	5.930	1.001	0.783
Previous Forecast	0.002	-0.017	0.029	0.004	0.001
Stock Index Firm	0.500	-3.121	4.749	0.748	0.322
Stock Returns Firm	0.000	-0.008	0.013	0.001	0.000
Stock Index S&P500	-0.500	-8.589	10.744	1.504	-0.239
Stock Returns S&P500	-0.094	-0.754	0.965	0.266	-0.122
Corrected estimates	**Average**	**Minimum**	**Maximum**	**StDev**	**Median**
Average Active Forecast	1.080	0.459	1.491	0.161	1.086
Delta Previous Forecast	0.603	-0.026	1.261	0.198	0.606
Constant	-0.050	-0.460	0.246	0.092	-0.039
Δ Average Active Forecast	0.842	-0.314	2.643	0.463	0.807
Previous Forecast	0.002	-0.008	0.017	0.002	0.001
Stock Index Firm	0.428	-0.292	1.820	0.352	0.343
Stock Returns Firm	0.000	-0.001	0.002	0.000	0.000
Stock Index S&P500	-0.366	-2.382	1.115	0.493	-0.291
Stock Returns S&P500	-0.105	-0.505	0.461	0.156	-0.116

Table 4: Aggregated results on the standardized estimates for testing Hypothesis 1, both raw and corrected.

Uncorrected estimates	Average	Average of Absolute	$\frac{Average}{Average\ of\ Absolute}$
Average Active Forecast	1.267	1.269	0.999
Delta Previous Forecast	0.124	0.131	0.946
Δ Average Active Forecast	0.029	0.035	0.819
Previous Forecast	0.057	0.064	0.893
Stock Index Firm	0.021	0.027	0.778
Stock Returns Firm	0.004	0.030	0.133
Stock Index S&P500	-0.008	0.017	-0.484
Stock Returns S&P500	-0.073	0.285	-0.255
Corrected estimates	**Average**	**Average of Absolute**	$\frac{Average}{Average\ of\ Absolute}$
Average Active Forecast	1.183	1.183	1.000
Delta Previous Forecast	0.105	0.105	1.000
Δ Average Active Forecast	0.018	0.018	0.992
Previous Forecast	0.033	0.035	0.947
Stock Index Firm	0.016	0.016	0.985
Stock Returns Firm	0.004	0.010	0.410
Stock Index S&P500	-0.005	0.006	-0.794
Stock Returns S&P500	-0.109	0.168	-0.648

Table 5: Aggregated results on the t-Statistic of testing Hypothesis 1.

Variable	Average	Average of Absolute	Percentage significant
Average Active Forecast	14.535	14.537	97.6%
Delta Previous Forecast	8.315	8.392	91.0%
Constant	-1.531	2.715	55.5%
Δ Average Active Forecast	2.738	3.003	58.8%
Previous Forecast	3.307	3.797	68.3%
Stock Index Firm	2.794	3.037	57.6%
Stock Returns Firm	0.643	2.288	47.9%
Stock Index S&P500	-0.901	1.760	37.9%
Stock Returns S&P500	-1.605	2.726	55.0%

regressors are much more regular in the direction of their estimate, based on the third column of this panel.

In the bottom half of Table 4, the lower spread due to the correction is mostly visible in the third column, in which every proportion is closer to either one or minus one than before. The average standardized estimates and average absolute standardized estimates remain comparable to the uncorrected case, even though all values are slightly lower. Now, the contribution to the fit of the two main regressors equals 97.9%.

Table 3 shows the results for the t-statistic in testing whether the single parameter is significantly different from zero. The two main regressors have the largest average t-statistic, which shows that these estimates deviate most from zero, in a statistical sense. These two regressors also have the highest proportion of being significantly different from zero, using a 5% significance level, which is shown in the final column of the table, as discussed above.

In sum, the two main regressors stand out among the regressors, both considering their economic and their statistical significance. This suggests that, roughly speaking, one can forecast an individual earnings forecast by using the following rule of thumb:

$$Earnings\,Forecast = Average\,Active\,Forecast + 0.6\,Delta\,Previous\,Forecast \quad (7)$$

This rule of thumb can be improved for a specific firm by using the estimated coefficients.

Concerning the hypothesis, the high R^2 shows that individual earnings forecasts can be forecasted well using publicly available information. Also, of most practical and statistical use for this are the average of the active forecasts and the difference between the forecaster's previous forecast and the average at that time. These findings support Hypothesis 1.

Hypothesis 2 states that the forecast error is larger in size for the cases in which the earnings forecasts deviate more from their predicted forecast. For this, we first calculate standardized versions of the forecast error (the error in the earnings forecast) and the forecast residual (the residual of the model used for testing Hypothesis

1), while we take the time until the earnings announcement into account. We need to correct for this, as forecasts closer to the announcement most likely can be estimated with smaller forecast errors. We first regress the absolute forecast error or absolute forecast residual on an intercept and some transformations of the time until announcement; the transformations used are the linear, the centralized quadratic and the logarithmic transformation. Then we use the fit of this regression to calculate weights to standardize the errors and residuals.

Then we can test Hypothesis 2 by again making use of a regression. As the variable to be explained we will be using the log quadratic standardized forecast error of the earnings forecast, and the predictors will be an intercept and two transformations of the standardized residual. The transformations are in both cases the linear and the centralized quadratic transformation.

We have run both regressions for each individual firm and the aggregated results of these regressions are shown in Table 6. This table shows that on average, the squared standardized forecast error increases if the squared standardized residual increases, which means that forecasts that are far off from their predictable values on average perform worse. The negative-valued average parameter for the standardized residual indicates that forecasts above what is expected perform slightly better than forecasts below what is expected. The significance of these results is mostly statistical, as the R^2 is on average just below 10 %. In almost 68% of the cases, both parameters together are statistically significant, which is mostly due to the parameter of the squared standardized residual. These findings together provide enough support for Hypothesis 2 for the general case.

Table 6: Regression results to test Hypothesis 2 (significance level = 5%).

	Estimate			Standard Error		Percentage
Variable or Statistic	Average	Median	Standard Deviation	Average	Standard Deviation	significant
Intercept	0,651	0,650	0,089	0,040	0,022	77,7%
Standardized residual	-0,009	-0,013	0,096	0,026	0,017	46,0%
Squared standardized residual	0,026	0,022	0,026	0,007	0,008	62,9%
R^2	0,099	0,073	0,089			
p-value of model	0,050	0,000	0,156			67,8%

Table 7: Regression results to test Hypothesis 3.

	Q1		Q2		Q3		Q4		All	
Variable or Statistic	Average	StDev	Average	StDev	Average	StDev	Average	StDev	Average	StDev
Intercept	-0.036	2.803	-0.083	4.389	-0.126	7.496	0.434	1.826	0.009	4.282
Earnings Forecast	0.934	0.197	0.922	0.227	0.875	0.303	0.818	0.389	0.883	0.227
R^2	0.908	0.172	0.874	0.203	0.789	0.251	0.707	0.311	0.807	0.228
p-value of model	0.009	0.074	0.004	0.042	0.010	0.083	0.023	0.115	0.007	0.064
% p-value < 0.05	76.2%		76.7%		75.8%		73.3%		76.7%	

To test the third hypothesis, we regress the actuals on the earnings forecasts as quoted by the analysts. We do this for all forecasts in one regression, but we also estimate separate regressions per quarter of each year.

The results of these regressions are shown in Table 7. For all subsamples, the forecasts can be used to explain a large part in the variation in the actual earnings, considering the average values for R^2. The closer the subsample is to the realization, the higher the R^2 becomes, as could be expected.

Concerning the hypothesis, we need to compare the parameter estimates of Earnings Forecasts to 1. On average, this estimate is lower than 1 for all subsamples, which means that on average, the forecasters overshoot. They do this the most in the initial quarter of the year (Q4). When more information becomes available to correct their forecasts, the bias of the forecasts decreases. These results show support for Hypothesis 3, on average.

To test the fourth hypothesis, we regress the actuals on both the explained and the unexplained part of the earnings forecasts, using the forecasting equation as stated in (3) to construct the explained part. We do this for all forecasts in one regression, but we also estimate separate regressions per quarter.

The results of these regressions are shown in Table 8. For all subsamples, splitting the forecasts up into two parts increases the R^2 a bit compared to the results in Table 7. As before, the closer the subsample is to the realization, the higher the R^2 becomes.

Concerning the pattern in the estimates, we see for each subsample higher parameters for FOF than for ROF. Taking into account that the variation in FOF is also much higher than the variation in ROF (considering the high R^2 found while testing Hypothesis 1), this means that FOF has a much larger impact on the forecast than ROF. This is confirmed by what percentage FOF contributes to the fit, also shown in the table. This percentage varies between 80% and 95%. This difference in contribution is the largest for the quarter closest to the realization (Q1). Note that, while FOF contributes the most to the fit, it is not the case that ROF does not contribute. In an economic sense, the contribution of ROF is the largest the furthest away from the realization (Q4), while statistically, the parameters of ROF are more often significant for the quarters closer to the realization. Either way, there is always a reason to say that ROF contributes somewhat, although it is not much compared to FOF.

While the estimates of FOF are approximately similar to the estimates in Table 7 used for testing Hypothesis 3, this is not the case for the estimates of ROF. The FOF component only needs to be slightly moved towards its mean, but in the meantime the ROF component often needs to be changed by a large percentage. There are even quite some firms for which the estimates for ROF are negative, suggesting that you should do the reverse compared to FOF as what the average earnings forecaster is doing. This suggests that the earnings forecasters do not clearly improve the forecasts on top of FOF, individually.

Table 8: Regression results to test Hypothesis 4.

	Q1		Q2		Q3		Q4		All	
Variable or Statistic	Average	StDev	Average	StDev	Average	StDev	Average	StDev	Average	StDev
Intercept	-0.106	3.059	-0.139	5.303	-0.251	10.517	-0.054	5.914	-0.116	5.890
FOF	0.969	0.187	0.944	0.323	0.883	0.309	0.861	0.403	0.913	0.224
ROF	0.325	0.327	0.660	0.486	0.666	0.605	0.601	1.260	0.522	0.375
R^2	0.930	0.163	0.885	0.185	0.803	0.246	0.768	0.258	0.837	0.267
p-value of model	0.010	0.082	0.006	0.065	0.010	0.075	0.024	0.111	0.005	0.048
% p-value < 0.05	76.2%		76.7%		74.8%		58.7%		76.5%	
p-value of ROF	0,104	0,223	0,062	0,179	0,107	0,232	0,523	0,321	0,056	0,174
% p-value < 0.05	56,9%		65,4%		54,5%		6,5%		67,8%	
Relative contribution of FOF	94,9%		90,1%		86,5%		81,9%		92,1%	
% estimate of FOF is negative	0,8%		0,8%		1,5%		2,2%		1,0%	
% estimate of ROF is negative	8,7%		3,9%		6,5%		14,4%		2,7%	

In sum, the residual of the forecast predication equation can be used to improve the earnings forecast, but only slightly compared to the fit of the forecast prediction. In statistical sense, this improvement is most often significant for short forecast horizons, while economically, the gains are the largest for the larger forecast horizon. This confirms Hypothesis 4.

