UNIVERSITY OF CALIFORNIA, SAN DIEGO

Essays on Financial Analysts' Forecasts

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Economics

by

Marius del Giudice Rodriguez

Committee in charge:

Professor Allan Timmermann, Chair
Professor Graham Elliott
Professor Nir Jaimovich
Professor Bruce N. Lehmann
Professor Rossen Valkanov

2006

The dissertation of Marius del Giudice Rodriguez is approved, and it is acceptable in quality and form for publication on microfilm:

Chair

University of California, San Diego

2006

To my parents

Fernando and Cyntia Rodriguez

"*Successful investing is anticipating*

the anticipation of others"

– Sir John Maynard Keynes (1883–1946)

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

At UCSD, I had the opportunity to interact and learn from exceptional researchers. Allan Timmermann has provided me with great guidance. Allan taught me not only a great deal of econometrics and finance, but also that any serious academic work should be presented with clarity and precision. I thank Bruce, for always posing challenging questions and for instigating profitable discussions at his Starbucks office. I also would like to express a deep gratitude to the other committee members, Graham Elliot, Nir Jaimovich and Rossen Valkanov.

In San Diego, a lot of friends turn the hard life of a graduate student into something much nicer than what I expected. Jennifer Poole and Marco Aiolfi deserve a special thanks for their friendship and understanding. I also thank many of my classmates and colleagues at UCSD and in San Diego, in particular, Juri Marcucci, Carlos Capistran, George Monokrossous, Kevin King, Giuseppe Ragusa, Francesca Mazzolari, Gary Hardiman, Celine Charat, Daniel Lima, Marcelo Reginato and Kevin Sheppard. I also would like to thank all the staff members of the Department, in particular, Devaney Kerr for her assistance.

I would have not reached this far if it was not for the encouragement and support of my parents and family. To my parents Fernando and Cyntia Rodriguez I owe the greatest thanks. My sisters, Fernanda and Renata Rodriguez, and Catia Gieler all have been a part of this endeavor and deserve special thanks as well.

Finally, I would like thank the University of California at San Diego, for hosting me all these years and for providing me with the solid pillars of my future professional career.

Chapter 2 is based on *Managing Earnings Expectations: Persistence, Asymmetry and Predictability in Analysts' Earnings Forecasts* joint with Marco Aiolfi and Allan Timmermann.

Chapter 3 is based on *Financial Analysts' Forecast Revisions and Macroeconomic Information* joint with Marco Aiolfi.

VITA

1996	B. A., Universidade de Brasília
1999	M. A., Universidade de Brasília
2001–2006	Teaching assistant, Department of Economics, University of California San Diego
2004	C. Phil., University of California San Diego
2006	Ph. D., University of California San Diego

ABSTRACT OF THE DISSERTATION

Essays on Financial Analysts' Forecasts

by

Marius del Giudice Rodriguez
Doctor of Philosophy in Economics

University of California San Diego, 2006

Professor Allan Timmermann, Chair

This dissertation contains three self-contained chapters dealing with specific aspects of financial analysts' earnings forecasts.

After recent accounting scandals, much attention has turned to the incentives present in the career of professional financial analysts. The literature points to several reasons why financial analysts behave overoptimistically when providing their predictions. In particular, analysts may wish to maintain good relations with firm management, to please the underwriters and brokerage houses at which they are employed, and to broaden career choice. While the literature has focused more on analysts' strategic behavior in these situations, less attention has been paid to the implications these factors have on financial analysts' loss functions. The loss function dictates the criteria that analysts use in order to build their forecasts. Using a simple compensation scheme in which the sign of prediction errors affect their incomes differently, in the first chapter we examine the implications this has on their loss function. We show that depending on the contract offered, analysts have a strict preference for under-prediction or over-prediction and the size of this asymmetric behavior depends on the parameter that governs the financial analyst's preferences over wealth. This is turn affects the bias in their forecasts. Recent developments in the forecasting literature allow for the estimation of asymmetry parameters after observing data on forecasts. Moreover, they allow for a more general test of rationality once asymmetries are present. We make use of forecast data from financial analysts, provided by I/B/E/S, and present evidence of asymmetries and weak

evidence against rationality.

In the second chapter we study the evolution over time in the revisions to financial analysts' earnings estimates for the 30 Dow Jones firms over a 20 year period. If analysts' forecasts used information efficiently, earnings revisions should not be predictable. However, we find strong evidence that earnings revisions can in fact be predicted by means of the sign of the last revision or by using publicly available information such as short interest rates and past revisions. We propose a three-state model that accounts for the very different magnitude and persistence of positive, negative and 'no change' revisions and find that this model forecasts earnings revisions significantly better than an autoregressive model. We also find that our forecasts of earnings revisions predict the actual earnings figure beyond the information contained in analysts' earnings estimates.

Finally, the empirical literature on financial analysts' forecast revisions of corporate earnings has focused on past stock returns as the key determinant. The effects of macroeconomic information on forecast revisions is widely discussed, yet rarely tested in the literature. In the third chapter, we use dynamic factor analysis for large data sets to summarize a large cross-section of macroeconomic variables. The estimated factors are used as predictors of the average analyst's forecast revisions for different sectors of the economy. Our analysis suggests that factors extracted from macroeconomic variables do, indeed, improve on the current model with only past stock returns. In trying to explain what drives financial analysts' forecast revisions, the factors representing the macroeconomic environment must be considered to avoid a potential omitted variable problem. Moreover, the explanatory power and direction of such factors strongly depend on the industry in question.

1

Financial Analysts' Incentives and Forecast Biases

1.1 Introduction

The properties of forecasts made by financial analysts have been the subject of extensive analysis. They represent a measure of market expectations and provide important information about the fundamentals of listed firms. Therefore, in order to determine the added value of such information, it is necessary to investigate whether or not such predictions make use of all available information. If, for example, forecasts fail to use all information, this might serve as an explanation for market behavior that does not fit the usual assumptions of asset pricing. For instance, Dreman and Berry (1995) argue stock prices and earnings surprises are linked in such a way that very small percentage forecast errors by financial analysts may cause large changes in prices. Indeed, many market professionals consider a forecast error magnitude of plus or minus ten percent of forecast earnings enough to trigger a major reaction in stock prices. So, it is of great importance to uncover how these predictions are constructed and how financial analysts' incentives drive their forecasting behavior.

In this chapter, we attempt to uncover the effect incentives have on financial analysts' loss functions, and their forecast biases. We then ask the question: are financial analysts rational, given these incentives? This chapter's contributions to the extensive literature on the rationality of financial analysts' forecasts are manifold. First, we develop a wage contract between the forecast provider and the forecast user in which the forecaster is assumed to be penalized asymmetrically for forecast errors of different signs. Second, this compensation scheme gives rise to an asymmetric loss function for the financial analyst. Third, following Elliott, Komunjer, and Timmermann (2005b), we propose a methodology to estimate the asymmetric loss and jointly test for rationality. Finally, we demonstrate that analysts with asymmetric incentives, indeed appear rational once asymmetries are accounted for.

Previous literature has tested the rationality of financial analysts using quadratic and symmetric or linear and symmetric loss functions. However, financial analysts' incentives may imply a loss function that does not satisfy these properties. In particular, analysts may wish to maintain good relations with firm management, to please the underwriters and brokerage houses at which they are employed, and to broaden career choice. The evidence presented in Dugar and Nathan (1995), Dechow, Hutton, and Sloan (1999) and Michaely and Womack (1999) suggests that analysts whose forecasts are used by brokers with an underwriting relationship with the firm whose earnings are being predicted seem to provide more optimistic forecasts. This is easily translated into a situation in which the forecaster is endowed with an asymmetric loss function when performing his estimation and forecast selection.

Another type of incentive has been suggested by Lin and McNichols (1998) and Cowen, Groysberg, and Healy (2003) who find that analysts hired by brokerage firms involved in sales and trading provide more optimistic forecasts. Hong and Kubik (2003) relate analysts' optimism to career concerns. In their study, analysts who are more optimistic, relative to the consensus, experience favorable job separations. Finally, from a corporate finance point of view, Lim (2001) and Francis and Philbrick (1993) point to the fact that, in order to get better access to the firm's managers, analysts might be tempted to issue favorable opinions. The classic case is of an Enron analyst who was fired after Enron management put pressure on the analyst's firm (see McLean and Elkind (2003)).[1]

[1]Asymmetric loss is not the only explanation for biased forecasts. There are several results relating forecast biases to strategic interactions, such as Ehrbeck and Waldman (1996), Lamont (1995), Scharfstein and Stein (1990), Trueman (1994) and Laster, Bennett, and Geoum (1999). Elliott, Komunjer, and

If analysts are, in fact, faced with these incentives, it is naive to assume that their loss function penalizes forecast errors of different signs symmetrically. By allowing for a more flexible loss function, forecast errors may display some degree of bias which should not be linked to the lack of rationality, as shown in Zellner (1986), Christoffersen and Diebold (1997), Granger (1969), Granger (1999), Patton and Timmermann (2002), and Varian (1974).

Using a loss function that penalizes forecast errors of different signs in distinct ways has practical implications for testing rationality. The vast literature devoted to this task uses ordinary least squares (OLS)[2] or least absolute deviation (LAD)[3] procedures. By using these methodologies, the loss function of the analyst is constrained to be quadratic or linear with symmetric penalties. In order to test rationality for a flexible loss function, Elliott, Komunjer, and Timmermann (2005b) propose to construct a test by evaluating a transformation of the forecast errors reflecting asymmetries in the loss function.[4] However, this approach lacks general economic intuition.

We propose a compensation scheme that represents a contractual relationship between the forecast provider and the forecast user. The compensation mechanism is such that analysts' incomes are a function of the magnitude and the sign of the forecast error, plus a fixed income component. We demonstrate that, for all levels of risk aversion, an asymmetric compensation scheme, in which analysts face higher penalties if they under-predict the value of earnings per share leads to a loss function which displays an asymmetric relationship. Using the loss function derived from analysts' risk preferences for shocks to their wealth, we investigate the implications these asymmetric incentives have for the forecaster's biases with two numerical exercises. As expected, we find that higher values of the contractual parameter are associated with higher positive forecast biases implying over-prediction. Furthermore, for a fixed value of the contractual parameter, the size of the bias is directly related to the level of the analysts' risk aversion. More specifically, the magnitude of the forecast bias is greater for more risk averse ana-

Timmermann (2005a) suggest that the two explanations may be complements.

[2]See, for example, Keane and Runkle (1998), De Bondt and Thaler (1990), Abarbanell and Bernard (1992), Easterwood and Nutt (1999), Elliott, Philbrick, and Weidman (1995), Mendehall (1991), Ali, Klein, and Rosenfeld (1992), Klein (1990), Aitken, Frino, and Winn (1996), and Lys and Sohn (1990).

[3]Basu and Markov (2003)

[4]Rodriguez (2004) uses this methodology to test for rationality under an asymmetric and linear or quadratic loss function and fails to reject rationality.

lysts. Then, following the methodology in Elliott, Komunjer, and Timmermann (2005b), we can recover the asymmetry parameter and test for financial analyst rationality using data from I/B/E/S on forecasts reported by financial analysts. The estimates suggest that the loss function becomes more asymmetric, the more risk averse the analyst. Finally, when accounting for a more realistic and asymmetric loss function, we fail to reject the test for rationality among financial analysts. Therefore, once we allow for asymmetric incentives, financial analysts appear to present weaker evidence against rationality. Moreover, the size and direction of the asymmetric incentives strongly depend on the assumed risk aversion of the analyst.

The chapter is organized as follows. In the next section, we present a compensation scheme that affects analysts' wealth in a way that depends on the sign of the forecast error. We then go on to show how an analyst's loss function depends on risk preferences over shocks to wealth and how the asymmetry in the loss function, derived from the contractual parameter, translates into biased forecasts. Section 1.3 proposes a method to identify the asymmetry parameter using only available forecast information. We then offer a test for rationality conditioned on the family of loss functions assumed. We describe the financial analyst forecasting data provided by I/B/E/S in Section 1.4 and report descriptive statistics. In Section 1.5, we present results from the empirical estimation performed by the Generalized Method of Moments (GMM) of Hansen (1982), including a separate analysis for financial analysts in high status brokerage houses. Following the accounting literature, we focus on lagged forecast errors, lagged earnings changes, and lagged U.S. GDP growth as distinct choices for instruments. The final section concludes the chapter.

