

TI 2012-067/4

Tinbergen Institute Discussion Paper



What drives the Quotes of Earnings Forecasters?

Bert de Bruijn

Philip Hans Franses

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

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

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 is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

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

What drives the quotes of earnings forecasters?*

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

July 11, 2012

Abstract

Earnings forecasts can be useful for investment decisions. Research on earnings forecasts has focused on forecast performance in relation to firm characteristics, on categorizing the analysts into groups with similar behaviour and on the effect of an earnings announcement by the firm on future earnings forecasts. In this paper we investigate the factors that determine the value of the forecast and also investigate to what extent the timing of the forecast can be modeled. We propose a novel methodology that allows for such an investigation. As an illustration we analyze within-year earnings forecasts for AMD in the period 1997 to 2011, where the data are obtained from the I/B/E/S database. Our empirical findings suggest clear drivers of the value and the timing of the earnings forecast. We thus show that not only the forecasts themselves are predictable, but that also the timing of the quotes is predictable to some extent.

Keywords: Earnings Forecasts; Earnings Announcements; Financial Markets; Financial Analysts.

JEL classifications: G17, G24, M41.

*The authors would like to thank Koen Bel and Richard Paap for helpful comments and suggestions.

[†]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. What drives forecasters to give their quotes, that is, when do they do so and what value will they quote? Answers to these questions are relevant as the part that can be predicted from factors that are also observable to the user of the forecast is perhaps not the most interesting part of an earnings forecast. In fact, it is the unpredictable component of the earnings forecast that amounts to the forecaster's true added value, based on latent expertise and knowledge. Hence, the evaluation of the quality of earnings forecasts should also 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 performance of earnings forecasts, but interestingly enough, there is almost no research available on the drivers of the timing of such forecasts. What makes it that forecasters give their quotes today and not on any other day, and is the timing perhaps linked with the forecast quote itself? In this paper we fill this gap by proposing two models for earnings forecasts, one for the forecast itself and, more importantly, one for the timing of the quote. We illustrate the methodology for AMD earnings forecasts. Even though this concerns only a single firm, we are tempted to draw a few generalizing conclusions. Two key drivers of the timing of earnings forecasts appear to be the time since or to earnings announcements and the forecasting behavior of other individual forecasters. For example, the longer ago the forecast was given, the more likely it becomes that a new quote will be released. All in all, we document that earnings forecasts are to some extent predictable.

The outline of our paper is as follows. In Section 2 we concisely review the relevant literature. In Section 3 we discuss the AMD data. In Section 4 and 5 we present the models for the forecast quotes and for the timing. Estimation will be based on Bayesian Gibbs techniques as we also want to incorporate the days when no forecasts are quoted, and hence for which some of the relevant variables are unobserved. Section 6 explores if these two models can somehow be connected.

Finally, Section 7 concludes with 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 forecast research has focused on relationships between forecast performance and forecaster characteristics. Performance can be measured in several ways, such as forecast accuracy and forecast impact on stock markets. 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).

There has also been an interest to document how the value of an earnings forecast is related to what other analysts have forecasted. In particular, herding behavior is considered, which occurs when forecasters produce forecasts that converge towards the average of the other forecasters. In this line, 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 have different amounts of information which can be useful for investors to incorporate into their investment decisions. A leading or innovating forecast might on average be more useful than a herding forecast, as the latter will be close to what is already known. 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, aggregation of leading forecasts seems to be 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 grouping the forecasters. Instead, they consider how the first forecast revision after a forecast announcement has been 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 existing set of earnings forecasts.

In sum, the value of an earnings forecast has been studied concerning its performance and some of its drivers. Concerning the timing of a forecast, there has been some theoretical research (Guttman, 2010). An application to empirical data has been limited to investigating the effect of timing on performance (Kim et al., 2011). There are no studies which jointly consider the drivers of either the value or the timing of the earnings forecast. In the present paper we provide such an examination.

3 Data

We analyze the unadjusted forecasts for yearly earnings per share for the firm Advanced Micro Devices, Inc. (which is better known by its acronym AMD). A stylized version of the data format is shown in Figure 1. In this figure, an x corresponds to the occurrence of a forecast on a particular day by an analyst, while one of the explanatory variables (returns) is observed each day. EA indicates the earnings announcements, and these are the targeted events of the forecasts. After each EA, the next EA is the new target of the forecasts, as we limit ourselves to within-year forecasts, which are the forecasts that are produced for the yearly earnings announcement of the current year. The figure further shows several typical properties of the data, that is (i) the occurrence of forecasts is irregular; (ii) some forecasters produce forecasts more often than others; (iii) some forecasters produce on a more regular

basis than others; (iv) some forecasters only become active at a later stage; and (v) some forecasters stop making forecasts.

We only use within-year data, so we discard forecasts that target at future yearly earnings announcements instead of the first upcoming announcement. We do so to avoid additional (and complicated) correlation between several forecasts from the same forecaster for different earnings announcements.

For AMD, we investigate the period from January 9th 1997 to September 14th 2011. In this period 129 forecasters have produced at least one yearly earnings forecast for AMD. The period amounts to 5373 days for which we observe whether a forecaster produces a forecast or not. In total 2589 forecasts have been produced. Figure 2 shows the distribution of the number of the forecasts per forecaster. The minimum number of forecasts per single forecaster is 1, while the maximum number is 92. The number of forecasts produced by a forecaster is heavily skewed. Many forecasters produce few forecasts, while a few forecasters produce many. The forecasts and their dates have been taken from the I/B/E/S database. The time it takes to convert the data from this source into daily observed variables is the main reason that we only consider AMD in the present study.

We also use stock market related variables, which are based on the NYSE quotes, normalized to 1 on the first day of the sample, January 9th 1997. As a measure of returns, we use log-returns.

4 Modeling the value of earnings forecasts

In this section we present a model for the value of the earnings forecasts. We also present the explanatory variables and then apply the model to our data as an illustration.

Model specification

Let $y_{i,t}$ denote the forecast of forecaster i quoted on day t , if actually observed. This indicates that we have different time series with forecasts for every forecaster. As new information becomes available only gradually over time, the forecasts also will

differ over time. And, contemporaneous forecasts produced by different forecasters will most likely not be the same, see Figure 3. To visually highlight which forecasts are produced by the same forecaster, we interpolate linearly between two subsequent forecasts produced for the same earnings announcement. Figure 3 seems to show that the different forecasters follow a similar pattern over time. This seems to suggest that most of the variation in the forecasts might be captured by variables that only change over time, and not by individual-specific variation.