6. Conclusion

We have analyzed earnings forecasts retrieved from the I/B/E/S database concerning 596 firms for the sample 1995 to 2011, with a specific focus on whether these earnings forecasts can be predicted from available data. Our main result is that earnings forecasts can be predicted quite accurately using publicly available information. Second, we have shown that earnings forecasts that are less predictable are also less accurate. This confirms previous findings in the literature that herding can be valuable in obtaining accurate forecasts. We have also shown that earnings forecasters who quote forecasts that are too extreme need to correct these while the earnings announcement approaches. We have shown that the unpredictable component of earnings forecasts can contain information which we can use to improve the forecasts, and that the size of the gain is dependent on the forecast horizon.

For the end-user of the earnings forecasts, this has several implications. First, the end-user can already predict earnings forecasts to some extent, by either using the rule of thumb (7) or by estimating the equation for that specific firm. Then, the end-user can use our result that less predictable forecasts are less accurate. This means that the end-user can disregard forecasts that deviate a lot from its predicted forecast, and only focus on the forecasts that are closer to its prediction. In fact, a forecast only slightly above the predicted forecast is better news than a forecast much above the predicted forecast, as the first case is more trustworthy.

To further expand upon the notion that unpredictable earnings forecasts are less accurate, future research could try to categorize forecasts or forecasters into two types: the type that is unpredictable, possibly because of wanting stand out, and the type that is more predictable, possibly because that type is really aiming for a high accuracy. This would make it more clear where the line is that divides the forecasts that should be ignored and the forecasts that should be used. Also, this could be used to determine whether within the category of 'predictable forecasts', the more extreme forecasts are in fact more accurate, because these do contain new information, in contrast to the unpredictable forecast type. A different type of categorization that could be manifested within the forecasters is that there are two or more levels of optimism towards the specific firm. Some forecasters might always forecast higher than others, and it is interesting to see if this type of behavior is something that is significantly present in the data and how large its influence is.

A. Figures

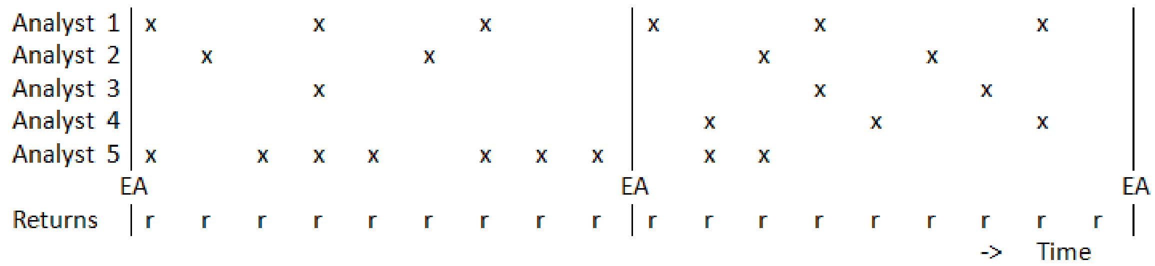

Figure 1: An example of the data format, with an x indicating an earnings forecast and EA indicating when a new yearly earnings announcement takes place. This figure shows for five forecasters for two years a variety of hypothetical patterns of forecasts.

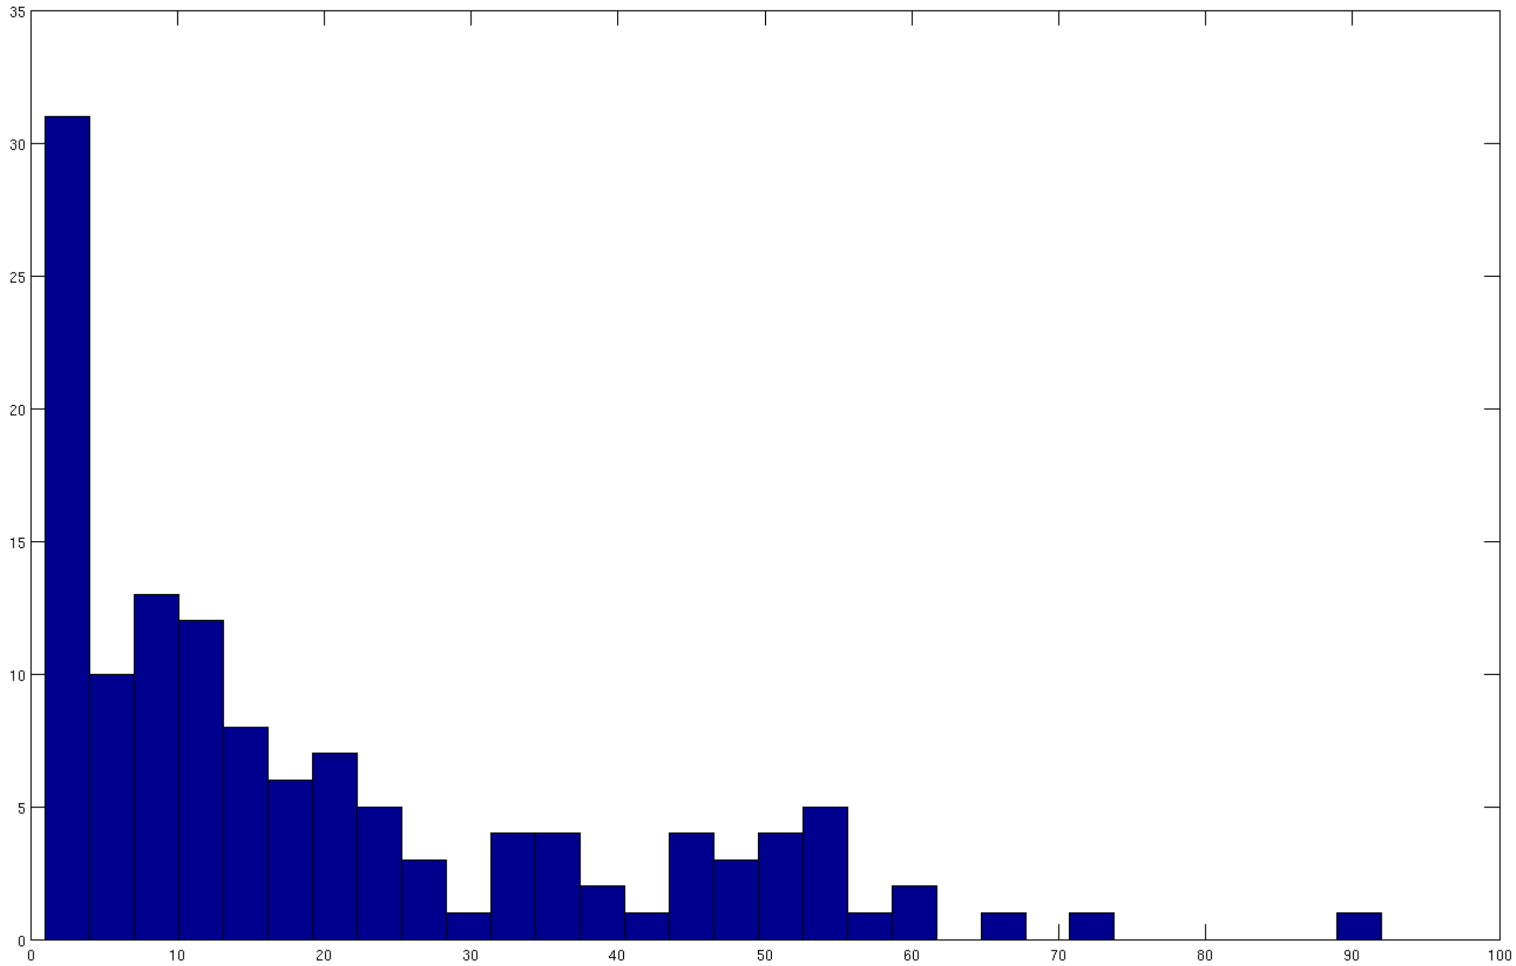

Figure 2: The distribution of the number of observed forecasts per forecaster.