1.2 Analysts' Incentives and Their Loss Functions

One of the main tasks of financial analysts is to provide information about firms' fundamentals. Analysts are hired to study how the macroeconomic climate and industry-specific variables affect a company's performance. That is, based on their information set, analysts must provide predictions of the firm's economic health. The measure most commonly used to express firm performance is the level of earnings per share.

We consider that analysts are trying to forecast a random variable, Y_{t+1}. They use a set of instruments or information variables to compute their forecast, f_{t+1}, which gives rise to a forecast error $e_{t+1} = Y_{t+1} - f_{t+1}$. A positive forecast error, $e_{t+1} > 0$, suggests the analyst under-predicted earnings per share, while a negative forecast error, $e_{t+1} < 0$, implies over-prediction.

1.2.1 The Compensation Scheme

Agents from underwriting institutions, brokerage houses, and investment institutions hire analysts and use their forecasts for many purposes. These forecast users, therefore, have a direct interest in how estimates are constructed. For instance, they may have a clear incentive for the forecasts to be optimistic, due to the nature of the investment business. A simple way to achieve this goal is to offer the analyst a compensation scheme that penalizes over-prediction and under-prediction differently. In this chapter, we design a contract in which the analyst's wage is a function of his forecast errors, plus a fixed wage component. Furthermore, forecast errors of different signs are penalized asymmetrically.[5]

Suppose that $\overline{w}$ is the potential wage a financial analyst could earn. In order to give the right incentives, we assume that forecast users offer a wage given by

$$w_{t+1} = \overline{w} - \alpha|e_{t+1}|1_{\{e_{t+1}>0\}} - (1-\alpha)|e_{t+1}|1_{\{e_{t+1}<0\}}. \tag{1.1}$$

Under this compensation scheme, larger forecast errors of either sign are more strongly penalized when $\alpha \in (0, 1)$. It is in the interest of the employer to set contract parameters within this, interval. Also, the forecaster achieves the potential wage $\overline{w}$, only when he commits no errors. The parameter α dictates how forecast errors of different signs affect

[5]Bernhardt, Campello, and Kutsoati (2004) and Bernhardt and Kutsoati (1999) also propose a compensation scheme in which the analyst's real wage depends on his absolute predictive ability and his relative performance with respect to the consensus.

total wealth.[6] We can rewrite (1.1) as

$$w_{t+1} = \overline{w} - \alpha|e_{t+1}| - (1-2\alpha)|e_{t+1}|1_{\{e_{t+1}<0\}}. \tag{1.2}$$

If α is between 0.5 and 1, analysts have a stronger incentive for producing negative forecast errors, so it should be expected that forecasts reflect some degree of optimism (recall, $e_{t+1} = Y_{t+1} - f_{t+1}$). If on the other hand, α is between 0 and 0.5, analysts have a preference towards under-prediction. Compensation schemes such as these can be motivated by incentives that rely on the use of derivatives. The bonus associated with negative forecast errors mimics the payoff of a put option.

Finally, analysts have preferences over wealth described by a differentiable and concave utility function U, satisfying the conditions $U'(\cdot) > 0$ and $U''(\cdot) \leq 0$. Within this framework, it remains to identify the link between the incentive scheme offered by the forecast user, analysts' risk preferences, and the bias in the forecast. Next, we derive the asymmetric loss function analysts use to construct forecasts when faced with the above compensation scheme.

1.2.2 The Loss Function

Granger and Newbold (1986) argue that to obtain any value for a point forecast, we require a criterion against which various alternatives can be judged. A natural way to proceed is to define a loss function. Formally, Granger (1999) defines a loss function to be:

Definition 1.1. A loss function is a real valued function L of the forecast error e_{t+1} that satisfies the following properties:

- $L(0) = 0$
- $\min_{e_{t+1}} L(e_{t+1}) = 0$ so $L(e_{t+1}) \geq 0$

[6] Beyond the scope of this chapter, we could conceive of a contractual model which defines α as a function of several state variables, such as the likelihood of stock sales. In this chapter, we appeal to basic economic intuition and anecdotal evidence that over-prediction and under-prediction have different effects on analysts' incomes.

- L is nondecreasing for $e_{t+1} > 0$ and nonincreasing for $e_{t+1} < 0$, i.e. L is nondecreasing in $|e_{t+1}|$

The objective of this chapter is to study the forecasting behavior of analysts under the asymmetric incentives presented by the contractual relationship. In order to translate this contractual relationship into a loss function with the above properties, we introduce decision-based forecast theory. Manski (1991) suggests that the choice of a loss function should reflect the actual loss associated with mispredictions. With this in mind, we proceed, following the idea presented in Granger and Machina (2006), where for a forecaster with a particular utility function, loss is the difference between the utility under a zero forecast error and the utility under the actual forecast error. More formally, for the contractual relation described above, the loss function is defined as

$$L(e_{t+1}) = U(\overline{w}) - U(w_{t+1}) \tag{1.3}$$

When analysts are confronted with asymmetric wage contracts, they may use (1.3) as a criterion to estimate and judge the performance of their forecasting models. However, it remains to be shown that (1.3) satisfies the properties of a loss function as characterized by Definition 1.1.

Proposition 1.2. *The function defined in (1.3) satisfies the properties of a loss function presented in Definition 1.1 if* $\alpha \in (0, 1)$ *and* U *is such that* $U'(\cdot) > 0$.

Given appropriate values of α for the forecast user, the analyst's loss function conforms to the Granger (1999) properties of a loss function. More formally, Proposition 1.2 suggests (1.3) satisfies the properties of a loss function if forecast users do not choose contractual parameters in violation of their own interests, that is $\alpha \notin (0, 1)$. The underlying implication is that employers should not give incentives that allow analysts to ignore the precision of their estimates.[7]

The choice of utility function has important implications for the loss function in (1.3). We derive the loss function and the implied forecast bias under the assumption financial analysts are risk neutral and risk averse. If analysts are risk neutral with respect

[7] Anecdotal evidence and basic economic intuition suggest this is not untrue.

to negative shocks to their wealth, then the utility function should take the linear form, $U(w_{t+1}) = \varsigma + \gamma w_{t+1}$ where we assume ς and γ are strictly positive. This utility function implies a loss function of the form

$$L(e_{t+1}) = \gamma\alpha|e_{t+1}|1_{\{e_{t+1}>0\}} + \gamma(1-\alpha)|e_{t+1}|1_{\{e_{t+1}<0\}}. \tag{1.4}$$

This is the well known loss function used in quantile predictions proposed by Koenker and Basset (1978). The bias derived from this loss function equals the difference between the mean and the $\alpha - quantile$ of the random variable to be forecast. For the purpose of uncovering the asymmetry parameter, we restrict γ to be equal to one for the case of risk neutral financial analysts.

For the risk averse analyst, the result is not as simple. The degree of curvature of the utility function has direct consequences for how analysts are penalized. In this chapter, we restrict the preferences of risk averse forecasters to be negative exponential, $U(w_{t+1}) = 1 - e^{-\rho w_{t+1}}$. From this utility, we derive the loss function

$$L(e_{t+1}) = e^{-\rho w_{t+1}} - e^{-\rho \overline{w}}. \tag{1.5}$$

Under this utility specification, the fixed wealth component of the compensation scheme is distinct from the wage components determined by the forecast errors. This is needed to ensure the estimation since we do not observe an analyst's potential wage.[8]

Figure 1.1 plots the shape of the loss function for reasonable values of risk aversion, a contractual parameter equal to 0.5 and a fixed wealth equal to 0.2. Notice that for higher values of risk aversion, the loss is flatter for smaller forecast errors. This stems from the fact that analysts place more weight on the fixed wealth component of their wage contract and, therefore, are less penalized for small forecast errors. However, as the forecast errors increase in size, more risk averse analysts suffer larger penalties. The analyst's risk aversion directly translates into how responsive the loss function is to the magnitude of the forecast error.

In Figure 1.2, we increase the value of the fixed wealth component of the wage con-

[8]The unfortunate implication is that this approach would not work for other common utility functions, such as power utility.

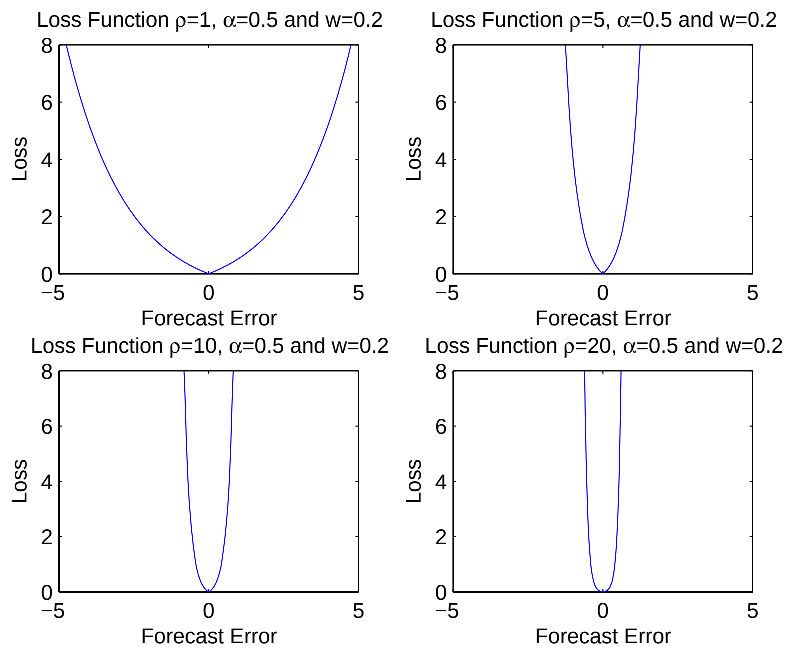

Figure 1.1: Shape of Symmetric Loss Functions for Risk Averse Analysts for $\overline{w} = 0.2$

tract to 0.3. The loss function becomes flatter in the neighborhood of zero for different levels of risk preference parameters. The forecast errors now have a smaller impact on the analyst's overall wealth. Therefore, analysts are less penalized for their mispredictions as the potential wage increases.

Finally, Figure 1.3 demonstrates the case in which over-prediction and under- prediction affect the loss function in distinct ways. For the case in which the contractual parameter, α, equals to 0.7, analysts face higher penalties if they under-predict the value of earnings per share. For all levels of risk aversion, analysts' loss functions display an asymmetric relationship. That is, analysts are penalized more for under-prediction than for over-prediction by the same amount. These properties, displayed in Figures 1.1 - 1.3, are summarized in the following proposition.