We propose a simple model relating $y_{i,t}$ to other variables Z_{t-1} , that is,

$$y_{i,t} = \gamma Z_{t-1} + \eta_{i,t} \tag{1}$$

with $\eta_{i,t} \sim N(0, \sigma_\eta^2)$ denoting the idiosyncratic shock. Z_{t-1} contains an intercept and relevant regressors that are highly correlated with the forecast and which are the same for all forecasters. Our choice of regressors will be discussed below.

If the parameter estimates for the model are known, the model can also be used to construct estimates of forecasts for days when no forecasts are quoted. For this, we use the notation $y_{i,t}^*$. The forecast revision for the same target is the difference $y_{i,t}^* - y_{i,s_t}$ when the last available forecast by forecaster i on day t has been produced on day s_t . This estimated variable will be used later as a regressor in the timing model, to be discussed in Section 5.

We will estimate the parameters of this model using a Bayesian approach (Gibbs sampling), using a diffuse prior on γ and a degenerate inverted Gamma-2 prior on σ_η^2 ($p(\sigma_\eta^2) \propto \sigma_\eta^{-2}$). Using the draws of γ and σ_η^2 we can then construct draws of $y_{i,t}^*$.

Choice of regressors in the value model

We now discuss which regressors we will use in the vector Z_{t-1} . We decide to consider the average of the forecasts that are observed and the stock price. Figure 4 seems to confirm our choice of regressors as the patterns are rather similar. As the individual interpolated forecasts are similar to each other and thus also to their average, the individual forecasts could very well be modeled using these regressors.

In sum, we define Z_{t-1} as

$$Z_{t-1} = [1 \quad MF_{t-1} \quad p_{t-1}] \tag{2}$$

with $MF_{t-1} = \frac{1}{n} \sum_{i=1}^n y_{i,s_{t-1}}$ denoting the average of the relevant forecasts of the forecasters at time $t - 1$ (which are the most recent forecasts made for the same earnings report) and with p_{t-1} denoting the stock price at time $t - 1$.

This means that (1) can be written as $y_{i,t} = \gamma_0 + \gamma_1 MF_{t-1} + \gamma_2 p_{t-1} + \eta_{i,t}$. If $\gamma_0 = \gamma_2 = 0$ and $\gamma_1 = 1$, then new forecasts are on average equal to the previous averaged forecasts for the same earnings announcement. When $\gamma_1 \neq 1$, the forecasts are on average moving away from their previous average values. For example, if $\gamma_1 > 1$, this would mean that new forecasts are on average higher than previous forecasts in the same year. If the new forecasts are also better forecasts, this information could be used to improve the previous forecasts. A similar situation occurs if $\gamma_1 = 1$, but $\gamma_0 > 0$ (while $\gamma_2 = 0$).

Empirical results for AMD

As an illustration, we will now present results of the Gibbs sampling to estimate the parameters of (1). Details can be found in Appendix A.1, which also contains the convergence criteria.

Results on the posterior density of the parameters are in Table 1. The estimated mean of the MF_{t-1} coefficient shows that (*ceteris paribus*) an increase in the average observed forecast results in about a similar-sized increase in the individual forecasts. Interestingly, this parameter of MF_{t-1} seems to be quite in excess of 1 (at least when using the 95% HPD region), indicating that for this data the forecast updates were on average positive. This could be due to an upward trend in the forecasts throughout the year, but it could also be that the more positive forecasters have updated their forecasts more often. Another possibility is that this $\gamma_1 > 1$ is only due to $\gamma_2 < 0$. The interpretation of this result of γ_2 is that a higher stock price is related to a decrease in the forecast, while keeping the average forecast constant.

In sum, the parameter estimates are not too far off from the case with $\gamma_0 = 0$ and $\gamma_1 = 1$. Apparently, earnings forecasters produce new forecasts that are about equal to the previous average observed forecast plus a negative effect of previous-day stock returns.

5 Modeling the timing of earnings forecasts

In this section, we present a model for the timing of the earnings forecasts. We also discuss the explanatory variables and then we apply the model to our data.

Model specification

Let $q_{i,t}$ be a binary variable indicating whether or not forecaster i has quoted a new forecast at time t . For modeling this $q_{i,t}$, we will use a dynamic probit model, which uses variables contained in a vector $X_{i,t}$ as regressor, that is,

$$P(q_{i,t} = 1) = \Phi(X_{i,t}\beta), \quad (3)$$

where Φ is the CDF of the normal distribution. This model can also be written as

$$w_{i,t} = X_{i,t}\beta + \varepsilon_{i,t} \quad (4)$$

$$q_{i,t} = I[w_{i,t} > 0] \quad (5)$$

with $\varepsilon_{i,t} \sim N(0, \sigma_\varepsilon^2)$ the idiosyncratic shock and $I[.]$ the indicator function which is equal to 1 if its argument is true and 0 otherwise.

For identification reasons, σ_ε^2 will be set equal to 1. The elements of β are estimated using a Gibbs-sampling procedure with a diffuse prior on β , just like for γ in Section 4. Data Augmentation will be applied to obtain simulated values of $w_{i,t}$.

Choice of regressors in the timing model

As regressors in $X_{i,t}$ we use various types of variables. We use general timing variables, stock market information for the firm under study, aggregate information from all analysts, individual track record of forecaster i , and the forecast update compared to the previous forecast by forecaster i . We will now discuss each of the regressors and hypothesize expected signs of their coefficients. An overview of all regressors can be found in Table 2.

Considering the general timing variables, we choose to incorporate the time distance of t to the previous and forthcoming yearly earnings report (TSR and TUR) and also the intermediate quarterly reports (TQ1, TQ2 and TQ3). Their corresponding variables have all been constructed as a measure of distance (in contrast

to closeness) to the respective event, namely by using $\log(|\text{number of days}|+1)$. The absolute value reflects that we expect a similar forecasting probability just before the event as just after. The logarithm reflects that the difference in effect between, for example, 2 and 3 days before the next event will be larger than between 45 and 46 days. We expect to see a larger forecasting probability in the vicinity of a new report, which corresponds to a negative sign for these regressors.