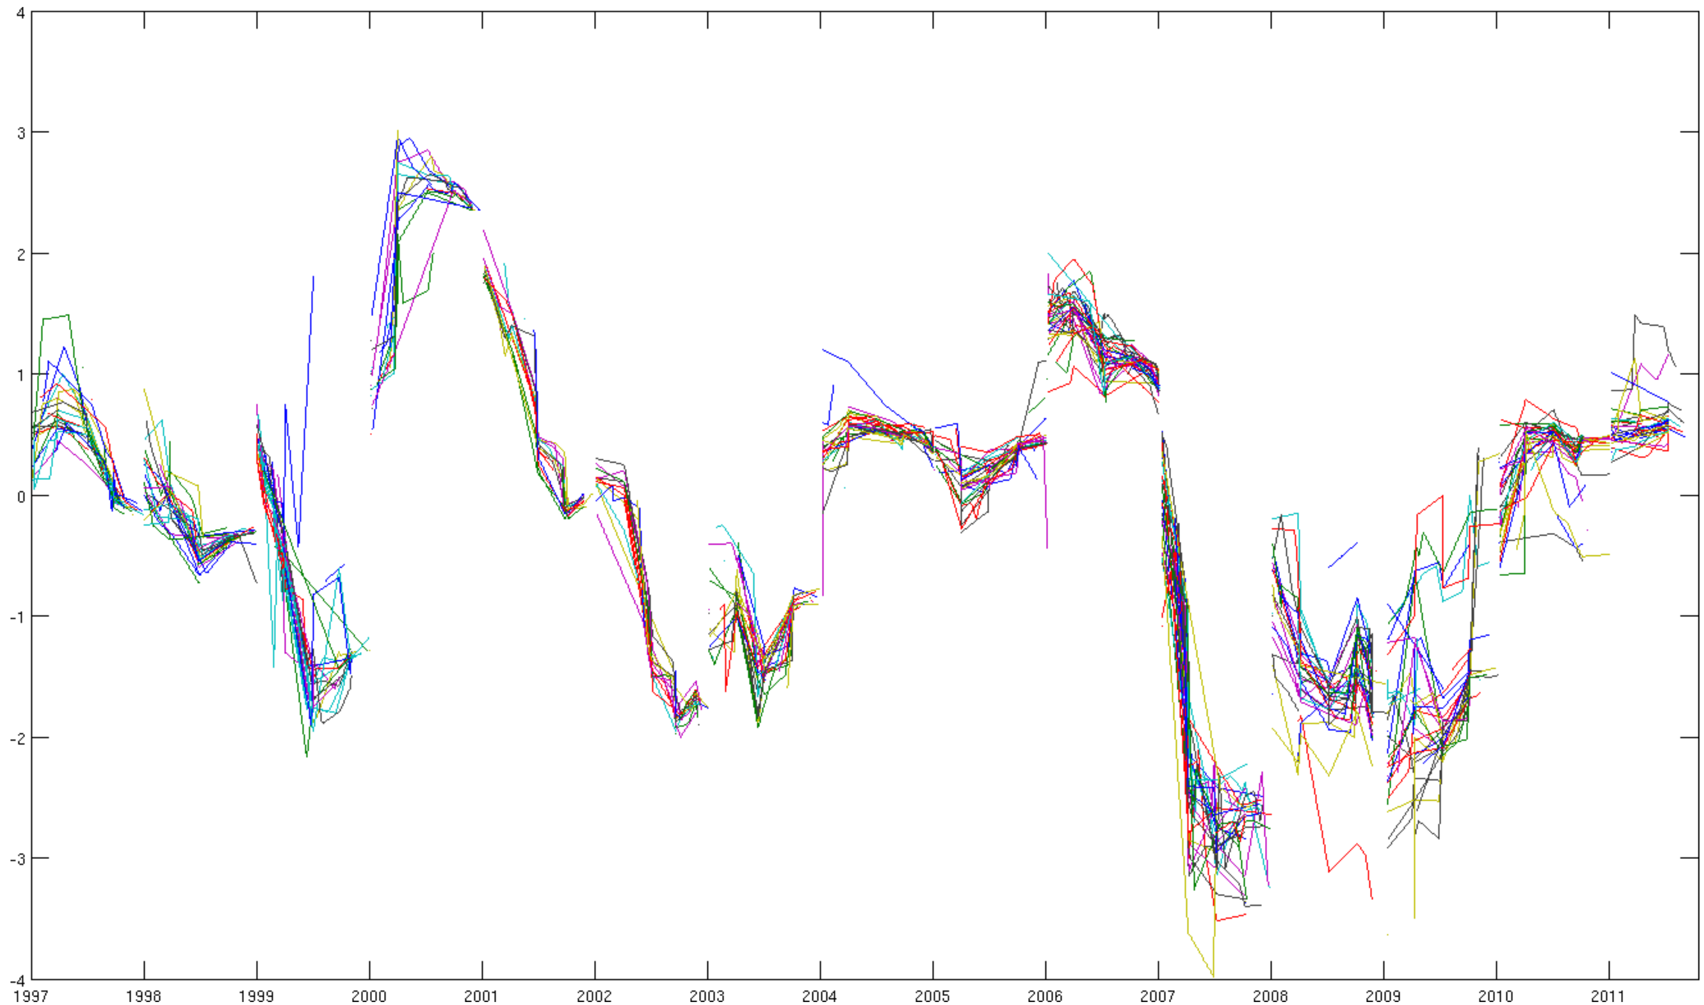

Figure 3: Linearly interpolated forecast per individual forecaster.

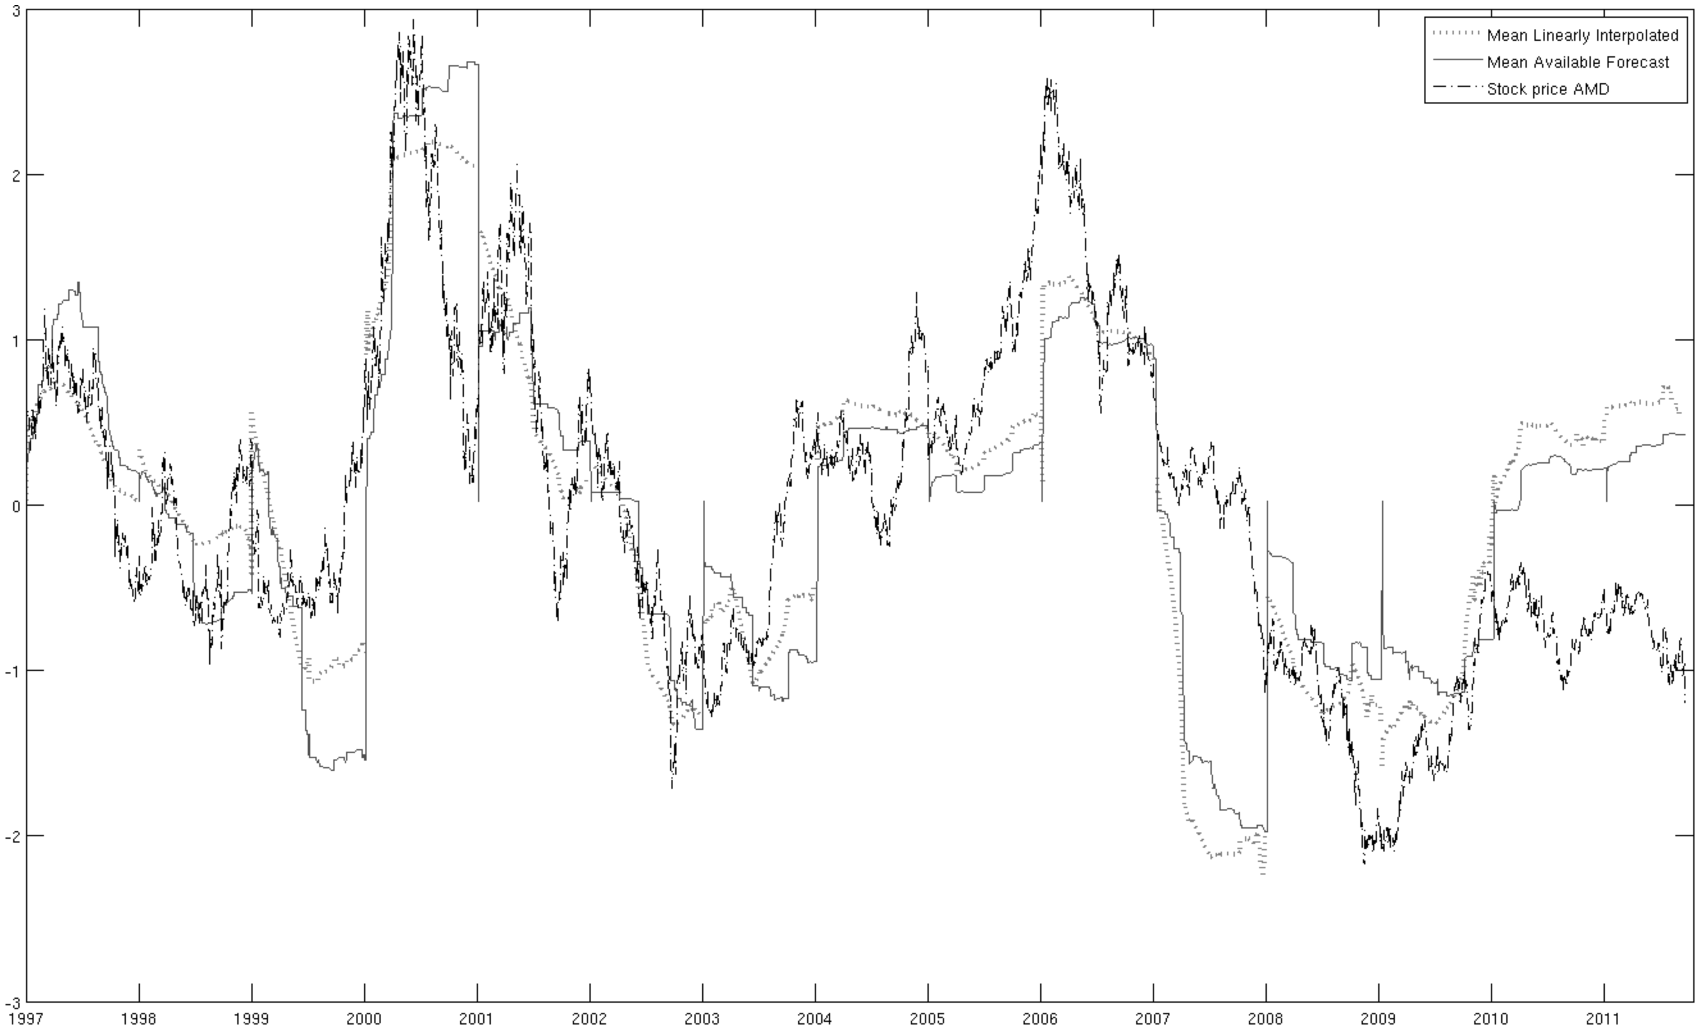

Figure 4: The average of the interpolated forecasts evaluated against the stock price of AMD and the average currently observed forecasts (all standardized).

References

Bolliger, G., 2004. The characteristics of individual analysts forecasts in Europe. Journal of Banking & Finance 28, 2283–2309. doi:10.1016/j.jbankfin.2003.09.011.

Brown, L.D., 1993. Earnings forecasting research: its implications for capital markets research. International Journal of Forecasting 9, 295–320. doi:10.1016/0169-2070(93)90023-G.

Brown, L.D., Hugon, A., 2009. Team earnings forecasting. Review of Accounting Studies 14, 587–607.

Clement, M.B., Hales, J., Xue, Y., 2011. Understanding analysts' use of stock returns and other analysts' revisions when forecasting earnings. Journal of Accounting and Economics 51, 279–299.

Clement, M.B., Tse, S.Y., 2005. Financial analyst characteristics and herding behavior in forecasting. The Journal of Finance 60, 307–341.

Cooper, R.A., Day, T.E., Lewis, C.M., 2001. Following the leader: a study of individual analysts earnings forecasts. Journal of Financial Economics 61, 383–416.

Easterwood, J.C., Nutt, S.R., 1999. Inefficiency in analysts' earnings forecasts: Systematic misreaction or systematic optimism? The Journal of Finance 54, 1777–1797.

Jegadeesh, N., Kim, W., 2010. Do analysts herd? An analysis of recommendations and market reactions. Review of Financial Studies 23, 901–937.

Kim, Y., Lobo, G.J., Song, M., 2011. Analyst characteristics, timing of forecast revisions, and analyst forecasting ability. Journal of Banking & Finance 35, 2158–2168.

Lamont, O.A., 2002. Macroeconomic forecasts and microeconomic forecasters. Journal of Economic Behavior & Organization 48, 265–280.

Landsman, W.R., Maydew, E.L., Thornock, J.R., 2012. The information content of annual earnings announcements and mandatory adoption of IFRS. Journal of Accounting and Economics 53, 34–54.