Proposition 1.3. *For the loss function defined in (1.5) we have:*

1. *For given values of contractual and preference parameters, (α, ρ), the higher the fixed wealth, the lower the loss.*

2. *For given levels of preference parameter, ρ, and fixed wealth, $\overline{w}$, if α is different*

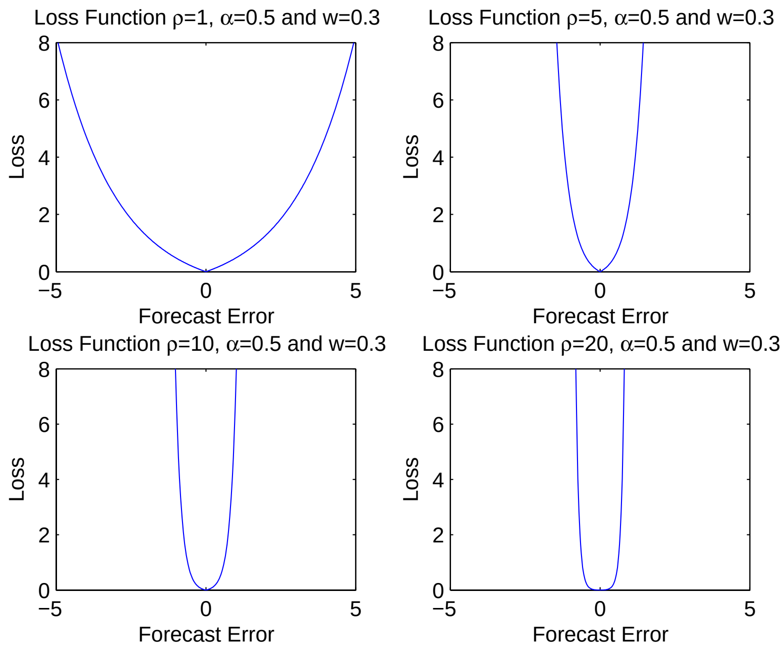

Figure 1.2: Shape of Symmetric Loss Functions for Risk Averse Analysts for $\overline{w} = 0.3$

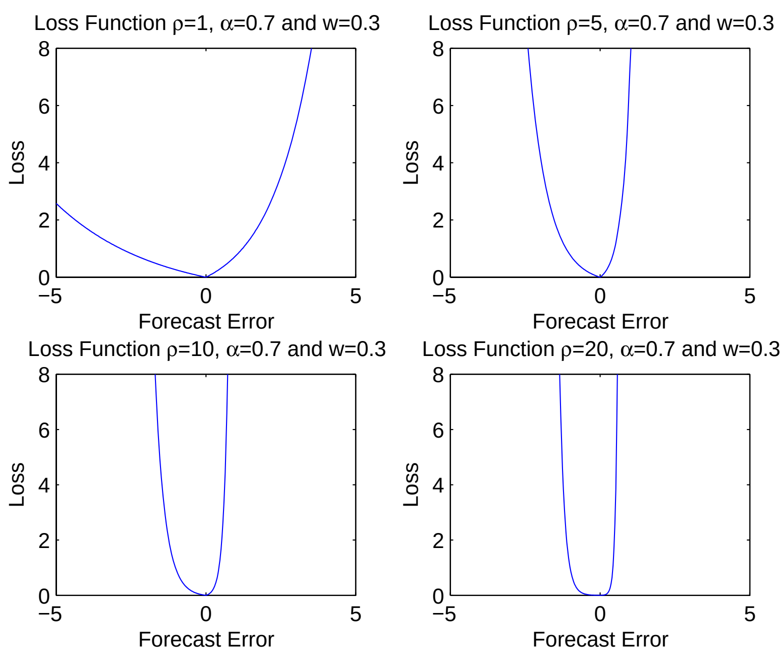

Figure 1.3: Shape of Asymmetric Loss Functions for Risk Averse Analysts where $\overline{w} = 0.3$

than 0.5, the loss is asymmetric.

1.2.3 Asymmetric Loss and Forecast Bias

As the previous section illustrates, analysts who are faced with asymmetric compensation contracts follow an asymmetric loss function when constructing their market predictions. We demonstrated that for all levels of risk aversion, analysts' wage rates in the next period will be penalized more for under-prediction than for over-prediction. It is clear this asymmetric relationship has important implications for analysts' forecast biases.

For any loss function, the forecaster chooses the point forecast that minimizes the expected loss with respect to a given information set. In other words, the chosen forecast should minimize statistical risk. Formally, we have

$$\min_{f_{t+1}} E_t[L(e_{t+1})]. \tag{1.6}$$

Zellner (1986), Christoffersen and Diebold (1997), Granger (1969), Granger (1999), Patton and Timmermann (2002), and Varian (1974) show how asymmetries in the loss function relate to forecast bias for linear, quadratic, and linear-exponential loss functions. Similarly, for our loss function defined in (1.5), the next proposition proves that if the loss function is asymmetric, the resulting forecasts will be biased.

Let the data generating process for Y_{t+1} take the form

$$Y_{t+1} \quad = \quad \mu_{t+1|t} + \sigma_{t+1|t}\varepsilon_{t+1}$$

where $\varepsilon_{t+1}|\mathcal{F}_t \sim G_{t+h|t}$, $E[\varepsilon_{t+1}|\mathcal{F}_t] = 0$, $E[\varepsilon^2_{t+1}|\mathcal{F}_t] = 1$ and restrict ε_{t+1} to have a symmetric conditional density, $g_{t+1|t}$.

Proposition 1.4. *For the general data generating process above, the forecast derived from the loss function defined in (1.5) is biased as long as* $\alpha \neq 0.5$*.*

The main problem with decision-based forecasting when analysts may be risk averse, is that there is no closed form solution to (1.6) for commonly assumed distributions. The exception is the trivial case of the linear loss function defined in (1.4). So, we rely on numerical experiments to derive the relationship between the asymmetric loss function and analysts' forecast biases for risk averse analysts. We define two separate measures of expected forecast bias: the forecast minus the sample mean, and the forecast minus the $\alpha - quantile$ of the sample distribution.

First, we obtain 50 samples of size 1000 from a standard normal distribution. Recall, analysts' risk preferences are assumed to be negative exponential, leading to the loss function defined in (1.5). We restrict the risk aversion parameter to four different values: 1, 5, 10 and 20. Then, given a fixed preference parameter, we let the contractual parameter, α, vary from 0.25 to 0.75. By performing a grid search over all possible combinations of the preference and contractual parameters, we compute the optimal prediction as the value that minimizes the statistical risk in (1.6). Finally, for each sample, we compute the bias in the forecast. We use the average of the bias across the 50 samples as the measure of expected bias for each combination of preference and contractual parameters.

Theory suggests that the more asymmetric the compensation contract offered to the analyst, the more biased will be his forecasts. Recall from (1.2), if α is between 0.5 and 1, analysts have an incentive for optimism in their predictions. If on the other hand, α is between 0 and 0.5, analysts have a preference towards under-prediction. Therefore, we expect the results of the numerical exercise to show that for higher values of α, since analysts will be more penalized for under-prediction, independent of their levels of risk aversion, they will tend to over-predict the variable. Similarly, for lower levels of α, we expect analysts will tend to under-predict the variable.

The expected forecast bias defined by the forecast minus the sample mean measures the differences in forecasts between risk averse analysts and those implied by a mean square error loss function. We compute the optimal forecast made by the risk averse analyst for the four different levels of risk aversion and different contractual parameters. We then define the forecast bias as the difference between the forecast made by the risk averse analyst and the sample mean. Finally, we compute the mean of this difference and define it to be our measure of expected bias.

Figure 1.4 presents this measure of expected bias as a function of the contractual

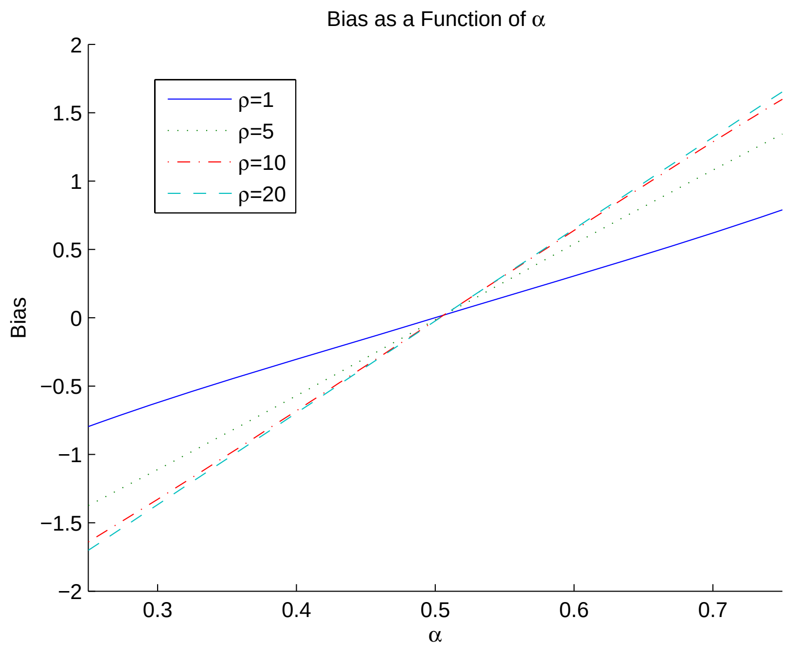

Figure 1.4: Asymmetric Loss, Risk Aversion and Expected Forecast Bias (forecast minus mean)

parameter, α, for each of the four parameter values of risk aversion, ρ. As expected, higher values of α are associated with higher positive forecast biases, implying over-prediction. Similarly, lower values of α are associated with a larger negative bias. That is, the contractual parameter drives analysts' forecast biases, independent of their preferences for risk. However, for a fixed value of the contractual parameter, the size of the bias is directly related to the level of risk aversion. More specifically, the magnitude of over- and under-prediction is greater for more risk averse analysts.

The expected forecast bias defined by the forecast minus the $\alpha - quantile$ of the sample distribution measures the differences in forecasts between risk averse and risk neutral analysts, given contractual parameters. As before, we compute the optimal forecast made by a risk averse analyst for the four different levels of risk aversion and different contractual parameters. At the same time, we compute the forecast made by a risk neutral analyst for different values of α; recall, this forecast is just the $\alpha - quantile$ of the distribution. We then define the measure of bias to be the difference between the forecast made by the risk averse analyst and the forecast made by the risk neutral analyst for each value of α. Finally, we compute the mean of this difference across samples and define it to be our measure of expected bias.

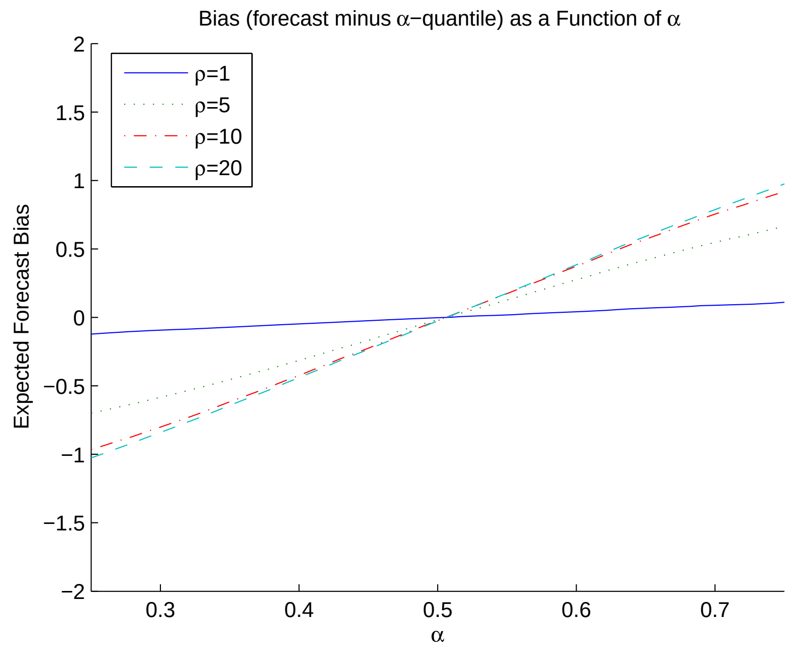

Figure 1.5: Asymmetric Loss, Risk Aversion and Expected Forecast Bias (forecast minus $\alpha - quantile$)

Figure 1.5 presents a plot of this measure of expected forecast bias as a function of α, for each of the four preference parameters. We notice that the difference between the forecasts made by risk averse analysts and risk neutral analysts is larger for higher values of risk aversion and extreme values of contractual parameters. For instance, when α is higher, more risk averse analysts provide forecasts far greater than the risk neutral analyst. The intuition for this relies on the fact that analysts with high levels of risk aversion are more sensitive to under-prediction and, therefore, put more weight on over-predicting the variable. For the analyst with a lower degree of risk aversion, however, this difference is smaller since he is more willing to accept forecast errors.

These results are highly intuitive and coincide with anecdotal evidence from the market. Analysts with higher values of preference parameters are more penalized for large forecast errors. The greater the asymmetry in their compensation contract, that is, the further α is from 0.5, the more extreme will be their forecasts and, therefore, the more biased. The main conclusion derived from these experiments is that if forecast users want to ensure that analysts' forecasts are more optimistic, they should choose higher values of α, whereas if they want to forecasters to under-predict they should choose lower values for the contractual parameter.

1.3 Backing Out the Asymmetry

The extensive literature dedicated to testing the rationality of financial analysts generally constrains analysts' loss functions to be mean square error (MSE) or least absolute deviation (LAD). The analysis in this section builds on the work of Elliott, Komunjer, and Timmermann (2005a) and Elliott, Komunjer, and Timmermann (2005b) in order to study how analysts use the asymmetric loss function derived in the previous section to construct their forecasts. Within this framework, we develop an estimation procedure to jointly uncover the asymmetry parameter and test for financial analyst rationality using data on past forecasts.