For the stock market, we consider the daily returns of the AMD stock on the NYSE and its absolute value as a measure of volatility. We expect large returns (in absolute sense) to be a conveyor of new information on the firm, so we expect a positive sign for AbsRET. The sign of the coefficient of RET then depends on whether positive or negative returns have a larger impact on the updating probability. Due to optimism (Easterwood and Nutt, 1999), we expect the forecasters to react more to positive shocks, and hence we conjecture a positive sign.

Concerning the aggregate actions of all analysts, we make use of the total number of forecasts on the previous day, together with the change in the mean of all the available forecasts and the absolute value of the change in mean. It could be that forecasters respond more due to general activity of others (which corresponds with SQL, the number of forecasts on the previous day), but it could also be that they only react more if these actions are of a large size (corresponding with AbsDMF, the absolute change in the mean forecast). In either case, we expect both of these to have positive coefficients. For DMF it could go either way, but as before for RET, we expect a positive sign due to optimism.

For the variables related to the track record of the forecaster, we include whether a forecaster has already been active for this firm this year (ACT). We expect that forecasters may eventually stop following a firm after following it during previous years. In such a situation, the associated parameter is negative. Also, if a forecaster has already produced a forecast yesterday (QL), he will not be inclined to do it again today since that would make him seem unreliable. In general, the effect of the time since the last forecast (TSU) is less clear. It could be a positive effect as the forecaster might want to regularly update his forecast, but it could also be a negative effect if after a long time (but still within the same year) the forecaster

decides to stop following the firm. Given the irregularity of the forecasts, we expect the latter option to be more likely and thus we expect a negative sign for TSU.

The variables corresponding to the forecast updates of the forecaster are both based on our conjecture that forecasters are more likely to update their forecasts if their current forecast is very different from the one that is now publicly available. This would be reflected by a positive sign for AbsDYY, while for DYY we again expect a positive sign due to optimism. The forecast update will be based on the draws of $y_{i,t}^*$ as discussed in Section 4, which makes the update equal to $y_{i,t}^* - y_{i,s_t}$.

Empirical results for AMD

As an illustration, we will now present the results of the Gibbs sampling to estimate the parameters of (4). Details can be found in Appendix A.2, which also contains the convergence criteria.

Table 3 contains the results of the Gibbs estimation. We present the mean and several other properties of the distribution of the parameters. Values of the mean in boldface indicate that 0 is outside the 95% HPD region and that the sign is as expected. If the values of the mean are in italics, then 0 is outside the HPD region, but the sign is not as expected.

Also shown in the figure is a column with the header *StCoef*. This standardized coefficient would have been the mean of the distribution of the coefficient if we would have standardized the regressors beforehand. This is needed to compare the explanatory power of the different regressors, as larger values of StCoef (in an absolute sense) correspond to more explanatory power.

Parameter estimates and fit

The results show that there is no single variable that explains the bulk of the variation. The two variables with the largest coefficient estimates (StCoef) are ACT and TSU, which are two of the three variables that associate with the individual track record. The estimates confirm our expectation (i) that forecasters that have not yet been active in this current year have a smaller probability of producing a forecast than those who have been active already and (ii) that the more time has passed

since the last forecast, the smaller is the probability of a new forecast.

Other variables with much explanatory power are the general timing variables that are the same for all forecasters. These are the time since the last report and until the next report, and the time distance to the quarterly reports. These all have negative posterior estimates as we expected, which shows that forecasters become more active just before and after an official firm event.

The other variables have varying explanatory power. SQL and QL have almost opposite estimates (with signs as expected). This shows that the day after his own forecast a forecaster will not quickly produce a new forecast, but he will produce a new forecast the day after many forecasts by all other forecasters. This might be due to reacting to the other forecasters, but it could also occur because there has been some unexpected firm information and that some forecasters are a bit slower to respond than others. Both stock market related variables do not contain much explanatory power. It can be seen that large returns tend to be followed by a forecast a day later, in particular when large returns are positive. The change in the aggregate forecast has a significant but small effect on the probability of producing a new forecast. The parameter of DMF is the only (significant) estimate that has an unexpected sign.

Surprisingly, the two variables for which 0 is contained in the HPD region (DYY and AbsDYY) are the two forecast update variables. This indicates that forecasters ignore the *value* of their own past forecast in deciding when to produce a new one, even though they do not ignore the *timing* of their past forecast (indicated by the coefficients for ACT, TSU and QL).

Overall, the highest value of McFadden R^2 across all draws is equal to 0.118, which is high for probit models, certainly for sample sizes like ours. We conclude that the timing of the earnings forecasts can be modeled quite well.

Fit per category

Table 4 shows the joint explanatory power per cluster of variables, calculated by summing the absolute values of the standardized coefficients. It can be seen that the majority of the explained timing of the forecasts is explained by general timing

variables and variables corresponding to the individual track record (together 82.7%). The activity of all forecasters and the stock market contain only little information, and the effect of the forecast update itself is negligible.

6 Correlations between the shocks of both models

In the previous sections, we have discussed a linear model for the value of a forecast ($y_{i,t} = \gamma Z_{t-1} + \eta_{i,t}$) and a dynamic probit model for the timing of a forecast (slightly rewritten to $q_{i,t} = I[w_{i,t} = X_{i,t}\beta > -\varepsilon_{i,t}]$). These models have been discussed independently, but it might well be that they are related. In this section, we discuss two extensions that both capture an interaction between the two models. The first is a correlation between the shocks in both models and the second is a relation between the occurrence of tail values for both models.

Correlation between the errors

In this subsection we allow for correlation between $\eta_{i,t}$ and $\varepsilon_{i,t}$ using a joint model. If this correlation is positive, this suggests that optimistic forecasts (forecasts above what is to be expected on the basis of (1)) are faster quoted than pessimistic forecasts. If the correlation is negative, the reverse is true: pessimistic forecasts will have a larger probability of being quoted than optimistic forecasts.