Laster, D., Bennett, P., Geoum, I.S., 1999. Rational bias in macroeconomic forecasts. The Quarterly Journal of Economics 114, 293–318.

Ramnath, S., Rock, S., Shane, P., 2008. The financial analyst forecasting literature: A taxonomy with suggestions for further research. International Journal of Forecasting 24, 34–75.

Schipper, K., 1991. Analysts forecasts. Accounting Horizons 5, 105–121.

Sheng, X., Thevenot, M., 2012. A new measure of earnings forecast uncertainty. Journal of Accounting and Economics 53, 21–33.

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Forecasting Earnings Forecasts*

Bert de Bruijn[a,b,†]
Philip Hans Franses[a,b,‡]

[a] Econometric Institute, Erasmus School of Economics,
Erasmus University Rotterdam, the Netherlands
[b] Tinbergen Institute, the Netherlands

August 22, 2013

We analyze earnings forecasts retrieved from the I/B/E/S database concerning 596 firms for the sample 1995 to 2011, with a specific focus on whether these earnings forecasts can be predicted from available data. Our main result is that earnings forecasts can be predicted quite accurately using publicly available information. Second, we show that earnings forecasts that are less predictable are also less accurate. We also show that earnings forecasters who quote forecasts that are too extreme need to correct these as the earnings announcement approaches. Finally, we show that the unpredictable component of earnings forecasts can contain information which we can use to improve the forecasts.

Keywords: Earnings Forecasts; Earnings Announcements; Financial Markets; Financial Analysts.
JEL classifications: G17, G24, M41.

*Wharton Research Data Services (WRDS) was used in preparing this paper. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers.

†Econometric Institute, Erasmus School of Economics, PO Box 1738, NL-3000 DR Rotterdam, The Netherlands; phone: +3110 4088902; email: `debruijn@ese.eur.nl` (corresponding author)
‡Econometric Institute, Erasmus School of Economics, PO Box 1738, NL-3000 DR Rotterdam, The Netherlands; phone: +3110 4081273; fax: +3110 4089162; email: `franses@ese.eur.nl`

1. Introduction

Earnings forecasts can provide useful information for investors. When investors in part rely on such forecasts, it is important to have more insights into how such earnings forecasts are created. A key research subject therefore concerns the drivers of the forecasts of earnings analysts. Such knowledge is relevant as the part that can be predicted from factors that are also observable to the end user of the forecast might not be the most interesting part of an earnings forecast. Indeed, it is the unpredictable component of the earnings forecast that amounts to the forecaster's true added value, based on latent expertise and domain-specific knowledge. As a consequence, in our perspective, the evaluation of the quality of earnings forecasts should mainly focus on that unpredictable part, as that is truly the added value of the professional forecaster.

There is much literature on the properties and accuracy of earnings forecasts, but there is no research that focuses on the prediction of such forecasts. Which variables are the most relevant drivers of earnings forecasts? Can we use the unpredictable part of the forecast to improve forecasts? In this paper we answer these questions using appropriate models. We apply these models to the earnings forecasts for a large number of firms which constitute the S&P500. Using this large sample of firms, we are confident to draw a few generalizing conclusions.

A key predictor of the earnings forecasts appears to be the average of all available earnings forecasts concerning the same forecast event. As an example, consider a forecaster who has produced his most recent forecast some period ago. If in the meantime information has been provided on the firm that has driven the forecasts of all (other) forecasters down, this forecaster will also on average produce a lower-valued forecast than before. A second predictor is the most recent difference between the individual forecaster's forecast and the average of the available contemporaneous forecasts. For example, a forecaster who previously was more optimistic about the earnings of a particular firm can be expected to persist in quoting above-average values. Other important conclusions that we draw from the data are that more unpredictable forecasts tend to be less accurate, and that the unpredictable component of the forecast can be used to improve the forecast. All in all, we document that earnings forecasts are quite predictable from data that are also available to the end user.

The outline of our paper is as follows. In Section 2 we develop several hypotheses to guide our empirical analysis, and we base these hypotheses on available studies, reviewed in Section 2. In Section 4 we discuss the data and in Section 5 we present our results. Section 6 concludes and provides various avenues for further research.

2. Literature review

Earnings forecasts have been the topic of interest for many researchers. For an extensive discussion of research on earnings forecasts in the period 1992-2007, see Ramnath et al. (2008). For earlier overviews we refer to Schipper (1991) and Brown (1993).

One stream of earnings forecasts research has focused on relationships between forecast performance and forecaster characteristics. Performance can be measured by forecast accuracy and forecast impact on stock market fluctuations. The characteristics of these performance measurements have been related to timeliness (Cooper et al., 2001; Kim et al., 2011), the number of firms that the analyst follows (Kim et al., 2011; Bolliger, 2004), the firm-specific experience of the analyst (Bolliger, 2004), age (Bolliger, 2004), the size of the firm being followed and of the firm at which the analyst works (Kim et al., 2011; Bolliger, 2004), and whether the analyst works individually or in a team (Brown and Hugon, 2009).

Another stream of research concerns the value of an earnings forecast and how it is related to what other analysts do. In particular, herding behavior is considered, which occurs when forecasters produce forecasts that converge towards the average of those of the other forecasters. There has been an effort to categorize earnings forecasters into two groups, corresponding to leaders and followers or to innovators and herders (Jegadeesh and Kim, 2010; Clement and Tse, 2005). This is interesting as different types of forecasters might consult different amounts of information which in turn can be useful for investors to incorporate into their investment decisions. A leading or innovating forecaster might on average be more useful to follow than a herding forecaster. This does not directly imply that leading forecasts are also more accurate, as accuracy and the type of forecast are not necessarily related. In fact, it has been documented that aggregation of leading forecasts is a fruitful tactic to produce accurate forecasts (Kim et al., 2011).

Recently, Clement et al. (2011) have studied the effect of stock returns and other analysts' forecasts on what analysts do. In contrast to Jegadeesh and Kim (2010) and Clement and Tse (2005), Clement et al. (2011) do not consider categorizing the forecasters into different groups. Instead, they consider how the first forecast revision after a forecast announcement is affected by how the stock market and other analysts have reacted to that forecast announcement. Landsman et al. (2012) also look at how earnings announcements affect the stock market, where these authors focus on how mandatory IFRS adoption has influenced this effect. Sheng and Thevenot (2012) propose a new earnings forecast uncertainty measure, which they use to demonstrate that forecasters focus more on the information in the earnings announcement if there is high uncertainty in the available set of earnings forecasts.

In sum, earnings forecasts have been studied concerning their performance and a few of their potential drivers. In this paper we extend the knowledge base by considering many more drivers of earnings forecasts, while we pay specific attention to the value of the unpredictable component of earnings forecasts.

3. Hypotheses

To guide our empirical analysis, we put forward several useful questions and hypotheses. We start with a general question, relate this question to previous research, and at the same time we introduce relevant notation and definitions.

Consider, for one analyst and one year, a series of earnings forecasts. This series is a single time series, with irregularly spaced observations. Time series usually exhibit serial correlation with lagged observations, which allows to forecast a future observation of the series using previous observations. This suggests to use the previous forecast of the analyst as a predictor for the current forecast. In practice this is not straightforward, as oftentimes analysts only produce a few forecasts within a year and also the firm-specific data may change over time.

Ideally, we would want to have a daily-observed series of forecasts, but this is not the case in practice. Instead, we may use the forecasts of other individual analysts, who might have produced a forecast in the recent days. This would incorporate recent information as long as the total group of analysts is active. The herding literature (Jegadeesh and Kim, 2010; Clement and Tse, 2005) suggests that there is reason to believe that individual analysts can purposely follow the aggregate forecast, also because their forecasts are driven by a common factor, as the analysts share most of the information on the state of the firm.

Forecasters can be influenced by optimism (Easterwood and Nutt, 1999). This can result in an analyst giving a higher earnings forecast than other analysts. If this optimism is persistent over time, it can be used to forecast a next forecast. For example, an analyst who previously forecasted the earnings at a level higher than other analysts did at that time, can be expected to persist with above-average values. One reason for this might be a deliberate strategic approach to get attention from certain firms (Laster et al., 1999).

We expect that by relying on both predictors, which are both well documented as relevant for earnings forecasts, it is possible to achieve a good forecasting accuracy when forecasting earnings forecasts. Because of this, we arrive at the following hypothesis.

Hypothesis 1 *Individual earnings forecasts can be forecasted well by using publically available information. This information concerns (1) the average of the forecasts of all analysts and (2) the difference between the previous forecast of the analyst and the average of the available forecasts at that time.*

To test this hypothesis, we will regress the earnings forecast on the two explanatory variables amongst several other possible candidates. We can use the regression results across different firms to evaluate how much the effect of the different variables varies across firms.

Our second question concerns the relation between how accurate one can predict a particular earnings forecast and the forecast error associated with this earnings forecast. For example, one could argue that forecasts which cannot be predicted well apparently contain new analyst-specific information and this information might improve the forecast to have a smaller forecast error. This assumes that the analyst who makes the unpredictable forecast is well capable of correctly interpreting such new information for its impact in the future. On the other hand, it could also be argued that forecasts which are fairly unpredictable perform worse, as analysts may misinterpretate the information or have other reasons to forecast an unexpected forecast. The second option is supported by Kim et al. (2011), who document that aggregation of leading forecasts is a fruitful strategy to produce accurate forecasts. Lamont (2002) shows that this also holds for macroeconomic forecasts, as he finds that bolder forecasts turn out to be less accurate. Following these two studies, we hypothesize:

Hypothesis 2 *The forecast error is larger in size when earnings forecasts deviate more from their predictable component.*

We test this hypothesis by regressing the log squared forecast error on several transformations of the residual of the forecasting equation used for Hypothesis 1.

Our third question focuses on the variation in earnings forecasts and how that is correlated with the realization of the earnings. For example, in case of a high correlation between forecasts and realizations, it could well be that the forecasts fluctuate more than the realizations do. In this case, the forecast performance would improve if the forecasts were adjusted towards an average value. The reverse could also occur. The analysts produce forecasts that are too close to the average, even though publicly-available information suggests that more extreme forecasts would be relevant. If analysts do not only care about forecasting accuracy, but also about how much attention their forecasts get, they have an incentive to overreact (Laster et al., 1999). In sum, we hypothesize:

Hypothesis 3 *Earnings forecasts tend to be too extreme and need to be corrected towards average values to improve accuracy.*

To test the above hypothesis, we consider the following equation:

$$Actual = \beta_0 + \beta_1 Forecast + rest\ term \tag{1}$$

If $\beta_0 = 0$ and $\beta_1 = 1$, the forecasts are on average right. If $\beta_1 > 1$, then the analysts are too timid with their forecasts, while they quote too extreme values if $\beta < 1$. We will estimate this equation for different forecast horizons.