The forecaster is faced with a general stochastic process, $X = \{X_t : \Omega \to \mathbb{R}^{m+1}, m \in \mathbb{N}, t = 1, \ldots, n+1\}$, defined on a complete probability space $(\Omega, \mathcal{F}, P)$, where $\mathcal{F}$ is a filtration determined by the σ-field generated by $\{X_s, s \leq t\}$. We assume that the forecaster is interested in predicting one component, Y_t, given the remaining components which take values in $\mathbb{R}^m$. Y_t is assumed to be continuous with a density $f(\cdot)$. The analyst's problem is to forecast, Y_{t+s}, where s is the horizon of interest, $s \geq 1$.

In this chapter, we consider only one-step ahead forecasts, $s = 1$. We define f_{t+1} to be the one-step forecast conditional on the information at time t. We restrict the study to the class of linear forecasts, that is, $f_{t+1} = \theta' Z_t$, where θ is a k-vector of unknown parameters, $\theta \in \Theta$, Θ is a compact set in $\mathbb{R}^k$, and Z_t is a k-vector of variables belonging to the information set.

When constructing a forecast at time t, the forecaster chooses a model $M = \{f_{t+1}\}$ and the set of variables to be included in the Z_t vector. The model is misspecified if either it fails to incorporate all available information in Z_t, or if the linear forecasting model fails to represent the true data generating process. However, it is worth pointing out that the researcher does not need to know the variables included in Z_t nor the model

M in order to recover the asymmetry parameter.

When faced with the compensation scheme in (1.2), analysts will construct optimal predictions based on the loss function defined in (1.3). For the analysis that follows, we restrict the utility function to be negative-exponential, leading to the loss function in (1.5). Given values for the analyst's level of risk aversion and contractual parameter that jointly describe his loss function, the analyst minimizes expected loss by choosing the optimal one-step forecast of Y_{t+1}. Equivalently, the forecaster chooses f_{t+1}, such that $f^*_{t+1} = \theta^{*\prime} Z_t$. Formally, θ^* follows from the minimization problem

$$\min_{\theta \in \Theta} E[L(e_{t+1})]. \tag{1.7}$$

The following proposition presents the relevant optimality condition.[9]

Proposition 1.5. *Under assumption (A.1.1), given values* $(\alpha, \rho) \in (0,1) \times \mathbb{R}_+$ *if* θ^* *is the minimum of (1.7), then* θ^* *satisfies the first order condition*

$$E[Z_t(1_{\{e_{t+1}(\theta^*)<0\}} - \alpha)e^{-\rho w_{t+1}(\theta^*)}] = 0$$

where $e_{t+1}(\theta^*) = Y_{t+1} - \theta^{*\prime} Z_t$ *and* $w_{t+1}(\theta^*) = -\alpha|e_{t+1}(\theta^*)| - (1-2\alpha)|e_{t+1}(\theta^*)|1_{\{e_{t+1}(\theta^*)<0\}}$.

An implication of this proposition is that when forecasts are optimal, information that is correctly incorporated into the forecast is orthogonal to a transformation of the forecast errors. The appropriate transformation is defined by the shape of the loss function.

The researcher observes a vector of variables, V_t, that are also available to the forecaster. Under rational forecasts, V_t is a subset of Z_t. Given values for the parameters (α, ρ), a sufficient test of rationality checks the moment conditions:

$$E[V_t(1_{\{e_{t+1}(\theta^*)<0\}} - \alpha)e^{-\rho w_{t+1}(\theta^*)}] = 0. \tag{1.8}$$

However, we do not observe values for analysts' contractual and risk parameters. We show that if we observe past forecasts and realized values, we can identify and estimate

[9] Please see Appendix A.1 for details on assumptions.

the asymmetry parameter and jointly test for rationality, for risk neutral and risk averse analysts with negative-exponential utility functions.

For given values of α and ρ, analysts use the moment condition derived above to determine θ^*. As we do not observe the vector of variables used in the forecasting problem, Z_t, but only the analyst's resulting optimal forecast, we must show that the moment condition defined in (1.8) links the asymmetry parameter and the optimal forecast in a unique way. Only then can we use (1.8) to uncover the asymmetry parameter. For this analysis, we hold the value of ρ fixed at different reasonable levels. Before we proceed, we provide a set of sufficient conditions for which the solution to Proposition 1.5 is a strict local minimum of the risk minimization problem. More specifically, the next proposition proves that analysts' moment conditions uniquely identify the contractual parameter.

Proposition 1.6. *Under assumptions (A.1.1) and (A.1.2), and given $(\alpha, \rho) \in (0,1) \times \mathbb{R}_+$, if $\theta^* \in \dot{\Theta}$ is a solution to the first order condition of Proposition 1.5 then θ^* is a strict local minimum on $\dot{\Theta}$.*

Finally, with an application of the implicit function theorem, the asymmetry parameter and the optimal forecast are uniquely linked as the next proposition shows.

Proposition 1.7. *Under assumptions (A.1.1) and (A.1.2), given $\rho \in \mathbb{R}_+$, there exists an open set $G \subset \dot{\Theta}$, such that, for any $\alpha \in (0,1)$,*

$$E[Z_t(1_{\{e_{t+1}(\theta^*)<0\}} - \alpha)e^{-\rho w_{t+1}(\theta^*)}] = 0 \tag{1.9}$$

has a unique solution θ^* in G and the function $\theta^* = \theta_\rho(\alpha)$ defined implicitly by Proposition 1.5 is bijective and differentiable from $(0,1)$ to G.

Propositions 1.5 - 1.7 allow for the identification of the asymmetry parameter by Generalized Method of Moments (GMM) estimation as proposed by Hansen (1982). We jointly test for rationality within this framework using a $J - test$, which checks the validity of the moment conditions for the estimated asymmetry parameter. In order to proceed with this approach, we choose a consistent estimate of the spectral density matrix, e.g., the Newey and West (1987) estimator.[10]

[10]The results are robust across different estimates of the spectral density matrix, such as those by

1.4 Data

The empirical estimation uses data on forecasts reported by financial analysts of earnings per share to uncover the asymmetry parameter and test for rationality.

The forecast data are provided by the Institutional Brokers Estimate System (I/B/E/S). I/B/E/S is a matched analyst-firm database consisting of forecasts of earnings per share as reported by financial analysts from all market institutions, covering the years 1985 through 2005. The main variables of interest are financial analysts' forecasts of earnings per share, forecast date, report date, forecast revisions, actual earnings per share, the firm's identification code, the analyst's identification code and the analyst's employer's identification code.

1.4.1 Data Restrictions

For this analysis, the data are restricted as follows. First, I/B/E/S provides analysts' forecasts for fiscal years and fiscal quarters. We focus on one-step quarterly forecasts. Next, as fiscal periods differ across firms, we consider only companies whose fiscal quarters end in March, June, September, and December, following U.S. government fiscal quarters. The final sample period ranges from the first quarter of 1985 through the last quarter of 2005, with at most 81 observations per firm.

Analysts do not provide forecasts consistently; that is, analysts do not display a continuity of forecasts over the whole sample period. This may be due to a strategic decision, as is suggested in Scherbina (2004) and Chen, Francis, and Jiang (2004). The authors argue that analysts avoid opinions that go against the consensus opinion. In fact, they may only report a forecast if the forecasted earnings per share is above a certain threshold.[11] Therefore, we restrict the set of analysts to those who report forecasts for at least 50 quarters over the entire sample period. Though this criterion may be stringent, we argue it will provide better precision in the estimates. With this restriction, we consider only analysts who are consistently involved in the forecasting profession and, thus, are actively competing for market share.

Andrews (1991) and White (1980).

[11] Scherbina (2004) notes an "unwritten rule" of financial analyst forecasting: "If you can't say something positive, don't say anything at all."

Analysts update forecasts as new information arrives. The theory outlined in the previous sections suggests analysts use all available information in generating forecasts. Moreover, in order to uncover the asymmetry parameter, forecasts must be optimal. Using the I/B/E/S data on forecast revisions, when these occur, we adopt the rule applied by Keane and Runkle (1998) to determine which forecast to consider in the analysis for each analyst. Keane and Runkle (1998) restrict the sample of forecasts to those for which they believe, with some level of certainty, that the firm's announcement from the previous quarter was known at the time the analyst made the forecast. They select only forecasts whose report dates were dated at least seven days after the firm's announcement date.

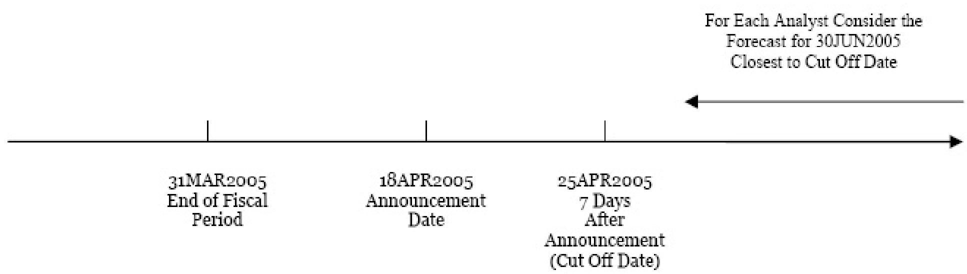

Figure 1.6: Forecast Selection Rule

Figure 1.6 illustrates the individual forecast selection rule for the I/B/E/S data. For a given fiscal period, say 30JUN2005, we select only forecasts reported seven days after the firm's announcement for the previous period, 31MAR2005. For each analyst, we consider the next period's forecast to be the forecast closest to the cutoff date. Using the seven-day cutoff rule, the final sample consists of a total of 107 analysts reporting forecasts for 202 firms. As some analysts provide forecasts for more than one firm, the full time series consists of 318 forecasts spanning 20 years.

1.4.2 Forecast Biases

The purpose of our analysis is to estimate the rationality of financial analysts' forecasts under asymmetric loss functions derived from the compensation scheme defined in (1.2). We need analysts' forecast biases in order to accomplish this task. I/B/E/S reports

Table 1.1: Forecast Errors Descriptive Statistics for the Entire Sample

This table reports descriptive statistics of the distribution of the mean forecast errors and mean forecast errors in percentage terms across analysts. The sample ranged from 31 March 1985 through 31 March 2005. We consider only analysts who provide more than 50 forecasts.

Forecast Errors:		
	Value	Percentage (%)
Mean	-0.0210	-0.7747
Median	-0.0042	-0.4817
Std.Dev.	0.0677	32.3471
Min	-0.5799	-142.7398
Max	0.0916	189.2982
Skewness	-3.8069	0.9713
Kurtosis	25.5542	10.9426

actual earnings as soon as they are released into the market. These earnings reports are then adjusted for comparability with analysts' forecasts.[12] [13] Using forecasts and actual earnings per share values, we compute the forecast errors, defined as the realized value minus the forecast.

Table 1.1 reports descriptive statistics for the full sample. The left panel presents the overall mean forecast errors. First, we call attention to the fact that both the mean and the median are negative, suggesting that analysts have preferences for over-prediction. Furthermore, the minimum and maximum demonstrate that analysts commit larger errors in absolute value when they over-predict than when they under-predict. The negative skewness suggests that the left tail of the distribution is fatter than the right tail, or large negative forecast errors occur more often than positive forecast errors. Finally, the kurtosis measure shows that the tails are fatter than a normal distribution. We plot the same data in Figure 1.7. The histogram suggests there is a small tendency towards a positive

[12]I/B/E/S does not require analysts to forecast earnings per share in the basic or diluted format, but I/B/E/S lets the majority rule. Then, in the case where an analyst follows a firm on a basis that is different from the consensus, I/B/E/S adjusts his estimates to conform with the majority.

[13]Since Philbrick and Ricks (1991), many studies, including Keane and Runkle (1998), use data for actual earnings from *Compustat* due to what they call "data alignment problems" in I/B/E/S. However, I/B/E/S (2000) states, "...I/B/E/S strives to report actual earnings as soon as they are released into the market place. For US and Canada, earnings reports are culled directly from the newswires, adjusted for comparability with estimates and reported to subscribers via the Intra Day Surprise Report, which is delivered five times each trading day..." p.7. Moreover, "...I/B/E/S adjusts reported earnings to match analysts' forecasts on both annual and quarterly basis. This is why I/B/E/S actuals may not agree with other published actuals; i.e. *Compustat*." p.8. With this in mind, we argue I/B/E/S actual earnings data are valid and appropriate for the purposes of this study.