Allowing for a correlation ρ between both shocks is equivalent to defining the covariance matrix $\Sigma = \begin{bmatrix} \sigma_\varepsilon^2 & \sigma_{\varepsilon\eta} \\ \sigma_{\varepsilon\eta} & \sigma_\eta^2 \end{bmatrix}$ with $\sigma_{\varepsilon\eta} = \rho\sigma_\varepsilon\sigma_\eta$ the covariance. Similar as in Section 5, we need to restrict σ_ε to 1 for identification. This suggests using $\Sigma = \begin{bmatrix} 1 & \sigma_{\varepsilon\eta} \\ \sigma_{\varepsilon\eta} & \sigma_\eta^2 \end{bmatrix}$, but as it is not obvious how to sample the covariance matrix using this specification, we instead follow the reparametrization of McCulloch, Polson and Rossi (2000): $\Sigma = \begin{bmatrix} 1 & \rho \\ \rho & \phi + \rho^2 \end{bmatrix}$. In the sampling routine, we use the degenerate inverted Wishart prior $p(\Sigma) \propto |\Sigma|^{-3/2}$. Details can be found in Appendix A.3.

As can be seen in Table 5, the results show that there is no clear evidence of a correlation between the two models, that is, 0 is almost in the center of the HPD region of $\rho_{\eta\varepsilon}$. This means that positive and negative forecasts are produced with the same probability.

As there is not much correlation, the results of the other parameters are also almost the same as in the two separate models.

A posteriori evaluation of correlation of tails

Another possibility is that tail forecasts are produced with a different probability than forecasts that do not differ much from the general expectation. To see whether this is the case, we calculate the correlation between $\eta_{i,t}^2$ and $\varepsilon_{i,t}^2$. Alternatively, we also calculate the correlation between $|\eta_{i,t}|$ and $|\varepsilon_{i,t}|$. In both cases, a positive correlation indicates that the tail forecasts are produced at unexpected moments, while normal forecasts are produced more often at normal forecasting moments. A negative correlation indicates that tail forecasts tend to be produced at normal forecasting moments, while normal forecasts would occur more often at unexpected forecasting moments.

We do not incorporate this possibility directly into the model, but calculate these correlations a posteriori for each iteration in the Gibbs sampler. This means that we base this measure on the estimated residuals. As we calculate it for each iteration, this procedure provides us with a density of the tail correlations.

We calculate the above measure only for $q_{i,t} = 1$, as for $q_{i,t} = 0$ the $\eta_{i,t}$ is not observed but simulated instead. This simulation is of course independent of the $\varepsilon_{i,t}$ as we do not model this type of correlation. Because of this, we only focus on the values of $\eta_{i,t}$ and $\varepsilon_{i,t}$ for when $q_{i,t} = 1$.

Table 6 contains the correlations between the tails of the two forecast residuals. It can be seen that there is virtually no correlation between the size of the residuals of the two models, no matter if one uses the absolute or the squared measure. This means that tail values of the forecasts do not coincide with unexpected timing of the forecasts. It is also not the case that they are contrasting each other, such as would be the case if unexpected timing results in forecasts that are closer to the general expectation than otherwise.

Next, there is also no significant autocorrelation in the residuals (bottom row of Table 6). This means that there is no sign that extreme-valued or unexpected observations occur two days in a row.

7 Conclusion

The methodology that we presented in this paper can be used to elicit the drivers of earnings forecast values and of the timing of such quotes. We illustrated our models for earnings forecasts for AMD, and we found that both value and timing are quite predictable.

Of course, our findings concerned only a single firm, and future work should indicate whether our current findings hold true more in general. A second issue for further research concerns the models themselves. At present, we have assumed that the model parameters are the same across all forecasters. However, it might well be that there are clusters of forecasters who display similar behavior, and for who different drivers might prevail. This would build on the recent studies in Jegadeesh & Kim (2010) and Clement & Tse (2005).

A third, and more involved, avenue for further research relates to the evaluation of the quality of earnings forecasts. As we have demonstrated in this paper, apparently an earnings forecast can be decomposed into a part that is predictable using an econometric model which includes publicly available information, and a part that is not predictable as such. Intuitively, one would reward the qualities of an earnings forecaster more if higher accuracy is associated with a larger unpredictable part. This leads to a new look at evaluating the accuracy of earnings forecasts, and we relegate this to future research.




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




Figure 3: Linearly interpolated forecast per individual forecasters.




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

	Mean	St.Dev.	0.5%	5%	Median	95%	99.5%	HPD-LB	HPD-UB
Intercept	0.049	0.003	0.042	0.044	0.049	0.053	0.055	0.044	0.054
MF_{t-1}	1.055	0.001	1.053	1.054	1.055	1.055	1.056	1.053	1.056
p_{t-1}	-0.108	0.003	-0.116	-0.113	-0.108	-0.104	-0.101	-0.114	-0.103
σ_η	0.115	0.003	0.107	0.109	0.115	0.120	0.124	0.108	0.121

Table 1: Properties of the estimated posterior marginal densities of the coefficients and error standard deviation in the model for the value of the earnings forecasts. The properties are the mean, standard deviation, several percentiles (including the median) and the 95% HPD interval. These are based on 100000 after-burnin draws.

Cluster	Abbreviation	Short description	Expected sign
	Intercept		
General timing	TSR	Time since last yearly report	-
	TUR	Time until next yearly report	-
	TQ1	Time away from quarter 1 report	-
	TQ2	Time away from quarter 2 report	-
	TQ3	Time away from quarter 3 report	-
Stock market	RET	Previous day daily return of AMD stock	+
	AbsRET	Absolute value of RET	+
All analysts	SQL	Total new forecasts on previous day	+
	DMF	Change in the mean available forecast	+
Individual track record	AbsDMF	Absolute value of DMF	+
	ACT	Has analyst been inactive in this year until now?	-
	TSU	Time since previous forecast by analyst	-
	QL	Did analyst forecast on previous day?	-
Forecast update	DYY	Difference between forecast $y_{i,t}^*$ & previous forecast $y_{i,s_{t-1}}$	+
	AbsDYY	Absolute value of DYY	+

Table 2: The regressors in the timing model. They are clustered according to five categories: general time variables, stock market information, aggregate forecasters information, individual track record and forecast updates.