Our fourth research question focuses on whether we can seperate the forecasts into two parts and evaluate these parts concerning their forecasting contribution. The two parts are constructed using our forecast of the forecast. One part is the predictable component of the forecast, and the other part is what is left, that is, the unexplained part or residual of the forecast. This could lead to several interesting situations. For example, it could be that both parts are equally important in their contribution to the forecasting quality. It could however also be that one contains more information than the other. As aggregating forecasts can lead to superior forecasts (Kim et al., 2011), we expect the predictable part of the forecast to be most important. This does not necessarily mean that the unpredictable component contains no additional information. In fact, we hypothesize:

Hypothesis 4 *The unpredictable component of the earnings forecast can be used to improve forecast performance relative to using only the predictable component of the forecast. This improvement will be largest for the longest forecast horizon.*

We test this hypothesis by estimating an equation similar to (1):

$$Actual = \beta_0 + \beta_1 FOF + \beta_2 ROF + rest\ term \tag{2}$$

in which FOF stands for Forecast of Forecast and ROF denotes the Residual of Forecast, the two parts that together constitute the original earnings forecast. The regression results can be used to compute the contribution of FOF and ROF to the total forecasting power. Again, we estimate this equation for different forecast horizons of the earnings forecasts.

4. Data and sample selection

Data has been collected from WRDS[1], using the I/B/E/S database for the analyst forecasts and the CRSP data for the stock prices and returns.

[1]http://wrds-web.wharton.upenn.edu/wrds/

Concerning the earnings forecasts, we have collected data for all firms which have been part of the S&P500 during the period 1995 to 2011. This amounts to 658 firms due to mergers, name changes and entry and exit of firms. We focus on the within-year yearly earnings forecasts, that is, the forecasts that are produced to forecast the earnings of the current year. The structure of the data is characterized by Figure 1. This figure shows a cross for the moment an analyst makes a forecast available, which is not at the same moment or with the same frequency for all analysts. Next, this figure shows that there are variables which we measure at the highest frequency. As an example, the returns are shown, which we measure daily. Finally, this figure shows vertical lines depicting the moment of the earnings announcement, at which point the realization occurs of the variable that is to be forecasted by the analysts. We only use within-year earnings forecasts, which means that we only include forecasts that are forecasting the variable announced at the next upcoming yearly earnings announcement.

We have linked the earnings data to the stock data where possible. As this link could not be established for all firms, our initial sample is cut down to 596 firms. Some descriptives of the remaining sample are shown in Table 1. This table shows that the number of forecasters per firm is asymmetricly distributed. This is also the case for the number of forecasts per firm, and the number of forecasts per forecaster per firm. This asymmetry shows that there relatively few firms with high earnings forecast activity, and many firms with low earnings forecast activity. One can expect this to be linked to the size of the firm, that is, that larger firms also receive more attention from earnings analysts.

Table 1: Descriptive statistics of the 596 firms with earnings and stock market data, for the entire sample period.

Number of:	Forecasters per firm	Forecasts per firm	Forecasts per forecaster per firm
Average	100	1177	11.32
Median	94	1003	10.54
Minimum	11	84	3.43
Maximum	310	4890	27.14
Standard Deviation	52.35	807.39	3.92

5. Empirical Results

Hypothesis 1 states that individual earnings forecasts can be forecasted using (1) the average of the available forecasts and (2) the difference between the previous forecast of the analyst and the average forecast at that time. We will test this hypothesis by regressing the earnings forecast on several explanatory variables, and we expect the regression coefficients to be positive and significant for both these variables. These two variables are depicted in the top panel of Table 2, along with other variables that we include in the regression (bottom panel), to be discussed below. We will describe the regression by using the notation

$$y_{i,j,t} = X_{i,j,t}\beta_j + \varepsilon_{i,j,t}, \tag{3}$$

with subscript i denoting the individual forecaster, j the firm for which the earnings are forecasted and t the day on which the forecast is produced. The parameter coefficients are denoted by β_j, which is a vector consisting of $\beta_{j,k}$ for $k = 1, .., K$, one parameter for each variable in $X_{i,j,t}$. We will let the vector of parameter coefficients differ per firm, but not per individual nor for different time periods. Also, the error variance $\sigma^2_{\varepsilon,j}$ differs per firm.

Additional to the two earlier-mentioned variables, we also include the first difference in the average of the active forecasts. Forecasters tend to herd (Jegadeesh and Kim, 2010; Clement and Tse, 2005), but not every forecaster will respond during the same day, so that leads us to suspect that some forecasters will respond one day later. We expect these herders to follow the trend and move in the same direction as the change in the previous day, so we expect the associated parameter to be positive.

Next, we also include the previous forecast, on top of already including the difference between the previous forecast and the average forecast at that time. Some forecasters might not so much be influenced by what other forecasters do. Therefore, we do not want their relative forecast (compared to the average forecast), but the forecast itself as an additional predictor.

Finally, we also include some information about the stock market. If the stock market in general, or the market for the firm-specific stocks, is healthy, forecasters might be more positive on the future than if the situation is unhealthy. This also holds in the short-term case, which is why we expect the forecasts to be higher if the daily returns have been higher. This implies that we expect all associated signs to be positive.

For estimating this regression, we will start with the standard Ordinary Least Squares (OLS). There might be some firms for which the results will differ much from the other firms due to outliers, especially if the number of forecasts for such a firm is not high. Extreme cases will be left out of the sample, for which we use the criterion that none of the regression estimates should be more than four times the standard deviation away from the mean of that parameter. Also, firms with less than 50 data points in the regression are left out. If we would include these firms (with estimates based on a low number of data points, or with very outlying estimates), we would add noise to our results.

For the remaining firms we introduce a latent variable model for β_j. We can use this latent variable model to correct estimates that have been estimated with just over 50 data points and which are thus less accurate and more prone to outliers. These estimates can be adjusted towards the overall mean of that respective parameter, and we do that in such a way that estimates based on more than thousand observations are hardly affected. As necessary assumption for this model, we use

$$\beta_j \sim N(\beta^*, \Sigma_\beta) \tag{4}$$

which means that the latent parameter vector β_j (the estimated parameters for firm j) is related to the overall mean parameter vector β^*. For simplicity, we will assume the covariance matrix Σ_β to be diagonal. Then we employ the following steps:

1. The elements of β^* and Σ_β are estimated by taking the weighted average and weighted variance of all individual estimates.

2. We update each individual estimate by taking a weighted average:

$$\beta_{j,k}^{(u)} = w_{j,k}\beta_k^* + (1 - w_{j,k})\beta_{j,k} \tag{5}$$

$$w_{j,k} = \frac{\frac{1}{\sigma_{\beta,k}}}{\frac{1}{\sigma_{\beta,k}} + \frac{n_k}{\sigma_{\varepsilon,j}}} \tag{6}$$

   The weights are calculated using the inverses of the latent variable standard deviation and the standard error of the regression, as these determine how accurate both sources of information on the $\beta_{j,k}$ estimate are.

We will repeat (5) and (6) until convergence.

Table 2: The variables that enter the regression used to test Hypothesis 1. All regressors use one-day lagged information.

Variable	Description	Expected sign
Average Active Forecast	The average of all most recent forecasts of every forecaster, until the previous day	+
Delta Previous Forecast	The difference between the previous forecast of the forecaster, and the average active forecast at that time	+
Constant		
Δ Average Active Forecast	First difference in Average Active Forecast	+
Previous Forecast	The previous forecast of the individual forecaster	+
Stock Index Firm	The stock market index of the firm for which the earnings are fore-casted	+
Stock Returns Firm	The stock market daily returns of the firm for which the earnings are forecasted	+
Stock Index S&P500	The stock market index of the S&P500 index	+
Stock Returns S&P500	The stock market daily returns of the S&P500 index	+

The results of this estimation process are summarized in Table 3 to 5. Of the 596 firms that have both earnings forecast and stock market data, 133 are left out due to less than 50 observations in the regression. Also, 43 firms have at least one estimate that is more than 4 standard deviations away from the mean of all other firms. This means that 420 firms are left after our filtering approach.

Tables 3 to 5 contain the results of applying above methodology. Table 3 shows the aggregated results on the parameter estimates of (3), both with and without applying the correction method based on (4). Similarly, Table 4 shows results for the standardized estimates, which are the estimates one finds after first standardizing all regressors. Table 5 contains results for the t-statistic and related statistics. As the corrected estimates are not t-distributed, we only report results in this table for the uncorrected estimates. While Table 4 can be used to evaluate economic significance of the estimates, Table 5 is the basis for a statistical evaluation.

We now discuss the results per table. The results for the uncorrected estimates are shown in the top half of Table 3. These results show that both main regressors have an effect in the expected direction, while only four of the remaining seven estimated parameters have on average the expected sign. The only variable for which zero is not included in the 95 % interval of parameters is Average Active Forecast.

In the bottom half of Table 3, the parameter estimates have been corrected using the estimates for all firms. This does not really affect the average or median estimate, but the spread is highly affected: the standard deviation is lower for every variable, and also the minimum and maximum estimate are closer to the average. Because of this, now both the Average Active Forecast and the Delta Previous Forecast have 95 % estimate intervals that do not include zero.

The top half of Table 4 shows the results for the standardized estimates. The two main regressors have a larger average standardized estimate. Looking at the contribution to the fit, both main regressors again perform well. The returns of the S&P500 is the only other relevant variable, but the direction of the parameter is not stable (third column of second panel) and thus the effect of the variable is not predictable. We find the contribution to the fit of the two main regressors to be 94.8%. Not shown is the average R^2, which equals 96.2%. Also, the two main

Table 3: Aggregated results on the estimates for testing Hypothesis 1, both raw and corrected.

Uncorrected estimates	Average	Minimum	Maximum	StDev	Median
Average Active Forecast	1.067	-0.076	1.763	0.274	1.090
Delta Previous Forecast	0.589	-0.735	1.622	0.303	0.608
Constant	-0.056	-1.248	0.887	0.198	-0.037
Δ Average Active Forecast	0.852	-4.100	5.930	1.001	0.783
Previous Forecast	0.002	-0.017	0.029	0.004	0.001
Stock Index Firm	0.500	-3.121	4.749	0.748	0.322
Stock Returns Firm	0.000	-0.008	0.013	0.001	0.000
Stock Index S&P500	-0.500	-8.589	10.744	1.504	-0.239
Stock Returns S&P500	-0.094	-0.754	0.965	0.266	-0.122
Corrected estimates	Average	Minimum	Maximum	StDev	Median
Average Active Forecast	1.080	0.459	1.491	0.161	1.086
Delta Previous Forecast	0.603	-0.026	1.261	0.198	0.606
Constant	-0.050	-0.460	0.246	0.092	-0.039
Δ Average Active Forecast	0.842	-0.314	2.643	0.463	0.807
Previous Forecast	0.002	-0.008	0.017	0.002	0.001
Stock Index Firm	0.428	-0.292	1.820	0.352	0.343
Stock Returns Firm	0.000	-0.001	0.002	0.000	0.000
Stock Index S&P500	-0.366	-2.382	1.115	0.493	-0.291
Stock Returns S&P500	-0.105	-0.505	0.461	0.156	-0.116

Table 4: Aggregated results on the standardized estimates for testing Hypothesis 1, both raw and corrected.