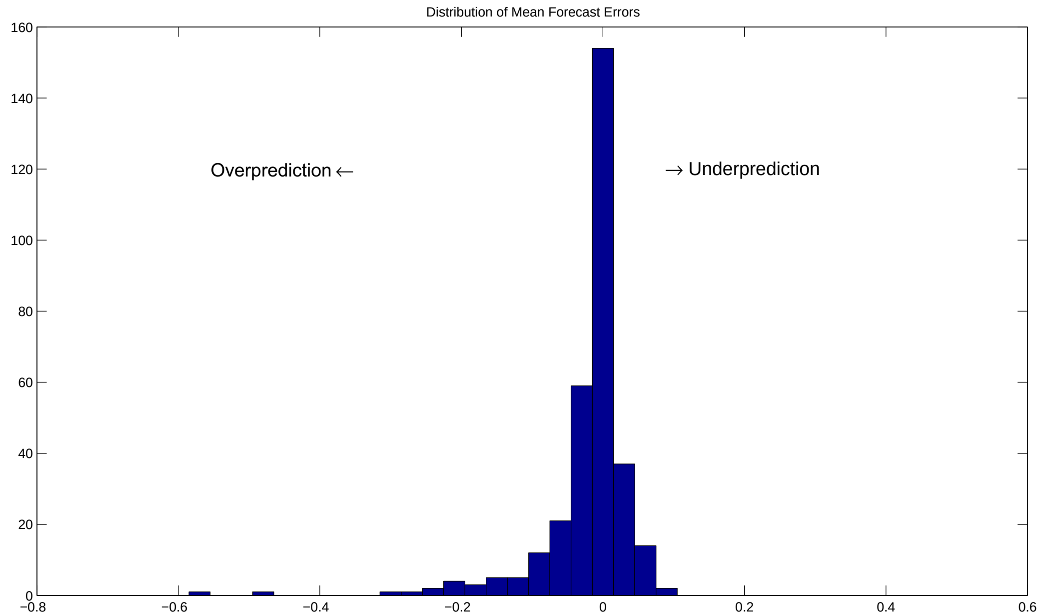

Figure 1.7: Histogram of Mean Forecast Errors

bias; that is, analysts tend to over-predict. Also, we notice there is not a tendency for extreme forecasts in either direction. The right panel of Table 1.1 reports descriptive statistics for the distribution of mean forecast errors measured in percentage terms. Here, the forecast error is computed as a percentage of the realization. Contrary to our focus on the level of the mean forecast error, the minimum and maximum for the percentage mean forecast errors demonstrate that analysts tend to under-predict. Also, the distribution approaches a normal distribution in terms of skewness and the there is a lower level of kurtosis.

1.4.3 Analysts in High Status Brokerage Firms

I/B/E/S tracks analysts and brokerage firms by numeric codes. With the permission of I/B/E/S, we were granted access to the data file translating the firm's numeric code into the institutional name. With this detailed information, we were able to identify the brokerage house that the analyst was working for at the time he reported his forecast. Carter, Dark, and Singh (1998) construct a list of high status brokerage houses in the United States, based on relative IPO performance. Hong and Kubik (2003) use

Table 1.2: List of High Status Brokerage Houses Using Carter-Manaster Ranking
An alphabetical list of the ten brokerage houses classified as high status using Carter-Manaster rankings of Carter, Dark and Singh (1998).

	Carter-Manaster Ranks
Alex Brown & Sons	8.88
Drexel Burham Lambert	8.83
First Boston Corporation	9.00
Goldman Sachs & Company	9.00
Hambrecht & Quist	9.00
Merrill Lynch	8.88
Morgan Stanley & Company	8.88
Paine Webber	8.75
Prudential-Bache	8.75
Salomon Brothers	9.00

the Carter-Manaster ranking to investigate how biased forecasts relate to favorable job separations. Table 1.2 presents the Carter-Manaster ranking of High Status Brokerage Houses from Carter, Dark, and Singh (1998).

From the full sample, we extract a sub-sample of analysts who are constrained to work in one of the top ten high status brokerage houses. Analysts often move across different houses when providing their forecasts often. Therefore, we need a criterion for selecting this sub-sample. First, we consider only analysts whose final forecast was made for one of these houses. Also, we restrict the data to analysts providing more than 20 forecasts for that house. The final sample includes a total of 54 analysts reporting a total of 174 forecasts.

Table 1.3 reports descriptive statistics for the sample of analysts from high status brokerage firms. The pattern is not remarkably different from the full sample. The one difference lies in the percentage mean forecast errors; high status analysts still tend to over-predict.

1.5 Empirical Results

In Section 1.2 of the chapter, we demonstrated that if an analyst is compensated asymmetrically, incurring different penalties for different signs of the forecast errors, independent of his preferences towards risk, the analyst will display some degree of

Table 1.3: Forecast Errors Descriptive Statistics for Analysts in High Status Brokerage Houses

This table reports descriptive statistics of the distribution of the mean forecast errors and mean forecast errors in percentage terms across analysts who work for the top ten brokerage houses. The sample ranged from 31 March 1985 through 31 March 2005. We considered only analysts who provide more than 50 forecasts.

Forecast Errors:		
	Value	Percentage (%)
Mean	-0.0299	-2.9134
Median	-0.0057	-0.7785
Std.Dev.	0.0792	31.5578
Min	-0.5799	-142.7398
Max	0.0594	141.6635
Skewness	-3.8972	-0.0386
Kurtosis	22.5714	8.5876

asymmetry in the loss function used for forecast selection. Then, in Section 1.3 we showed how to obtain an estimate of this asymmetric pattern for a given level of risk tolerance, and, moreover, how to jointly test for rationality by evaluating the moment conditions.

1.5.1 Literature Review

Until now, the literature dedicated to the study of the rationality of financial analysts has used OLS or LAD pooled estimation. Results hinge on the particular set of instruments chosen.

With the exception of Keane and Runkle (1998), the first generation of studies, using only mean square error loss, consistently rejected the rationality of financial analysts' forecasts. De Bondt and Thaler (1990), for example, show that analysts do not to make efficient use of the information given by prior earnings. They regress actual earnings changes on forecasted earnings changes and find a coefficient that is significantly less than one. Their interpretation is that analysts are overreacting to the actual earnings change. Similarly, Abarbanell and Bernard (1992) and Easterwood and Nutt (1999) point to evidence against rationality when prior earnings changes are used as instruments. Abarbanell and Bernard (1992) use a simple pooled OLS procedure where the forecast error is regressed on past earnings changes. Easterwood and Nutt (1999) use a more general specification in which they allow for different intercepts and slopes depending on the direction of the sign and the magnitude of the prior earnings change. Elliott, Philbrick, and Weidman (1995) report that analysts do not seem to efficiently process the information embedded in past forecast revisions. They regress forecast error on forecast revision and find a significant relationship implying a rejection of rationality. Studies such as Mendehall (1991) and Ali, Klein, and Rosenfeld (1992) report that analysts do not process the informational content that is present in previous forecast errors. Klein (1990), Aitken, Frino, and Winn (1996) and Lys and Sohn (1990) indicate that analysts do not efficiently use the information present in prior stock price movements nor the information available in stock returns.

Keane and Runkle (1998) perform a pooled OLS regression allowing for a specific covariance structure that accounts for shocks that affect individual analysts and specific industries. They also argue that most of the previous results do not properly address the issue of discretionary asset write-downs, which affect earnings but are ignored by analysts when constructing earnings forecasts. Using past forecast errors as the set of instrumental variables and taking account of these special charges, they conclude that, in fact, financial analysts' forecasts appear rational.

An important idea, not further explored in Keane and Runkle (1998), is that false rejections may be due to the misrepresentation of financial analysts' preferences, or the strategic environment in which predictions are made. That is, analysts may have asymmetric loss functions or analysts may report their forecasts on the basis of predictions made by other analysts in the market. This chapter addresses the issue of an asymmetric loss function through the contractual relationship between the forecast user and the forecast provider.

Contrary to the previous literature focused on mean square error loss, Gu and Wu (2003) argue analysts minimize their absolute forecast errors and find earnings present a skewed distribution. Based on this work, Basu and Markov (2003) investigate the rationality of financial analysts by estimating a pooled LAD. They argue that financial analysts' preferences are better described using this functional form. Using a combination of instrumental variables, such as past earnings levels and changes, past forecast revisions, past forecast errors, and past stock returns, the authors conclude there is no evidence of forecast inefficiency using LAD regressions and linear loss.

1.5.2 Evidence for Asymmetric Loss Functions

We recover the asymmetry parameter using the Generalized Method of Moments (GMM) estimation technique proposed by Hansen (1982). For a consistent parameter estimate, we need a set of instruments that we believe were available to the forecaster at the time he constructed his prediction. In accordance with the literature, we use three sets of instruments: 1) two lags of forecast errors, 2) two lags of earnings changes, and 3) a combination of one lag each of the forecast error and U.S. GDP growth rates.[14] All sets of instruments also include a constant.

Table 1.4 reports results for the asymmetry parameter, for different levels of risk preferences, for the full sample. The top panel presents the GMM estimates for the first set of instrumental variables. First, we note that for the risk neutral analyst, we do not find strong evidence in favor of asymmetry. The mean of the asymmetry estimates is 0.43 and only 36.48 percent of the estimates are statistically different from 0.5. Of the estimates different from 0.5, 29.56 percent are statistically less than 0.5 (6.92 percent are statistically greater than 0.5). In fact, these estimates confirm the findings in Gu and Wu (2003) and Basu and Markov (2003), given that for 63.52 percent of the estimates, we cannot reject a loss function that is linear and symmetric (LAD).[15]

Next, as we increase the level of risk aversion, the distribution of estimates shifts in favor of a contractual parameter which implies asymmetric behavior. For instance, for analysts with a risk parameter equal to 20 (the highest level of risk aversion), 57.23 percent of estimates are statistically different than 0.5. Contrary to the risk neutral analyst, of those estimates different from 0.5, 34.59 percent are statistically greater than 0.5, suggesting a loss function which more strongly penalizes under-prediction.[16] Thus, this evidence suggests that we can only believe analysts truly benefit from the bonus in (1.2) if we accept that they are highly risk averse. Asymmetry estimates using the second and

[14]Data on U.S. GDP growth rates are from the International Monetary Fund's *International Financial Statistics*.

[15]Extreme values, such as the minimum and the maximum of the estimates, are imprecisely estimated due to large standard errors, not presented in the table.

[16]A possible explanation may be that estimation procedures for the most risk averse analysts are more sensitive to the presence of outliers.

Table 1.4: Descriptive Statistics for Asymmetry Parameter Estimates for Different Preference Parameters for the Entire Sample

This table presents some descriptive statistics of the distribution of asymmetry parameter estimates found for different analysts. Also, the last columns to the right present the percentage of estimates that are statistically different than 0.5, that are statistically greater than 0.5 and statistically smaller than 0.5, all at the 5% level. We entertain three possibilities, one in which the set of instruments is composed of a constant and two forecast error lags, a second where we use a constant plus two past earnings change and finally, we consider a constant, a lagged forecast error and previous US GDP growth rate. The sample ranged from 31 March 1985 through 31 March 2005. We considered analysts who provided more than 50 forecasts. All missing values were dropped in the estimation procedure.