	Mean	stdev	0.50%	5%	Median	95%	99.50%	HPD-LB	HPD-UB	St.Coef
Intercept	1.053	0.257	0.382	0.633	1.051	1.476	1.704	0.561	1.566	
TSR	-0.135	0.019	-0.183	-0.166	-0.135	-0.104	-0.086	-0.172	-0.098	-0.131
TUR	-0.201	0.019	-0.249	-0.233	-0.201	-0.169	-0.151	-0.239	-0.164	-0.168
TQ1	-0.147	0.010	-0.172	-0.163	-0.147	-0.131	-0.122	-0.166	-0.128	-0.162
TQ2	-0.096	0.010	-0.121	-0.112	-0.096	-0.080	-0.070	-0.116	-0.077	-0.096
TQ3	-0.089	0.011	-0.116	-0.106	-0.089	-0.072	-0.062	-0.110	-0.068	-0.098
RET	2.848	0.426	1.766	2.152	2.848	3.552	3.950	2.032	3.695	0.044
AbsRET	5.440	0.552	3.968	4.531	5.443	6.346	6.840	4.360	6.525	0.068
SQL	0.049	0.003	0.042	0.044	0.049	0.054	0.056	0.043	0.055	0.100
DMF	<i>-0.629</i>	0.270	-1.329	-1.075	-0.629	-0.188	0.065	-1.155	-0.097	-0.021
AbsDMF	0.716	0.350	-0.204	0.137	0.720	1.286	1.607	0.042	1.409	0.023
ACT	-0.823	0.037	-0.919	-0.884	-0.822	-0.762	-0.726	-0.895	-0.749	-0.337
TSU	-0.071	0.008	-0.092	-0.085	-0.072	-0.058	-0.051	-0.087	-0.056	-0.212
QL	-1.077	0.084	-1.295	-1.215	-1.076	-0.940	-0.860	-1.242	-0.916	-0.085
DYY	0.006	0.005	-0.006	-0.001	0.006	0.014	0.018	-0.003	0.015	0.007
AbsDYY	-0.008	0.006	-0.023	-0.017	-0.008	0.001	0.007	-0.019	0.003	-0.006

Table 3: Properties of the estimated posterior marginal densities of the coefficients in the model for the timing of the forecasts. The properties are the mean, standard deviation, several percentiles (including the median), the 95% HPD interval and the mean in case the variable had been standardized. Boldface indicates 0 outside the HPD region and the sign is as expected; italics indicate 0 outside the HPD region with an unexpected sign. The results are based on 100000 after-burnin draws. Across all draws, the maximum value for *McFadden R²* is 0.118, with a total number of observations equal to 411831.

Cluster	Joint explanatory power
General timing	0.654
Stock market	0.112
All forecasters	0.144
Individual track record	0.634
Forecast update	0.013

Table 4: Sum of the (absolute values of the) standardized coefficients per cluster of variables.

Regressor	Mean	St.Dev	0.5%	5%	Median	95%	99.5%	HPD-LB	HPD-UB	StCoef
<i>Value Model</i>										
Constant	0.049	0.003	0.042	0.044	0.049	0.053	0.055	0.044	0.054	
MF_{t-1}	1.055	0.001	1.053	1.054	1.055	1.055	1.056	1.053	1.056	
P_{t-1}	-0.108	0.003	-0.115	-0.113	-0.108	-0.104	-0.101	-0.114	-0.103	
<i>Timing Model</i>										
Constant	1.060	0.259	0.385	0.631	1.061	1.483	1.722	0.544	1.564	
TSR	-0.135	0.019	-0.184	-0.167	-0.136	-0.103	-0.085	-0.173	-0.098	-0.132
TUR	-0.202	0.019	-0.251	-0.233	-0.202	-0.170	-0.152	-0.240	-0.165	-0.169
TQ1	-0.147	0.010	-0.172	-0.163	-0.147	-0.131	-0.122	-0.166	-0.128	-0.162
TQ2	-0.096	0.010	-0.121	-0.113	-0.096	-0.080	-0.070	-0.116	-0.077	-0.096
TQ3	-0.089	0.011	-0.117	-0.107	-0.089	-0.071	-0.062	-0.109	-0.068	-0.098
RET	2.853	0.431	1.756	2.145	2.853	3.564	3.976	2.013	3.698	0.044
AbsRET	5.461	0.552	4.043	4.548	5.462	6.366	6.861	4.378	6.538	0.068
SQL	0.049	0.003	0.042	0.044	0.049	0.054	0.056	0.043	0.055	0.100
DMF	-0.637	0.266	-1.335	-1.073	-0.638	-0.199	0.056	-1.167	-0.119	-0.021
AbsDMF	0.716	0.350	-0.205	0.133	0.720	1.286	1.592	0.022	1.392	0.023
ACT	-0.824	0.037	-0.918	-0.884	-0.824	-0.762	-0.728	-0.897	-0.751	-0.337
TSU	-0.072	0.008	-0.091	-0.084	-0.072	-0.059	-0.051	-0.087	-0.056	-0.212
QL	-1.076	0.083	-1.300	-1.216	-1.075	-0.942	-0.872	-1.240	-0.914	-0.085
DYY	0.006	0.005	-0.006	-0.001	0.006	0.014	0.018	-0.003	0.015	-0.006
AbsDYY	-0.008	0.006	-0.022	-0.017	-0.008	0.002	0.007	-0.019	0.003	0.007
<i>Covariance</i>										
σ_η^2	0.113	0.003	0.105	0.108	0.113	0.119	0.123	0.107	0.120	
$\sigma_{\eta\varepsilon}$	0.006	0.021	-0.045	-0.029	0.004	0.043	0.057	-0.033	0.049	
$\rho_{\eta\varepsilon}$	0.017	0.064	-0.133	-0.086	0.013	0.127	0.167	-0.100	0.143	

Table 5: Properties of the estimated posterior marginal densities of the coefficients and parameters in the two models, allowing for correlation between $\eta_{i,t}$ and $\varepsilon_{i,t}$. The properties are the mean, standard deviation, several percentiles (including the median), the 95% HPD interval and the mean in case the variable had been standardized. These are based on 100000 after-burnin draws.

	Absolute, $q_{i,t} = 1$	Square, $q_{i,t} = 1$
$\eta_{i,t}$ & $\varepsilon_{i,t}$	0.002	0.006
$\eta_{i,t}$ & $\eta_{i,t-1}$	0.011	0.019
$\varepsilon_{i,t}$ & $\varepsilon_{i,t-1}$	0.018	0.023

Table 6: Correlation between absolute value and squares of residuals, calculated for the cases in which $q_{i,t} = 1$. An asterisk would have indicated that 0 is not contained in the 95% HPD region.