Uncorrected estimates	Average	Average of Absolute	$\frac{Average}{Average\ of\ Absolute}$
Average Active Forecast	1.267	1.269	0.999
Delta Previous Forecast	0.124	0.131	0.946
Δ Average Active Forecast	0.029	0.035	0.819
Previous Forecast	0.057	0.064	0.893
Stock Index Firm	0.021	0.027	0.778
Stock Returns Firm	0.004	0.030	0.133
Stock Index S&P500	-0.008	0.017	-0.484
Stock Returns S&P500	-0.073	0.285	-0.255
Corrected estimates	**Average**	**Average of Absolute**	$\frac{Average}{Average\ of\ Absolute}$
Average Active Forecast	1.183	1.183	1.000
Delta Previous Forecast	0.105	0.105	1.000
Δ Average Active Forecast	0.018	0.018	0.992
Previous Forecast	0.033	0.035	0.947
Stock Index Firm	0.016	0.016	0.985
Stock Returns Firm	0.004	0.010	0.410
Stock Index S&P500	-0.005	0.006	-0.794
Stock Returns S&P500	-0.109	0.168	-0.648

Table 5: Aggregated results on the t-Statistic of testing Hypothesis 1.

Variable	Average	Average of Absolute	Percentage significant
Average Active Forecast	14.535	14.537	97.6%
Delta Previous Forecast	8.315	8.392	91.0%
Constant	-1.531	2.715	55.5%
Δ Average Active Forecast	2.738	3.003	58.8%
Previous Forecast	3.307	3.797	68.3%
Stock Index Firm	2.794	3.037	57.6%
Stock Returns Firm	0.643	2.288	47.9%
Stock Index S&P500	-0.901	1.760	37.9%
Stock Returns S&P500	-1.605	2.726	55.0%

regressors are much more regular in the direction of their estimate, based on the third column of this panel.

In the bottom half of Table 4, the lower spread due to the correction is mostly visible in the third column, in which every proportion is closer to either one or minus one than before. The average standardized estimates and average absolute standardized estimates remain comparable to the uncorrected case, even though all values are slightly lower. Now, the contribution to the fit of the two main regressors equals 97.9%.

Table 3 shows the results for the t-statistic in testing whether the single parameter is significantly different from zero. The two main regressors have the largest average t-statistic, which shows that these estimates deviate most from zero, in a statistical sense. These two regressors also have the highest proportion of being significantly different from zero, using a 5% significance level, which is shown in the final column of the table, as discussed above.

In sum, the two main regressors stand out among the regressors, both considering their economic and their statistical significance. This suggests that, roughly speaking, one can forecast an individual earnings forecast by using the following rule of thumb:

$$Earnings\,Forecast = Average\,Active\,Forecast + 0.6\,Delta\,Previous\,Forecast \tag{7}$$

This rule of thumb can be improved for a specific firm by using the estimated coefficients.

Concerning the hypothesis, the high R^2 shows that individual earnings forecasts can be forecasted well using publicly available information. Also, of most practical and statistical use for this are the average of the active forecasts and the difference between the forecaster's previous forecast and the average at that time. These findings support Hypothesis 1.

Hypothesis 2 states that the forecast error is larger in size for the cases in which the earnings forecasts deviate more from their predicted forecast. For this, we first calculate standardized versions of the forecast error (the error in the earnings forecast) and the forecast residual (the residual of the model used for testing Hypothesis

1), while we take the time until the earnings announcement into account. We need to correct for this, as forecasts closer to the announcement most likely can be estimated with smaller forecast errors. We first regress the absolute forecast error or absolute forecast residual on an intercept and some transformations of the time until announcement; the transformations used are the linear, the centralized quadratic and the logarithmic transformation. Then we use the fit of this regression to calculate weights to standardize the errors and residuals.

Then we can test Hypothesis 2 by again making use of a regression. As the variable to be explained we will be using the log quadratic standardized forecast error of the earnings forecast, and the predictors will be an intercept and two transformations of the standardized residual. The transformations are in both cases the linear and the centralized quadratic transformation.

We have run both regressions for each individual firm and the aggregated results of these regressions are shown in Table 6. This table shows that on average, the squared standardized forecast error increases if the squared standardized residual increases, which means that forecasts that are far off from their predictable values on average perform worse. The negative-valued average parameter for the standardized residual indicates that forecasts above what is expected perform slightly better than forecasts below what is expected. The significance of these results is mostly statistical, as the R^2 is on average just below 10 %. In almost 68% of the cases, both parameters together are statistically significant, which is mostly due to the parameter of the squared standardized residual. These findings together provide enough support for Hypothesis 2 for the general case.

Table 6: Regression results to test Hypothesis 2 (significance level = 5%).

	Estimate			Standard Error		Percentage
Variable or Statistic	Average	Median	Standard Deviation	Average	Standard Deviation	significant
Intercept	0,651	0,650	0,089	0,040	0,022	77,7%
Standardized residual	-0,009	-0,013	0,096	0,026	0,017	46,0%
Squared standardized residual	0,026	0,022	0,026	0,007	0,008	62,9%
R^2	0,099	0,073	0,089			
p-value of model	0,050	0,000	0,156			67,8%

Table 7: Regression results to test Hypothesis 3.

	Q1		Q2		Q3		Q4		All	
Variable or Statistic	Average	StDev	Average	StDev	Average	StDev	Average	StDev	Average	StDev
Intercept	-0.036	2.803	-0.083	4.389	-0.126	7.496	0.434	1.826	0.009	4.282
Earnings Forecast	0.934	0.197	0.922	0.227	0.875	0.303	0.818	0.389	0.883	0.227
R^2	0.908	0.172	0.874	0.203	0.789	0.251	0.707	0.311	0.807	0.228
p-value of model	0.009	0.074	0.004	0.042	0.010	0.083	0.023	0.115	0.007	0.064
% p-value < 0.05	76.2%		76.7%		75.8%		73.3%		76.7%	

To test the third hypothesis, we regress the actuals on the earnings forecasts as quoted by the analysts. We do this for all forecasts in one regression, but we also estimate separate regressions per quarter of each year.

The results of these regressions are shown in Table 7. For all subsamples, the forecasts can be used to explain a large part in the variation in the actual earnings, considering the average values for R^2. The closer the subsample is to the realization, the higher the R^2 becomes, as could be expected.

Concerning the hypothesis, we need to compare the parameter estimates of Earnings Forecasts to 1. On average, this estimate is lower than 1 for all subsamples, which means that on average, the forecasters overshoot. They do this the most in the initial quarter of the year (Q4). When more information becomes available to correct their forecasts, the bias of the forecasts decreases. These results show support for Hypothesis 3, on average.

To test the fourth hypothesis, we regress the actuals on both the explained and the unexplained part of the earnings forecasts, using the forecasting equation as stated in (3) to construct the explained part. We do this for all forecasts in one regression, but we also estimate separate regressions per quarter.

The results of these regressions are shown in Table 8. For all subsamples, splitting the forecasts up into two parts increases the R^2 a bit compared to the results in Table 7. As before, the closer the subsample is to the realization, the higher the R^2 becomes.

Concerning the pattern in the estimates, we see for each subsample higher parameters for FOF than for ROF. Taking into account that the variation in FOF is also much higher than the variation in ROF (considering the high R^2 found while testing Hypothesis 1), this means that FOF has a much larger impact on the forecast than ROF. This is confirmed by what percentage FOF contributes to the fit, also shown in the table. This percentage varies between 80% and 95%. This difference in contribution is the largest for the quarter closest to the realization (Q1). Note that, while FOF contributes the most to the fit, it is not the case that ROF does not contribute. In an economic sense, the contribution of ROF is the largest the furthest away from the realization (Q4), while statistically, the parameters of ROF are more often significant for the quarters closer to the realization. Either way, there is always a reason to say that ROF contributes somewhat, although it is not much compared to FOF.

While the estimates of FOF are approximately similar to the estimates in Table 7 used for testing Hypothesis 3, this is not the case for the estimates of ROF. The FOF component only needs to be slightly moved towards its mean, but in the meantime the ROF component often needs to be changed by a large percentage. There are even quite some firms for which the estimates for ROF are negative, suggesting that you should do the reverse compared to FOF as what the average earnings forecaster is doing. This suggests that the earnings forecasters do not clearly improve the forecasts on top of FOF, individually.

Table 8: Regression results to test Hypothesis 4.

	Q1		Q2		Q3		Q4		All	
Variable or Statistic	Average	StDev	Average	StDev	Average	StDev	Average	StDev	Average	StDev
Intercept	-0.106	3.059	-0.139	5.303	-0.251	10.517	-0.054	5.914	-0.116	5.890
FOF	0.969	0.187	0.944	0.323	0.883	0.309	0.861	0.403	0.913	0.224
ROF	0.325	0.327	0.660	0.486	0.666	0.605	0.601	1.260	0.522	0.375
R^2	0.930	0.163	0.885	0.185	0.803	0.246	0.768	0.258	0.837	0.267
p-value of model	0.010	0.082	0.006	0.065	0.010	0.075	0.024	0.111	0.005	0.048
% p-value < 0.05	76.2%		76.7%		74.8%		58.7%		76.5%	
p-value of ROF	0,104	0,223	0,062	0,179	0,107	0,232	0,523	0,321	0,056	0,174
% p-value < 0.05	56,9%		65,4%		54,5%		6,5%		67,8%	
Relative contribution of FOF	94,9%		90,1%		86,5%		81,9%		92,1%	
% estimate of FOF is negative	0,8%		0,8%		1,5%		2,2%		1,0%	
% estimate of ROF is negative	8,7%		3,9%		6,5%		14,4%		2,7%	

In sum, the residual of the forecast predication equation can be used to improve the earnings forecast, but only slightly compared to the fit of the forecast prediction. In statistical sense, this improvement is most often significant for short forecast horizons, while economically, the gains are the largest for the larger forecast horizon. This confirms Hypothesis 4.