Preference Parameter	**Mean**	**Median**	**Std. Dev.**	**Min**	**Max**	**Quantiles**		**(%) α not 1/2**	**(%) α Significant and**	
						25%	**75%**		$> 1/2$	$< 1/2$
Case 1: Constant Plus Two Error Lags as Instruments										
Risk neutral	0.4300	0.4246	0.1344	0.0711	0.8500	0.3422	0.5178	36.48	6.92	29.56
$\rho = \mathbf{1}$	0.4377	0.4348	0.1339	0.0713	0.8495	0.3488	0.5232	34.59	8.18	26.42
$\rho = \mathbf{5}$	0.4685	0.4686	0.1414	0.0722	0.8477	0.3745	0.5620	38.05	14.47	23.58
$\rho = \mathbf{10}$	0.4974	0.4947	0.1516	0.0733	0.8902	0.3996	0.5921	46.86	23.58	23.27
$\rho = \mathbf{20}$	0.5283	0.5302	0.1629	0.0756	0.9114	0.4093	0.6296	57.23	34.59	22.64
Case 2: Constant Plus Two Past Earnings Change as Instruments										
Risk neutral	0.4310	0.4222	0.1296	0.0974	0.7857	0.3442	0.5184	36.48	7.86	28.62
$\rho = \mathbf{1}$	0.4388	0.4379	0.1288	0.0976	0.7902	0.3543	0.5278	35.22	8.49	26.73
$\rho = \mathbf{5}$	0.4688	0.4698	0.1388	0.0984	0.8365	0.3778	0.5585	40.57	15.72	24.84
$\rho = \mathbf{10}$	0.4962	0.4973	0.1542	0.0995	0.8949	0.3938	0.5983	47.17	24.53	22.64
$\rho = \mathbf{20}$	0.5259	0.5251	0.1685	0.0829	0.9095	0.4063	0.6280	54.72	33.33	21.38
Case 3: Constant Plus Lagged Error and Previous US GDP Growth Rate										
Risk neutral	0.4321	0.4288	0.1347	0.0774	0.8557	0.3448	0.5243	37.11	7.55	29.56
$\rho = \mathbf{1}$	0.4405	0.4374	0.1338	0.0775	0.8551	0.3505	0.5307	36.48	8.81	27.67
$\rho = \mathbf{5}$	0.4735	0.4711	0.1437	0.0776	0.8529	0.3759	0.5658	38.99	15.41	23.58
$\rho = \mathbf{10}$	0.5018	0.4955	0.1557	0.0778	0.8945	0.4008	0.5993	47.80	25.16	22.64
$\rho = \mathbf{20}$	0.5314	0.5325	0.1681	0.0782	0.9090	0.4115	0.6311	57.55	34.91	22.64

Table 1.5: Descriptive Statistics for Asymmetry Parameter Estimates for Different Preference Parameters for Analysts in High Status Brokerage Houses

This table presents some descriptive statistics of the distribution of asymmetry parameter estimates found for different analysts who work for high status brokerage houses. Also, the last columns to the right present the percentage of estimates that are statistically different than 0.5, that are statistically greater than 0.5 and statistically smaller than 0.5, all at the 5% level. We entertain three possibilities, one in which the set of instruments is composed of a constant and two forecast error lags, a second where we use a constant plus two past earnings change and finally, we consider a constant, a lagged forecast error and previous US GDP growth rate.The sample ranged from 31 March 1985 through 31 March 2005. We considered analysts who provided more than 50 forecasts. All missing values were dropped in the estimation procedure.

Preference Parameter	Mean	Median	Std. Dev.	Min	Max	Quantiles		(%) α not 1/2	(%) α Significant and	
						25%	75%		$> 1/2$	$< 1/2$
Case 1: Constant Plus Two Error Lags as Instruments										
Risk neutral	0.4358	0.4359	0.1344	0.0955	0.7526	0.3464	0.5202	35.06	7.47	27.59
$\rho = \mathbf{1}$	0.4441	0.4432	0.1346	0.1022	0.7495	0.3555	0.5245	33.33	9.20	24.14
$\rho = \mathbf{5}$	0.4726	0.4734	0.1449	0.1114	0.7937	0.3836	0.5647	38.51	16.09	22.41
$\rho = \mathbf{10}$	0.4985	0.5016	0.1537	0.1130	0.8265	0.4089	0.5809	43.68	22.41	21.26
$\rho = \mathbf{20}$	0.5293	0.5307	0.1620	0.1163	0.8812	0.4243	0.6296	56.32	35.63	20.69
Case 2: Constant Plus Two Past Earnings Change as Instruments										
Risk neutral	0.4349	0.4308	0.1310	0.0974	0.6971	0.3471	0.5294	35.06	7.47	27.59
$\rho = \mathbf{1}$	0.4428	0.4426	0.1301	0.0976	0.7035	0.3579	0.5318	34.48	8.62	25.86
$\rho = \mathbf{5}$	0.4722	0.4708	0.1437	0.0984	0.7707	0.3789	0.5667	41.95	17.82	24.14
$\rho = \mathbf{10}$	0.4983	0.4999	0.1587	0.0995	0.8600	0.3991	0.5886	44.83	22.99	21.84
$\rho = \mathbf{20}$	0.5283	0.5309	0.1699	0.1015	0.9094	0.4186	0.6200	52.87	33.91	18.97
Case 3: Constant Plus Lagged Error and Previous US GDP Growth Rate										
Risk neutral	0.4374	0.4432	0.1358	0.0774	0.7451	0.3508	0.5296	34.48	7.47	27.01
$\rho = \mathbf{1}$	0.4459	0.4523	0.1353	0.0775	0.7423	0.3613	0.5328	35.63	9.77	25.86
$\rho = \mathbf{5}$	0.4778	0.4688	0.1497	0.0776	0.8511	0.3781	0.5806	40.23	17.24	22.99
$\rho = \mathbf{10}$	0.5036	0.4955	0.1611	0.0778	0.8945	0.4012	0.5898	45.98	24.14	21.84
$\rho = \mathbf{20}$	0.5321	0.5313	0.1698	0.0782	0.9022	0.4211	0.6311	56.90	36.21	20.69

third sets of instruments are presented in the lower panels of Table 1.4. The results are robust across different specifications.

Then, we restrict the full sample to only those analysts employed by high status brokerage firms, as described in Section 1.4. Table 1.5 reports the descriptive statistics for the asymmetry estimates for the three sets of instruments and for different levels of risk aversion. Similar to the full sample, risk neutral analysts at high status brokerage houses present little evidence for asymmetry. Only 35.06 percent of the estimates are statistically different from 0.5, with 27.59 percent statistically less than 0.5. Again, the greater the level of risk aversion, the more analysts display a tendency to over-predict. Though this analysis reflects the case for a constant and two lagged forecast errors as instruments (the top panel of Table 1.5), the results are similar across all sets of instruments and are presented in the lower panels of Table 1.5.

1.5.3 Rationality for Analysts with Asymmetric Loss Functions

Table 1.6 presents the percentage of cases for which rationality is rejected at the 5 percent level for the three different sets of instruments used in the asymmetry GMM estimation. The table is divided into three panels, each corresponding to a different set of instruments. In each case, we also report results for rationality tests for a symmetric square error loss (MSE) and a symmetric and linear loss (LAD).

For the estimation that included a constant and two error lags as instruments, we find, in accordance with the literature, that for the MSE loss we reject rationality for 43.71 percent of the cases. When the loss function is assumed to be LAD, we find that rationality is rejected for 27.67 percent of the analysts' forecasts. This result confirms findings in the literature, such as Gu and Wu (2003) and Basu and Markov (2003).

However, the question we seek to answer is: are financial analysts still irrational once their asymmetric incentives are considered? The moment conditions derived in Section 1.3 allow us to estimate the asymmetry parameter and, conditioned on the estimated value, we can test if the moment condition holds. In a GMM setting, we achieve this by using Hansen's test ($J - test$) of over-identifying restrictions.

First, recall from the previous section that we did, indeed, find evidence of analysts' asymmetric loss functions. Given the estimated asymmetry parameter, we present the

Table 1.6: Percentage of Cases in which Rationality is Rejected for the Entire Sample

This table presents the percentage of cases in which rationality is rejected, all at the 5% level. We entertain three possibilities, one in which the set of instruments is composed of a constant and two forecast error lags, a second where we use a constant plus two past earnings change and finally, we consider a constant, a lagged forecast error and previous US GDP growth rate. We present the results for the usual MSE and LAD cases plus the cases implied by risk neutrality and four different risk aversion parameters. The sample ranged from 31 March 1985 through 31 March 2005. We considered analysts who provided more than 50 forecasts. A ll missing values were dropped during the estimation procedure.

	Case 1: Constant and Two Error Lags as Instruments	Case 2: Constant and Two Past Earnings Change as Instruments	Case 3: Constant, Lagged Error and Previous US GDP Growth Rate
MSE	43.71	37.42	41.51
LAD	27.67	16.35	27.67
Risk Neutral	3.14	3.13	5.97
$\rho = \mathbf{1}$	3.46	2.50	5.35
$\rho = \mathbf{5}$	4.40	2.51	5.97
$\rho = \mathbf{10}$	5.56	2.83	8.18
$\rho = \mathbf{20}$	6.29	2.84	7.86

percentage of cases for which, based on the $J - test$, rationality is rejected for different risk preference parameters. For the risk neutral analyst, we can only reject rationality for 3.14 percent of the cases. Even for the most risk averse analyst, we find that rationality is rejected for only 6.29 percent of forecasts. Therefore, we conclude that we cannot reject rationality for the risk neutral or risk averse analyst.

Indeed, analysts face asymmetric incentives in the form of compensation contracts. Given these incentives, they use information efficiently when constructing their forecasts. In fact, for the highly risk averse analyst, contracts must be highly asymmetric in order not to reject rationality. On the contrary, for the risk neutral analyst, contracts may still be rather symmetric (with asymmetry parameters between 0.4 and 0.5) in order not to reject rationality. A potential problem in this setting is the issue of weak instruments. In this case, weak instruments refer to weak identification. Stock, Wright, and Yogo (2002) argue that if identification is weak, then GMM estimates may be sensitive to the addition of instruments. It is clear from the three panels in Table 1.6 that results are robust across the choice of instruments. Therefore, if we apply the counterpositive of their argument, we conclude that there is evidence of strong identification. When strong identification is present, tests for symmetry and rationality have good power (relative to the case of weak identification).

Table 1.7 presents the rationality tests for the analysts associated with the high status brokerage houses. Again, the pattern is similar to the full sample. First, once analysts are restricted to have asymmetric and quadratic loss functions, the percent of rejections of rationality drops significantly when compared to the symmetric linear case. Finally, for the more realistic loss function that accounts for asymmetric incentives, analysts appear to make use of all available information when constructing predictions. Moreover, results are robust across the different sets of instruments.

1.6 Conclusion

The main objective of this chapter is to address the issue of rationality in financial analysts' forecasts. We contribute to the extensive literature by allowing for a loss function derived from economic incentives. As a result, contrary to the literature using

Table 1.7: Percentage of Cases in which Rationality is Rejected for Analysts in High Status Brokerage Houses

This table presents the percentage of cases in which rationality is rejected for analysts who work for high status brokerage houses, all at the 5% level. We entertain three possibilities, one in which the set of instruments is composed of a constant and two forecast error lags, a second where we use a constant plus two past earnings change and finally, we consider a constant, a lagged forecast error and previous US GDP growth rate. We present the results for the usual MSE and LAD cases plus the cases implied by risk neutrality and four different risk aversion parameters. The sample ranged from 31 March 1985 through 31 March 2005. We considered analysts who provided more than 50 forecasts. All missing values were dropped during the estimation procedure.

	Case 1: Constant and Two Error Lags as Instruments	Case 2: Constant and Two Past Earnings Change as Instruments	Case 3: Constant, Lagged Error and Previous US GDP Growth Rate
MSE	40.23	33.91	36.78
LAD	22.99	10.92	22.41
Risk Neutral	1.15	1.15	4.02
$\rho = \mathbf{1}$	1.72	1.15	3.45
$\rho = \mathbf{5}$	2.87	2.30	4.02
$\rho = \mathbf{10}$	2.87	4.02	5.75
$\rho = \mathbf{20}$	4.60	2.87	5.17

symmetric and linear or quadratic loss, we fail to reject that financial analysts efficiently use all information when constructing their predictions.

Analysts' forecasts often serve as a measure of market expectations and are consistently used in investment decisions. However, financial analysts are frequently faced with incentives that prevent them from making unbiased predictions. More specifically, we propose a compensation scheme in which analysts are penalized more or less depending on the sign of their errors. This feature implies analysts use an asymmetric loss function when determining forecast selection. We prove the resulting forecasts will be biased. The fact that these forecasts display some degree of bias has long been associated with a lack of rationality in many studies. However, these studies implicitly assume a linear or quadratic symmetric loss function, which may not reflect analysts' true incentives. We propose a flexible loss that accounts for the asymmetric incentives in the forecasting business. Then, we show how to uncover the asymmetry in the loss function and jointly test for rationality using data on reported forecasts and a set of information variables that we believe were also available to the forecaster at the time the forecast was reported.

Our main results suggest that analysts do, indeed, display asymmetric loss functions. The degree of asymmetry is linked to an analysts level of risk aversion; the higher the risk averse is the analyst, the greater is his tendency to over-predict the forecast. The results are robust across the choice of instruments and for a sub-sample of analysts from high status brokerage firms. Finally, contrary to previous studies, we find strong evidence in favor of rationality for financial analysts.