A Additional results of the Bayesian inference

In this appendix we present additional results of the Bayesian inference, including the relevant Gibbs samplers and the convergence criteria.

A.1 The model for the value of the forecast

Gibbs steps

For the model for the value of the forecast, we need to draw $\gamma|\sigma_\eta, y_{i,t}$ and $\sigma_\eta|\gamma, y_{i,t}$. Additionally, we also want to construct draws of $y_{i,t}^*|\gamma, \sigma_\eta$ for the cases in which we do not observe $y_{i,t}$.

As the formulation of (1) is that of a linear regression model, we can use standard results, such as documented in (Zellner, 1971). This means that $\gamma|\sigma_\eta, y_{i,t} \sim N(\hat{\gamma}, \sigma_\eta(Z'Z)^{-1})$ with $\hat{\gamma}$ the OLS estimate of (1) and Z the matrix containing the regressors. Also, $\sigma_\eta|\gamma, y_{i,t}$ is distributed as $IG(\sum (y_{i,t} - Z_{i,t}\gamma)^2, nobs)$, with IG denoting the Inverted Gamma distribution, with $Z_{i,t}$ containing the regressors and with $nobs = \sum q_{i,t}$ equal to the number of observations that we have for $y_{i,t}$.

To simulate values of $y_{i,t}^*$ given γ and σ_η , we just simply use $y_{i,t}^* \sim N(Z\gamma, \sigma_\eta)$, which follows directly from the formulation of (1).

Convergence criteria

Table A1 shows the convergence criteria for the value model. To find out whether the draws have converged, we split the data into two parts and test whether the means are equal. Due to the Markovian nature of the draws, one needs to include an AR(1) term in the testing equation and/or use standard errors that are consistent when there is serial correlation. As can be seen, for none of the parameters convergence is rejected at any reasonable significance level.

Also shown are the effective sample sizes for the different parameters, which corrects for the fact that new draws do not contain 100% information if there is autocorrelation. Due to the low autocorrelation in the draws, the effective sample size is equal to the true sample size in this case. In fact, one can see from the Gibbs steps for this model that the draws of γ are also in theory uncorrelated.

Variable	Mean 1st half	Mean 2nd half	P-value dummy	ϕ (SE)	k	ESS
Constant	0.049	0.048	0.176	-0.002 (0.003)	< 1	100000
MF_{t-1}	1.055	1.055	0.242	-0.006 (0.003)	< 1	100000
p_{t-1}	-0.109	-0.108	0.189	-0.002 (0.003)	< 1	100000
σ_η	0.115	0.115	0.232	0.001 (0.003)	< 1	100000

Table A1: Convergence criteria of the draws for the model of the value of the revision based on 100000 draws. Results are based on an AR(1)-model with a dummy indicating whether a draw belongs to the first half or the second half of the draws. As there might be serial correlation in the draws, we use Newey-West HAC standard errors. The table presents the mean of the first 50000 and the last 50000 draws, the p-value of the dummy and the estimated AR(1) parameter ϕ . The value of k indicates the lowest value for which $\phi^k < 0.10$, with which the effective sample size (ESS) can be calculated.

Variable	Mean 1st half	Mean 2nd half	P-value dummy	AR(1) parameter	k	ESS
Constant	1.05	1.05	0.72	0.95	43.43	2302.31
TSR	-0.14	-0.13	0.74	0.96	51.48	1942.49
TUR	-0.20	-0.20	0.15	0.95	44.14	2265.60
Q1	-0.15	-0.15	0.62	0.94	37.53	2664.81
Q2	-0.10	-0.10	0.26	0.95	45.63	2191.72
Q3	-0.09	-0.09	0.08	0.95	45.65	2190.63
RET	2.84	2.85	0.45	0.93	33.55	2980.61
ARET	5.44	5.44	0.87	0.94	35.40	2824.60
SQL	0.05	0.05	0.81	0.89	19.50	5129.27
DMF	-0.63	-0.63	0.19	0.84	13.18	7585.26
ADMF	0.71	0.73	0.01	0.85	13.71	7292.49
NF	-0.82	-0.82	0.93	0.97	65.45	1527.79
TSU	-0.07	-0.07	0.38	0.96	56.22	1778.62
QL	-1.08	-1.08	0.51	0.93	33.48	2986.85
DYY	0.01	0.01	0.84	0.86	14.88	6718.59
ADYY	-0.01	-0.01	0.83	0.74	7.67	13029.84

Table A2: Convergence criteria of the draws for the model of the timing of the revision based on 100000 draws. Results are based on an AR(1)-model with a dummy indicating whether a draw belongs to the first half or the second half of the draws. As there might be serial correlation in the draws, we use Newey-West HAC standard errors. The table presents the mean of the first 50000 and the last 50000 draws, the p-value of the dummy and the estimated AR(1) parameter ϕ . The value of k indicates the lowest value for which $\phi^k < 0.10$, using which the effective sample size (ESS) can be calculated.

A.2 The model for the timing of the forecast

Gibbs steps

For the model for the timing of the forecast, we need to draw $\beta|\sigma_\varepsilon, y_{i,t}^*, q_{i,t}$ and $\sigma_\varepsilon|\beta, y_{i,t}^*, q_{i,t}$ with $y_{i,t}^*$ taken directly from the draws of Appendix A.1. To do this, we apply Data Augmentation (Dempster et al., 1977), which means that we also construct draws of $w_{i,t}|\beta, \sigma_\varepsilon, y_{i,t}^*, q_{i,t}$ and use these to draw $\beta|\sigma_\varepsilon, y_{i,t}^*, w_{i,t}, q_{i,t}$ and $\sigma_\varepsilon|\beta, y_{i,t}^*, w_{i,t}, q_{i,t}$.

Given $w_{i,t}$, the formulation of (4) is that of a linear regression model, so we can use standard results from (Zellner, 1971). This means that $\beta|\sigma_\varepsilon, y_{i,t}^*, w_{i,t}, q_{i,t} \sim N(\hat{\beta}, \sigma_\varepsilon(X'X)^{-1})$ with $\hat{\beta}$ the OLS estimate of (4) and $X_{i,t}$ the matrix of regressors as in Table 2. Also, $\sigma_\varepsilon|\beta, y_{i,t}^*, w_{i,t}, q_{i,t}$ is then distributed as $IG(\sum (w_{i,t} - X_{i,t}\beta)^2, N)$, with $X_{i,t}$ containing the values of X for forecaster i at time t and with N equal to the total number of observations.