6. Conclusion

We have analyzed earnings forecasts retrieved from the I/B/E/S database concerning 596 firms for the sample 1995 to 2011, with a specific focus on whether these earnings forecasts can be predicted from available data. Our main result is that earnings forecasts can be predicted quite accurately using publicly available information. Second, we have shown that earnings forecasts that are less predictable are also less accurate. This confirms previous findings in the literature that herding can be valuable in obtaining accurate forecasts. We have also shown that earnings forecasters who quote forecasts that are too extreme need to correct these while the earnings announcement approaches. We have shown that the unpredictable component of earnings forecasts can contain information which we can use to improve the forecasts, and that the size of the gain is dependent on the forecast horizon.

For the end-user of the earnings forecasts, this has several implications. First, the end-user can already predict earnings forecasts to some extent, by either using the rule of thumb (7) or by estimating the equation for that specific firm. Then, the end-user can use our result that less predictable forecasts are less accurate. This means that the end-user can disregard forecasts that deviate a lot from its predicted forecast, and only focus on the forecasts that are closer to its prediction. In fact, a forecast only slightly above the predicted forecast is better news than a forecast much above the predicted forecast, as the first case is more trustworthy.

To further expand upon the notion that unpredictable earnings forecasts are less accurate, future research could try to categorize forecasts or forecasters into two types: the type that is unpredictable, possibly because of wanting stand out, and the type that is more predictable, possibly because that type is really aiming for a high accuracy. This would make it more clear where the line is that divides the forecasts that should be ignored and the forecasts that should be used. Also, this could be used to determine whether within the category of 'predictable forecasts', the more extreme forecasts are in fact more accurate, because these do contain new information, in contrast to the unpredictable forecast type. A different type of categorization that could be manifested within the forecasters is that there are two or more levels of optimism towards the specific firm. Some forecasters might always forecast higher than others, and it is interesting to see if this type of behavior is something that is significantly present in the data and how large its influence is.

A. Figures

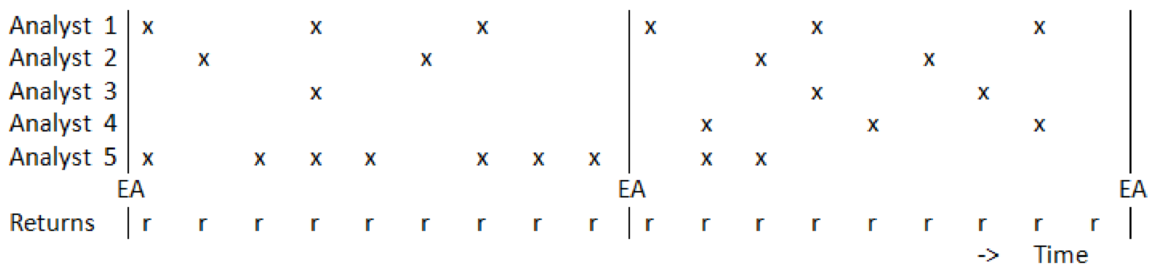

Figure 1: An example of the data format, with an x indicating an earnings forecast and EA indicating when a new yearly earnings announcement takes place. This figure shows for five forecasters for two years a variety of hypothetical patterns of forecasts.

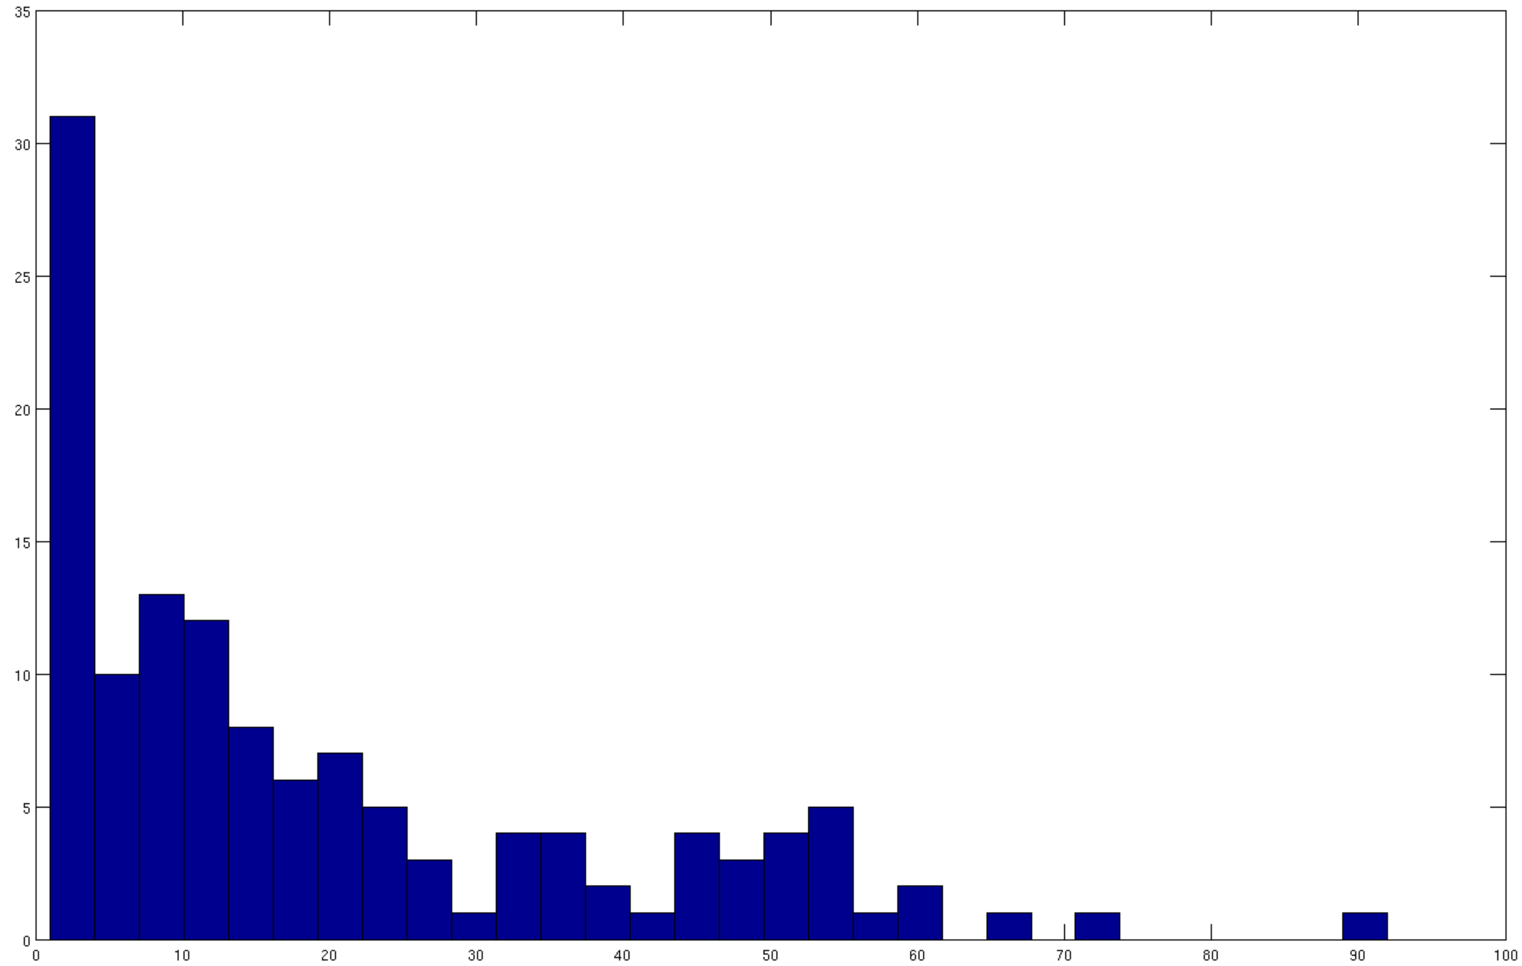

Figure 2: The distribution of the number of observed forecasts per forecaster.