Given that analysts have different preferences over the sign of the forecast error, an interesting question is how to optimally combine their predictions to arrive at a better consensus forecast that accounts for analysts' asymmetric incentives. Also, analysts consistently update their forecasts over time. Future work may consider what economic factors trigger these forecast revisions.

2

Managing Earnings Expectations: Persistence, Asymmetry and Predictability in Analysts' Earnings Forecasts

2.1 Introduction

Financial analysts issue earnings forecasts that tend to be systematically upwards biased.[1] This could be due to analysts' irrational behavior and inefficient use of information. An alternative explanation is that the bias reflects analysts' economic incentives which lead to asymmetric costs of over- and underpredictions of earnings. Analysts are rewarded in part on the basis of the precision of their forecasts which is an explicit factor in the Institutional Investor magazine's All-American rankings of analysts and also matters to investors who base their stock transactions on such forecasts.[2] However, analysts may also−implicitly or explicitly−be rewarded based on how favorable their

[1]See, e.g., Fried and Gilvoy (1982), O'Brien (1988), Butler and Lang (1991), Brous (1992), Kang, O'Brien, and Sivaramakrishnan (1994), Easterwood and Nutt (1999), Lim (2001) and Hong and Kubik (2003).

[2]Loh and Mian (2005) find that analysts with more accurate earnings forecasts also issue more profitable stock recommendations. Clement and Tse (2003) find that investors do not exploit analysts' forecasts efficiently since stock market returns correlate more strongly with the size of the analyst's brokerage house than with her past forecasting ability.

forecasts are seen to be by the firms. This matters to investment banking and trading relationships with corporate clients and may facilitate better access to firms' private information (Lim (2001)). It could also play a role in analysts' career prospects: Hong and Kubik (2003) find that analysts with more precise forecasts and (controlling for precision) more upward-biased earnings estimates are more likely to experience favorable job separations.

Such economic incentives create a trade-off and must be carefully balanced since clearly the more biased a forecast is, the less precise it can be. The nature of this trade-off and how it evolves over time is poorly understood, however. To address this point, we study in this chapter the term structure of earnings expectations which reflects how the trade-off in analysts' incentives depends on the forecast horizon. We propose a simple theoretical framework that links asymmetries in analysts' costs of over- and underpredicting earnings to the length of the forecast horizon. This framework has implications for how the bias and precision of the forecast depend on the forecast horizon.

We find that accuracy becomes increasingly important as the time to the end of the fiscal year approaches. The average forecast error goes from 12 cents per share to less than three cents per share as the forecast horizon gets reduced from 12 months to a single month. As pointed out by Hong and Kubik (2003), financial markets care about whether firms meet their earnings forecasts and so the accuracy needs to improve as the time to the actual earnings announcement date draws closer. Indeed, we find that analysts' forecasts change from being upward biased at longer forecast horizons to being slightly downward biased at shorter horizons.

If financial analysts use information efficiently and their incentives for over- and underpredicting earnings are evenly balanced, then revisions to their forecasts should not be predictable by means of information known to the analysts at the time of the original forecast. In contrast, asymmetric costs of over- and underpredictions of earnings should show up in the form of a systematic bias in the forecasts. The mere presence of a bias is therefore not well suited to test if analysts use information efficiently. In contrast, we propose a new test that studies whether the accuracy of the forecast increases as the forecast horizon is reduced. This provides a more robust test of forecast efficiency. Again our empirical results are consistent with increased precision in analysts' earnings forecasts as the end of the fiscal year is approached.

Our theoretical framework also suggests that revisions to analysts' earnings forecasts could be serially correlated and predictable. This implication seems to hold empirically. We find strong evidence of asymmetries and persistence in the magnitude and sign of revisions to analysts' earnings expectations. While it is well-known that negative revisions to earnings expectations occur more frequently than positive ones, it is less known that the probability of a negative revision goes up if the previous revision was also negative. Furthermore, the magnitude of positive and negative revisions to earnings expectations is very different, with the average negative revision roughly twice as large as the average positive revision. If, as suggested by the analysis of Easterwood and Nutt (1999), analysts underreact to bad news and overreact to good news and forecasts start out being heavily biased but get more precise over time, analysts must necessarily occasionally reduce their earnings expectations by a large amount. This is indeed what we see in our data.

We incorporate these findings in a dynamic three-state model that captures the persistence and asymmetry of positive, 'no change' and negative revisions to earnings expectations. In addition we study the effect on revisions to earnings expectations of past returns, past revisions, uncertainty about stock prices (captured by a 'realized volatility' measure) and the past T-bill rate. Three-state models that incorporate either past revisions to earnings expectations or the past T-bill rate are found to be capable of successfully predicting the direction of revisions to future earnings expectations (i.e. positive, no change or negative). In addition they generate a significant positive correlation with future revisions in out-of-sample forecasting experiments. Moreover, the ability of the three-state model to predict future revisions to earnings expectations is found to improve significantly over standard time-series forecasting models that include these variables. Finally, we find that the forecasts of the consensus revisions from the three-state model predict not only the revision to analysts' earnings estimates but also the actual earnings number announced by the firms once a year.

The chapter is structured as follows. Section 2.2 presents a simple theoretical framework to understand the behavior of analysts' earnings forecasts under asymmetric loss. Section 2.3 provides details of the data set and reports initial empirical results from the time-series of analysts' earnings forecasts. Section 2.4 introduces the three-state model used to capture the different dynamics associated with negative, 'no change' and positive revisions to earnings expectations and uses this to test for asymmetries and persistence in how analysts revise their earnings expectations. In turn this model is used in Section 2.5 to explore whether revisions to analysts earnings expectations are predictable and whether the *actual* earnings figures can be predicted by means of our forecasts of revisions in analysts' earnings expectations. Section 2.6 concludes.

2.2 Analyst Forecasts under Asymmetric Loss

In this section we explore implications of asymmetries in analysts' objectives for earnings forecasts. A large literature has documented that financial analysts have incentives to bias their earnings forecasts (see, e.g., Lim (2001) and Hong and Kubik (2003)). Brokerages employing sell side analysts may have investment banking relationships with firms whose earnings are being predicted, thus giving rise to an over-optimism bias in order to promote those firms' shares. Furthermore, analysts are likely to gain easier access to top executives if they present their firms' earnings prospects in a favorable light. Conversely, forecast accuracy has also been found to affect analysts' career prospects (Hong and Kubik (2003)).

We follow this literature by assuming that analysts' forecast objectives can be represented through a loss function, $L(e)$, that depends on the forecast error, e, and increases as the size of the forecast error ($|e|$) gets larger. The simplest way to account for asymmetries in analysts' objectives is to weight positive and negative forecast errors of equal size differently. For example, Rodriguez (2005) develops a wage contract between the forecast provider and the forecast user in which the forecaster is assumed to be penalized asymmetrically for forecast errors of different signs and shows how this compensation scheme gives rise to an asymmetric loss function for the financial analyst.

Let $f_{T,T-h}$ be the analyst's forecast of the actual earnings number, A_T, computed on the basis of information at time $T-h$, i.e. h periods prior to the earnings announcement date, T. The associated $h-$period forecast error is then given by $e_{T,T-h} = A_T - f_{T,T-h}$.

Consider the piece-wise linear (linlin) loss function:

$$L(e_{T,T-h}) = \begin{cases} (1-\theta)\left|e_{T,T-h}\right| & \text{if } e_{T,T-h} > 0 \\ \theta\left|e_{T,T-h}\right| & \text{if } e_{T,T-h} \leq 0 \end{cases}, \quad 0 < \theta < 1. \tag{2.1}$$

when $\theta = 1/2$, the cost of negative and positive forecast errors of equal size is identical. When $\theta < 1/2$, positive forecast errors (underpredictions) are more costly than negative errors (overpredictions), whereas when $\theta > 1/2$, negative forecast errors are costlier than positive ones. This loss function is convenient to work with but our results will not otherwise be dependent on this specific functional form.

Let $\mu_{e_{T,T-h}}$ and $\sigma^2_{e_{T,T-h}}$ be the conditional mean and variance of the $h-$period forecast error, $e_{T,T-h}$, given the analyst's information at time $T-h, \Omega_{T-h}$. By changing the forecast $f_{T,T-h}$ the analyst can vary the bias of the forecast error $\mu_{e_{T,T-h}} = E[A_T|\Omega_{T-h}] - f_{T,T-h}$.[3] The optimal bias minimizes the loss function (2.1) which is equivalent to solving (see Elliott and Timmermann (2004))

$$\min_{\mu_{e_{T,T-h}}} \left\{-\theta + 1_{e_{T,T-h}>0}\right\} e_{T,T-h}. \tag{2.2}$$

The expected loss to be minimized is thus

$$E\left[L(e_{T,T-h})\right] = \int_0^\infty e_{T,T-h} dF(e_{T,T-h}) - \theta\mu_{e_{T,T-h}}, \tag{2.3}$$

where F is the cumulative density function for e. Using the normalization $e_{T,T-h} = \mu_{e_{T,T-h}} + \sigma_{e_{T,T-h}} z_{T,T-h}$, we obtain the centered and standardized forecast error, $z_{T,T-h}$, whose density and cumulative density are denoted $f_z(z)$ and $F_z(z)$, respectively. z has unit variance and zero mean and it is natural to assume that $F_z(0.5) = 0$, so the effect of a non-zero mean for e is absorbed in $\mu_{e_{T,T-h}}$. The expected loss can now be expressed as (see Elliott and Timmermann (2004))

$$E\left[L(e_{T,T-h})\right] = \mu_{e_{T,T-h}}\left(1 - \theta - F_z\left(-\frac{\mu_{e_{T,T-h}}}{\sigma_{e_{T,T-h}}}\right)\right) + \sigma_{e_{T,T-h}} \int_{-\mu_e/\sigma_e}^{\infty} z_{T,T-h} dF_z(z_{T,T-h}). \tag{2.4}$$

[3]Conversely, the variance of the forecast error $e_{T,T-h}$ is determined by the variance of A_T conditional on Ω_{T-h} and hence is given.

To determine the optimal bias, $\mu^*_{e_{T,T-h}}$, we take the derivative of (2.4) with respect to $\mu_{e_{T,T-h}}$:

$$\left[1-\theta-F_z(-\mu_{e_{T,T-h}}/\sigma_{e_{T,T-h}})\right]+\frac{\mu_{e_{T,T-h}}}{\sigma_{e_{T,T-h}}}f_z(\frac{-\mu_{e_{T,T-h}}}{\sigma_{e_{T,T-h}}})-\frac{\mu_{e_{T,T-h}}}{\sigma_{e_{T,T-h}}}f_z(\frac{-\mu_{e_{T,T-h}}}{\sigma_{e_{T,T-h}}})=0.$$

Simplifying, we get

$$\mu^*_{e_{T,T-h}}=-\sigma_{e_{T,T-h}}F_z^{-1}(1-\theta), \tag{2.5}$$

where F_z^{-1} is the quantile function for z. This expression is easy to interpret. When $\theta = 0.5$, $F_z^{-1}(1-\theta)=0$, so the optimal bias is zero and the forecast is unbiased. In contrast, when $\theta < 1/2$, higher weight is placed on positive forecast errors, i.e. underpredictions of earnings. To lower the probability of large losses associated with positive forecast errors, in this situation the optimal forecast is biased in a way that makes the mean forecast error negative ($F_z^{-1}(1-\theta)>0$ and so $-\sigma F_z^{-1}(1-\theta)<0$) by centering the distribution of the forecast errors on a negative value. The opposite result holds when $\theta > 0.5$.

Figure 2.1 illustrates these the optimal bias when $\theta = 1/4$ so the forecaster prefers to overpredict earnings so as to avoid the costly positive forecast errors. Under low volatility ($\sigma_\varepsilon = 0.5$), the optimal bias is situated around -0.25. Doubling the uncertainty surrounding the earnings figure to $\sigma_\varepsilon = 1$, the bias becomes greater in magnitude (i.e. more strongly negative) and almost triples its value.