To simulate values of $w_{i,t}|\beta, \sigma_\varepsilon, y_{i,t}^*, q_{i,t}$, we use the formulation of (4): if $q_{i,t} = 1$, this means that $w_{i,t} > 0$ which is equivalent to $\varepsilon_{i,t} > -X_{i,t}\beta$. This means that in this case we need to draw $\varepsilon_{i,t}$ from $N(0, \sigma_\varepsilon)$, truncated from below by $-X_{i,t}\beta$. Similarly, if $q_{i,t} = 0$, $w_{i,t} < 0$ and thus $\varepsilon_{i,t} < -X_{i,t}\beta$. In this case, we need to draw $\varepsilon_{i,t}$ from $N(0, \sigma_\varepsilon)$ truncated from above by $-X_{i,t}\beta$. Using these draws of $\varepsilon_{i,t}$ we can then construct draws of $w_{i,t} = X_{i,t}\beta + \varepsilon_{i,t}$.

Convergence criteria

Table A2 shows the convergence criteria for the timing model. As above, we have split the data into two parts and tested whether the means are equal to find out whether the draws have converged, using again an AR(1) term in the testing equation. As can be seen, for one of the sixteen parameters, convergence is rejected at the 5% level. This is not directly a sign that the total system has not converged, as the probability of rejecting one of sixteen hypotheses at the 5% level is quite large.

Also shown are the effective sample sizes for the different parameters. The autocorrelation in the draws is quite high (between 0.74 and 0.97), so the effective sample size is quite a bit smaller than the number of draws. Still, the distributions of all

parameters are effectively based on at least 1500 almost-independent observations, which is large enough to have a satisfactory sample distribution.

A.3 The combined model with correlation

Gibbs steps

For the combined model with correlation, we again apply Data Augmentation. This means that we need draws of $\beta, \gamma | \Sigma, y_{i,t}^*, w_{i,t}$; of $\Sigma | \beta, \gamma, y_{i,t}^*, w_{i,t}$; of $y_{i,t}^* | \gamma, \Sigma, w_{i,t}$; and of $w_{i,t} | \beta, \gamma, \Sigma, y_{i,t}^*, q_{i,t}$. Most of the Gibbs steps are obvious extensions of the steps in the separate models.

First, β and γ can be drawn simultaneously, as shown in for example (Zellner and Ando, 2010). Define $\tilde{\beta} = \begin{bmatrix} \beta \\ \gamma \end{bmatrix}$. To combine X and Z into one large regressor matrix, we need to incorporate the correlation between both models, which can be done using the Cholesky decomposition $L'L = \Sigma^{-1}$. This results in the definition of $\tilde{X} = \begin{bmatrix} L_{1,1}X & L_{1,2}Z \\ L_{2,1}X & L_{2,2}Z \end{bmatrix}$. Similarly, $\tilde{y} = \begin{bmatrix} L_{1,1}w + L_{1,2}y^* \\ L_{2,1}w + L_{2,2}y^* \end{bmatrix}$. Then, $\tilde{\beta} | \Sigma, y_{i,t}^*, w_{i,t} \sim N(\hat{\tilde{\beta}}, (\tilde{X}'\tilde{X})^{-1})$ with $\hat{\tilde{\beta}}$ the OLS estimator $(\tilde{X}'\tilde{X})^{-1}\tilde{X}'\tilde{y}$.

For $\Sigma | \beta, \gamma, y_{i,t}^*, w_{i,t}$, the situation is not standard as the element corresponding to σ_ε^2 needs to be restricted to 1. We follow the solution of McCulloch et al. (2000), who propose the reparametrization $\Sigma = \begin{bmatrix} 1 & \rho \\ \rho & \phi + \rho^2 \end{bmatrix}$. Using their results, $\rho | \phi, \beta, \gamma, y_{i,t}^*, w_{i,t} \sim N((\sum_{i,t} \varepsilon_{i,t}^2)^{-1}(\sum_{i,t} \varepsilon_{i,t} \eta_{i,t}), \sqrt{\phi(\sum_{i,t} \varepsilon_{i,t}^2)^{-1}})$. To obtain a value for ϕ , use $\phi | \rho, \beta, \gamma, y_{i,t}^*, w_{i,t} \sim IG(\sum_{i,t} (\eta_{i,t} - \rho \varepsilon_{i,t})^2, N)$.

To draw values of the distribution of $w_{i,t} | \beta, \gamma, \Sigma, y_{i,t}^*, q_{i,t}$, we first calculate the full conditional distribution, which is $N(\mu_w^*, \sigma_w^*)$ with $\mu_w^* = X'\beta + \frac{\sigma_{\eta\varepsilon}}{\sigma_\eta^2}(y^* - Z'\gamma)$ and $\sigma_w^* = \sqrt{1 - \frac{\sigma_{\eta\varepsilon}^2}{\sigma_\eta^2}}$, using the conditional expectation and standard deviation in case of two normals and the indication assumption $\sigma_\varepsilon = 1$. Then, this conditional distribution is truncated using the information in $q_{i,t}$ as in the separate timing model, making sure that $w_{i,t} > 0$ if $q_{i,t} = 1$ and $w_{i,t} < 0$ otherwise.

Finally, to obtain draws of the distribution of $y_{i,t}^*$ given γ and σ_η , we again use a conditional normal distribution $y^* \sim N(\mu_y^*, \sigma_y^*)$, with $\mu_y^* = Z'\gamma + \sigma_{\eta\varepsilon}(w - X'\beta)$ and $\sigma_y^* = \sqrt{\sigma_\eta^2 - \sigma_{\eta\varepsilon}^2}$, again using that $\sigma_\varepsilon = 1$.

Convergence criteria

Table A3 shows the convergence criteria for the combined model. As before, we have split the data into two parts and tested whether the means are equal to see if the draws have converged, using again an AR(1) term in the testing equation. As can be seen, for one of the 21 parameters, convergence is rejected at the 5% level (as ρ is a function of σ_η^2 and $\sigma_{\eta\varepsilon}$). This is not directly a sign that the total system has not converged, as the probability of rejecting one of 21 hypotheses at the 5% level is quite large.