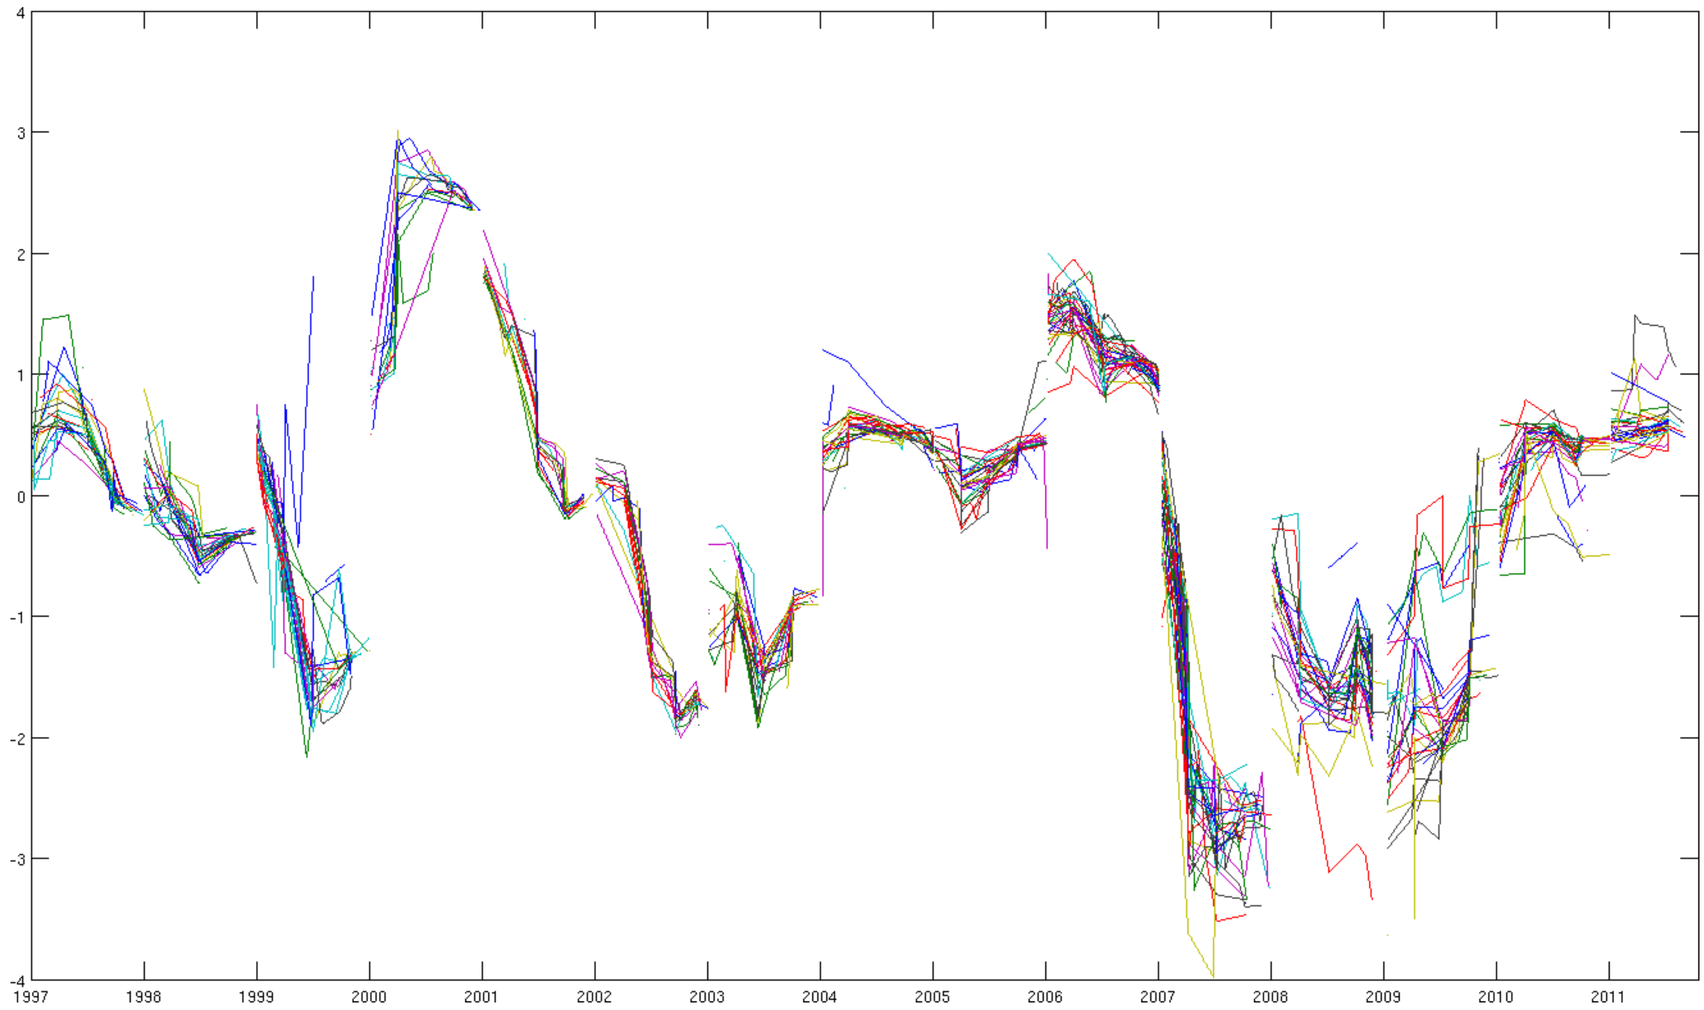

Figure 3: Linearly interpolated forecast per individual forecaster.

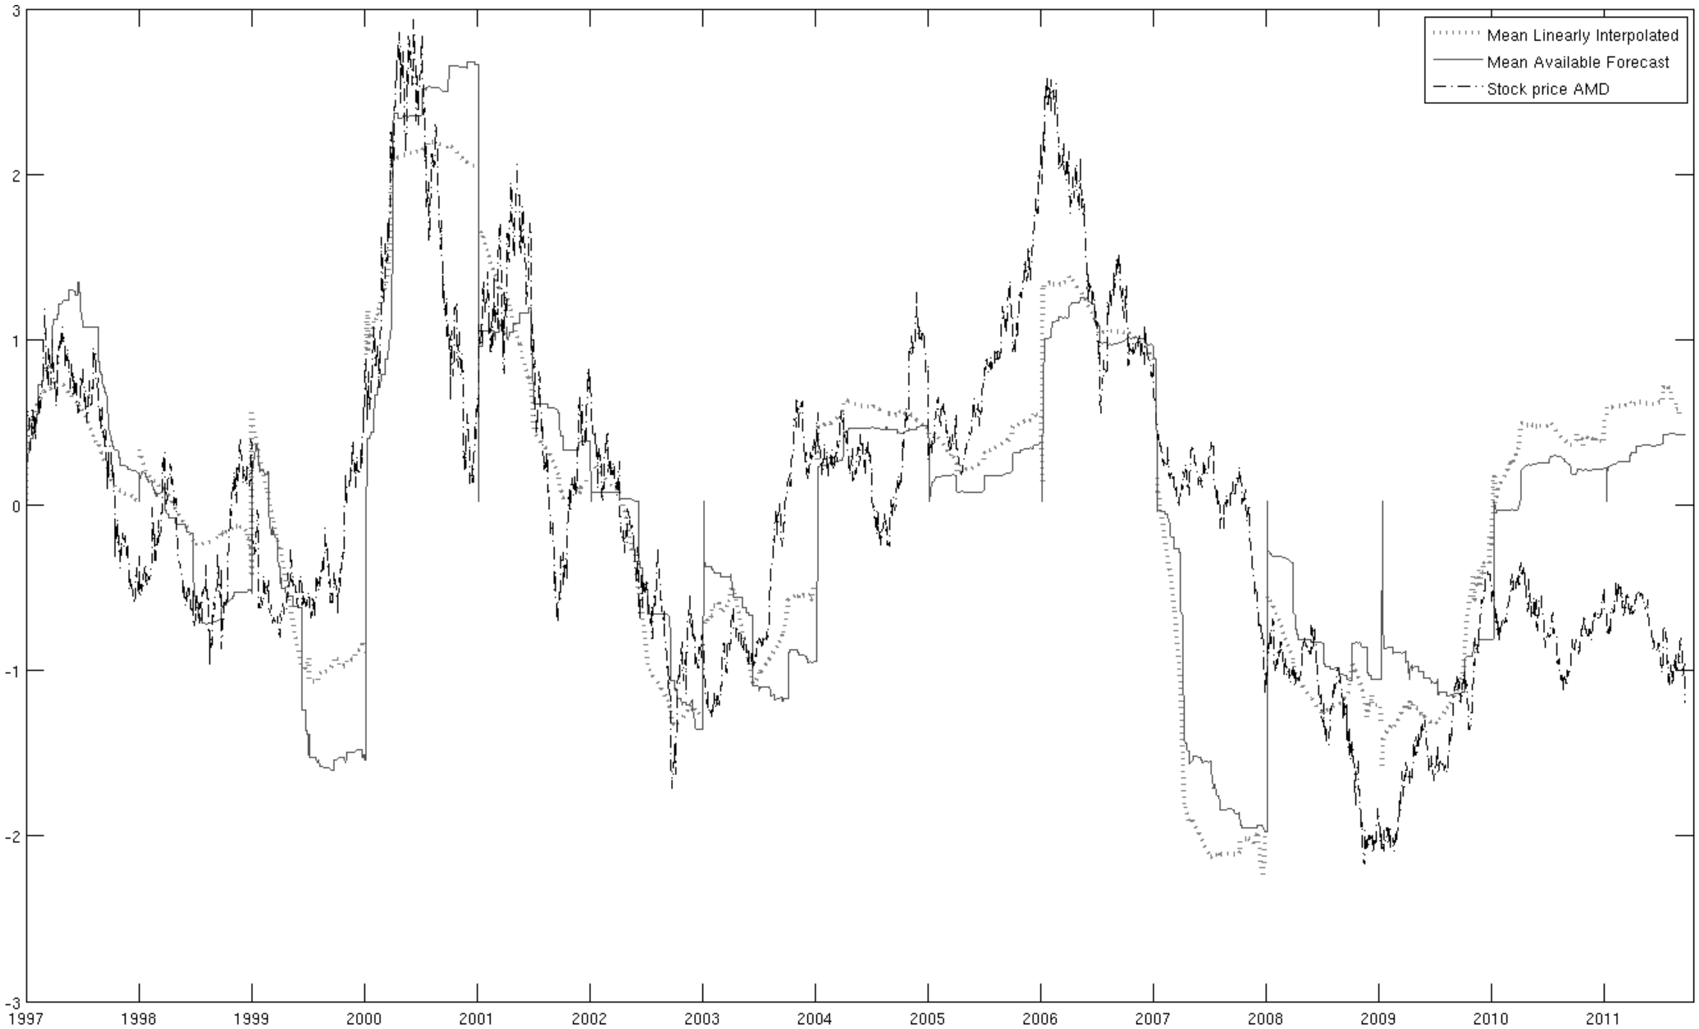

Figure 4: The average of the interpolated forecasts evaluated against the stock price of AMD and the average currently observed forecasts (all standardized).

References

Bolliger, G., 2004. The characteristics of individual analysts forecasts in Europe. Journal of Banking & Finance 28, 2283–2309. doi:10.1016/j.jbankfin.2003.09.011.

Brown, L.D., 1993. Earnings forecasting research: its implications for capital markets research. International Journal of Forecasting 9, 295–320. doi:10.1016/0169-2070(93)90023-G.

Brown, L.D., Hugon, A., 2009. Team earnings forecasting. Review of Accounting Studies 14, 587–607.

Clement, M.B., Hales, J., Xue, Y., 2011. Understanding analysts' use of stock returns and other analysts' revisions when forecasting earnings. Journal of Accounting and Economics 51, 279–299.

Clement, M.B., Tse, S.Y., 2005. Financial analyst characteristics and herding behavior in forecasting. The Journal of Finance 60, 307–341.

Cooper, R.A., Day, T.E., Lewis, C.M., 2001. Following the leader: a study of individual analysts earnings forecasts. Journal of Financial Economics 61, 383–416.

Easterwood, J.C., Nutt, S.R., 1999. Inefficiency in analysts' earnings forecasts: Systematic misreaction or systematic optimism? The Journal of Finance 54, 1777–1797.

Jegadeesh, N., Kim, W., 2010. Do analysts herd? An analysis of recommendations and market reactions. Review of Financial Studies 23, 901–937.

Kim, Y., Lobo, G.J., Song, M., 2011. Analyst characteristics, timing of forecast revisions, and analyst forecasting ability. Journal of Banking & Finance 35, 2158–2168.

Lamont, O.A., 2002. Macroeconomic forecasts and microeconomic forecasters. Journal of Economic Behavior & Organization 48, 265–280.

Landsman, W.R., Maydew, E.L., Thornock, J.R., 2012. The information content of annual earnings announcements and mandatory adoption of IFRS. Journal of Accounting and Economics 53, 34–54.

Laster, D., Bennett, P., Geoum, I.S., 1999. Rational bias in macroeconomic forecasts. The Quarterly Journal of Economics 114, 293–318.

Ramnath, S., Rock, S., Shane, P., 2008. The financial analyst forecasting literature: A taxonomy with suggestions for further research. International Journal of Forecasting 24, 34–75.

Schipper, K., 1991. Analysts forecasts. Accounting Horizons 5, 105–121.

Sheng, X., Thevenot, M., 2012. A new measure of earnings forecast uncertainty. Journal of Accounting and Economics 53, 21–33.