Also shown are the effective sample sizes for the different parameters. For the coefficients, the autocorrelation in the draws is quite high (between 0.74 and 0.97), so the effective sample size is quite a bit smaller than the number of draws. Still, the distributions of all parameters are effectively based on at least 1500 almost-independent observations, which is large enough to have a satisfactory sample distribution. For the covariance, the autocorrelation of the draws is even higher (almost 1), indicating that for this parameter the number of draws might have been too low.

Variable	Mean 1st half	Mean 2nd half	P-value dummy	AR(1) parameter	k	ESS
<i>Value model</i>						
Constant	0.049	0.049	0.283	-0.000 (0.003)	< 1	100000
MF_{t-1}	1.055	1.055	0.253	0.002 (0.003)	< 1	100000
p_{t-1}	-0.108	-0.108	0.313	-0.002 (0.003)	< 1	100000
<i>Timing model</i>						
Constant	1.056	1.065	0.334	0.950 (0.001)	44.71	2336.69
TSR	-0.135	-0.136	0.120	0.958 (0.001)	54.43	1837.30
TUR	-0.202	-0.202	0.442	0.950 (0.001)	44.87	2228.51
TQ1	-0.147	-0.147	0.777	0.940 (0.001)	37.33	2679.13
TQ2	-0.096	-0.097	0.699	0.951 (0.001)	45.91	2178.07
TQ3	-0.089	-0.089	0.443	0.952 (0.001)	47.02	2126.68
RET	2.856	2.849	0.741	0.936 (0.001)	34.77	2875.90
AbsRET	5.478	5.444	0.090	0.937 (0.001)	35.62	2807.60
SQL	0.049	0.049	0.628	0.889 (0.001)	19.51	5125.41
DMF	-0.637	-0.637	0.903	0.837 (0.002)	12.95	7722.84
AbsDMF	0.716	0.717	0.912	0.847 (0.002)	13.83	7233.10
ACT	-0.823	-0.825	0.342	0.966 (0.001)	66.03	1514.38
TSU	-0.072	-0.071	0.265	0.958 (0.001)	53.83	1857.83
QL	-1.077	-1.076	0.961	0.933 (0.001)	33.14	3017.38
DYY	0.006	0.006	0.328	0.857 (0.002)	14.91	6707.85
AbsDYY	-0.008	-0.008	0.324	0.740 (0.002)	7.64	13097.49
<i>Covariance</i>						
σ_η	0.113	0.114	0.851	0.994 (0.000)	371.92	268.88
$\sigma_{\eta\epsilon}$	0.005	0.007	0.010	0.999 (0.000)	3773.58	26.50
$\rho_{\eta\epsilon}$	0.014	0.021	0.030	0.999 (0.000)	3761.24	26.59

Table A3: Convergence criteria of the draws for the combined model which accounts for correlation, based on 100000 draws. Results are based on an AR(1)-model with a dummy indicating whether a draw belongs to the first half or the second half of the draws. As there might be serial correlation in the draws, we use Newey-West HAC standard errors. The table presents the mean of the first 50000 and the last 50000 draws, the p-value of the dummy and the estimated AR(1) parameter ϕ . The value of k indicates the lowest value for which $\phi^k < 0.10$, using which the effective sample size (ESS) can be calculated.

References

- Bolliger, G. (2004). The characteristics of individual analysts forecasts in Europe. *Journal of Banking & Finance*, 28(9):pp. 2283 – 2309.
- Brown, L. D. (1993). Earnings forecasting research: its implications for capital markets research. *International Journal of Forecasting*, 9(3):pp. 295 – 320.
- Brown, L. D. and Hugon, A. (2009). Team earnings forecasting. *Review of Accounting Studies*, 14:pp. 587–607.
- Clement, M. B., Hales, J., and Xue, Y. (2011). Understanding analysts’ use of stock returns and other analysts’ revisions when forecasting earnings. *Journal of Accounting and Economics*, 51(3):pp. 279 – 299.
- Clement, M. B. and Tse, S. Y. (2005). Financial analyst characteristics and herding behavior in forecasting. *The Journal of Finance*, 60(1):pp. 307–341.
- Cooper, R. A., Day, T. E., and Lewis, C. M. (2001). Following the leader: a study of individual analysts earnings forecasts. *Journal of Financial Economics*, 61(3):pp. 383 – 416.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum Likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):pp. 1–38.
- Easterwood, J. C. and Nutt, S. R. (1999). Inefficiency in analysts’ earnings forecasts: Systematic misreaction or systematic optimism? *The Journal of Finance*, 54(5):pp. 1777–1797.
- Guttman, I. (2010). The timing of analysts earnings forecasts. *The Accounting Review*, 85(2):pp. 513–545.
- Jegadeesh, N. and Kim, W. (2010). Do analysts herd? An analysis of recommendations and market reactions. *Review of Financial Studies*, 23(2):pp. 901–937.

- Kim, Y., Lobo, G. J., and Song, M. (2011). Analyst characteristics, timing of forecast revisions, and analyst forecasting ability. *Journal of Banking & Finance*, 35(8):pp. 2158 – 2168.
- Landsman, W. R., Maydew, E. L., and Thornock, J. R. (2012). The information content of annual earnings announcements and mandatory adoption of IFRS. *Journal of Accounting and Economics*, 53(12):pp. 34 – 54.
- McCulloch, R. E., Polson, N. G., and Rossi, P. E. (2000). A Bayesian analysis of the multinomial probit model with fully identified parameters. *Journal of Econometrics*, 99(1):pp. 173 – 193.
- Ramnath, S., Rock, S., and Shane, P. (2008). The financial analyst forecasting literature: A taxonomy with suggestions for further research. *International Journal of Forecasting*, 24(1):pp. 34 – 75.
- Schipper, K. (1991). Analysts forecasts. *Accounting Horizons*, 5(4):pp. 105–121.
- Sheng, X. and Thevenot, M. (2012). A new measure of earnings forecast uncertainty. *Journal of Accounting and Economics*, 53(12):pp. 21 – 33.
- Zellner, A. (1971). *An Introduction to Bayesian Inference in Econometrics*, volume 17. John Wiley and Sons.
- Zellner, A. and Ando, T. (2010). A direct Monte Carlo approach for Bayesian analysis of the seemingly unrelated regression model. *Journal of Econometrics*, 159(1):pp. 33 – 45.