Since the analyst's information set is expanding as time progresses and the forecast horizon, h, gets reduced, it follows from the convexity of the loss function (2.1) that, on average across different time periods, $\bar{\sigma}_{e_{T,T-h_s}} < \bar{\sigma}_{e_{T,T-h_L}}$, where $h_s < h_L$ represent short and long forecast horizons, respectively. To see this, take the derivative of (2.4)

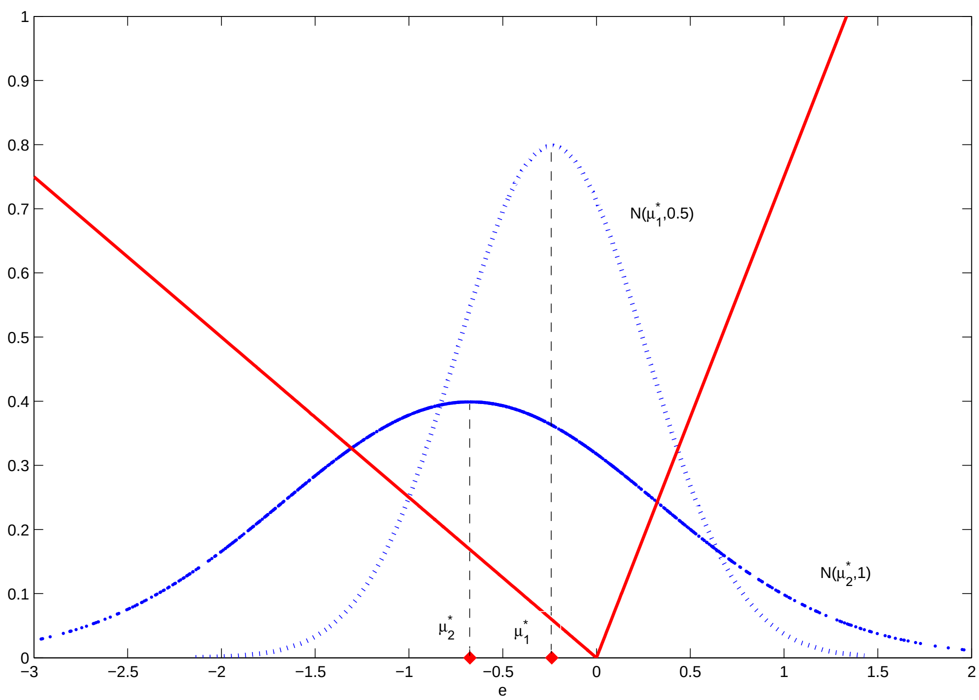

Figure 2.1: The figure reports the distribution of the forecast error, the lin-lin loss with $\theta = 0.25$, and the optimal bias $\mu_e^* = -\sigma_\varepsilon F_z^{-1}(1-\theta)$.

with respect to $\sigma_{e_{T,T-h}}$:

$$\begin{aligned}
\frac{\partial E\left[L(e_{T,T-h})\right]}{\partial \sigma_{e_{T,T-h}}} &= \frac{-\mu^2_{e_{T,T-h}}}{\sigma^2_{e_{T,T-h}}} f_z(\frac{-\mu_{e_{T,T-h}}}{\sigma_{e_{T,T-h}}}) + \int_{-\mu_e/\sigma_e}^{\infty} z_{T,T-h} dF_z(z_{T,T-h}) + \\
&\quad \frac{\mu^2_{e_{T,T-h}}}{\sigma^2_{e_{T,T-h}}} f_z(\frac{-\mu_{e_{T,T-h}}}{\sigma_{e_{T,T-h}}}) \\
&= \int_{-\mu_e/\sigma_e}^{\infty} z_{T,T-h} dF_z(z_{T,T-h}) > 0.
\end{aligned} \tag{2.6}$$

The last inequality can be seen as follows. If $\mu_{e_{T,T-h}} \leq 0$, then the integral is exclusively computed over positive values of z and hence holds automatically. On the other hand, if $\mu_{e_{T,T-h}} > 0$, then negative values of z are included in the integral. However, since $\int_{-\infty}^{\infty} z_{T,T-h} dF_z(z_{T,T-h}) = 0$, we have for this case $\int_{-\mu_e/\sigma_e}^{\infty} z_{T,T-h} dF_z(z_{T,T-h}) = -\int_{-\infty}^{-\mu_e/\sigma_e} z_{T,T-h} dF_z(z_{T,T-h}) > 0$.

It follows from this result that the bias becomes greater in times of increased uncertainty about corporate earnings. The reason for this is simple: Under asymmetric loss one side of the error distribution is particularly costly. When the dispersion of forecast errors goes up, the forecaster errs on the side of caution' by centering the forecast even further to the opposite side, i.e. by increasing the magnitude of the bias.

The finding that the bias increases as a function of the uncertainty surrounding earnings is intuitively appealing. A large bias relative to the earnings uncertainty would make it easy to detect biases in analysts' forecasts and conduct any bias-adjustments, thus mitigating its effect. Indeed our model has the desirable property that in the limit when corporate earnings uncertainty disappears, no bias remains.

Another consequence of (2.5) is that the average bias will increase as a function of the forecast horizon. This result holds quite generally independent of the specific process followed by actual earnings. Nevertheless, it is useful to study an explicit example. Suppose that actual earnings for fiscal year T, A_T, follow a random walk:

$$A_T = A_{T-1} + \varepsilon_T, \quad \varepsilon_T \sim (0, \sigma^2_\varepsilon). \tag{2.7}$$

As we are concerned with forecasts of annual earnings, we ignore seasonal components, but the model can easily be extended to account for additional components in earnings. Iterating h periods backwards on (2.7), we have

$$A_T = A_{T-h} + \varepsilon_{T-h+1} + ... + \varepsilon_T.$$

For simplicity assume that $\varepsilon_{T-h+1}, ..., \varepsilon_T$ are unpredictable given the analyst's information at time $T-h$, Ω_{T-h}, whereas A_{T-h} is known given Ω_{T-h}. Again this is easily modified and purely serves as a simplifying assumption.[4] For this example, $\sigma_{e_{T,T-h}} = \sigma_\varepsilon \sqrt{h}$ and so the optimal bias (2.5) under the random walk for earnings (2.7) becomes

$$\mu^*_{e_{T,T-h}} = -\sigma_\varepsilon \sqrt{h} F_z^{-1}(1-\theta). \tag{2.8}$$

We summarize these findings as follows.

Proposition 2.1. *Suppose that analysts' cost of overpredicting earnings exceeds their cost of underpredicting them according to the loss function (2.1) with $\theta < \frac{1}{2}$. Then*

1. *The optimal forecast is biased and the magnitude of the bias will on average be greater, the longer the forecast horizon.*

2. *The precision of the forecast (as measured by the inverse of the standard deviation of the forecast error) improves systematically as the forecast horizon is reduced;*

3. *Moreover, if the earnings follow the random walk process (2.7) then the optimal bias as well as the standard deviation of the forecast error is proportional to the square root of the forecast horizon.*

Proposition 2.1 has the important implication that the right way to test whether analysts efficiently incorporate new information into their earnings forecasts, is not through a test for unbiasedness in the forecast errors. There are good reasons to expect analysts' forecasts to be biased even when they utilize new information efficiently. Most obviously, a bias may simply reflect analysts' asymmetric costs of over- and underpredicting earnings.

[4] To the extent that ε_T is partially predictable and financial analysts receive a signal that has a correlation ρ with ε, it will be possible for them to reduce the variance of the forecast error associated with ε from σ_ε^2 to $\sigma_\varepsilon^2(1-\rho^2)$.

This does not mean that we cannot test analysts' forecast efficiency. In fact, Proposition 2.1 suggests two separate tests. One is to inspect whether the standard deviation of the forecast error gets reduced as the forecast horizon (i.e. the time to the earnings announcement) shrinks. Moreover, the bias is a function of the uncertainty in the earnings figure and so the magnitude of the bias should decline as the forecast horizon shrinks.

A somewhat surprising implication of asymmetric loss is that revisions to analysts' earnings forecasts, denoted $\Delta f_{T,h} = f_{T,T-h+1} - f_{T,T-h}$, may be serially correlated and (more generally) predictable over time whenever loss is asymmetric ($\theta \neq 1/2$). To see this note from (2.8) and (2.10) that the revision to the earnings forecast between periods $T-h$ and $T-h+1$ is given by

$$\Delta f_{T,h} = \varepsilon_{T-h+1} + (\sigma_{e_{T,T-h}} - \sigma_{e_{T,T-h+1}})F_z^{-1}(1-\theta). \tag{2.9}$$

When analysts' loss is symmetric, $\theta = 1/2$ and $F_z^{-1}(1-\theta) = 0$, so revisions will be serially uncorrelated provided that analysts use their information efficiently. Conversely, when $\theta \neq 1/2$, forecast revisions may be predictable if the uncertainty surrounding earnings forecasts, as measured by $(\sigma_{e_{T,T-h}} - \sigma_{e_{T,T-h+1}})$, changes over time in a way that is itself predictable. This could occur when the volatility of earnings news cluster through time and is consistent with a large literature on volatility clustering in many financial and macroeconomic variables.

To establish this result more formally—and to anticipate results from the empirical analysis—suppose that the forecast error is conditionally heteroskedastic with volatility driven by some underlying state variable S_t which can take two values, $s_t = 1$ or $s_t = 2$, transitions between which are governed by the transition probability matrix P with element $P[i,j] = P(s_t = i | s_{t-1} = j)$ and

$$P = \begin{pmatrix} p_{11} & 1-p_{22} \\ 1-p_{11} & p_{22} \end{pmatrix}.$$

Letting $\sigma^2_{s_T}$ be the conditional variance of ε in state s_T, the random walk process (2.7) now gets modified to

$$A_T = A_{T-1} + \sigma_{s_T} \upsilon_T, \quad \upsilon_T \sim IID(0,1). \tag{2.10}$$

Under these assumptions the conditional variance of the earnings forecast error given information at time $T-h$ is given by (assuming that $h>2$)

$$\begin{aligned}
\sigma^2_{e_{T,T-h}} &= E_{T-h}[(A_T - A_{T-h})^2] = E_{T-h}[\varepsilon^2_{T-h+1} + ... + \varepsilon^2_{T-1} + \varepsilon^2_T] \\
&= \frac{1}{2-p_{11}-p_{22}} \sum_{m=1}^{h} \pi'_{T-h} \times \\
&\quad \left[\begin{array}{cc} 1-p_{22}+\lambda_2^m(1-p_{11}) & 1-p_{22}-\lambda_2^m(1-p_{22}) \\ 1-p_{11}-\lambda_2^m(1-p_{11}) & 1-p_{11}+\lambda_2^m(1-p_{22}) \end{array}\right] \left(\begin{array}{c} \sigma_1^2 \\ \sigma_2^2 \end{array}\right),
\end{aligned}$$

where $\lambda_2 = -1 + p_{11} + p_{22}$ is an eigenvalue of P (the other one is unity) and π_{T-h} is the vector of initial state probabilities at time $T-h$. Similarly, at time $T-h+1$ we have

$$\begin{aligned}
\sigma^2_{e_{T,T-h+1}} &= E_{T-h}[\varepsilon^2_{T-h+2} + ... + \varepsilon^2_{T-1} + \varepsilon^2_T] \\
&= \frac{1}{2-p_{11}-p_{22}} \sum_{m=2}^{h} \pi'_{T-h} \times \\
&\quad \left[\begin{array}{cc} 1-p_{22}+\lambda_2^m(1-p_{11}) & 1-p_{22}-\lambda_2^m(1-p_{22}) \\ 1-p_{11}-\lambda_2^m(1-p_{11}) & 1-p_{11}+\lambda_2^m(1-p_{22}) \end{array}\right] \left(\begin{array}{c} \sigma_1^2 \\ \sigma_2^2 \end{array}\right).
\end{aligned}$$

It is clear from these expressions that $(\sigma_{e_{T,T-h}} - \sigma_{e_{T,T-h+1}})$—and hence from (2.9) the revision in analyst forecasts—is predictable since this difference depends on whether the current volatility is high, low or normal as reflected in the initial state probabilities, π_{T-h}. This is most easily seen by considering the expected value (given information at time $T-h$) of the difference $\sigma^2_{e_{T,T-h}} - \sigma^2_{e_{T,T-h+1}}$

$$\begin{aligned}
E_{T-h}[\sigma^2_{e_{T,T-h}} - \sigma^2_{e_{T,T-h+1}}] &= \frac{\pi'_{T-h}}{2-p_{11}-p_{22}} \times \\
&\quad \left[\begin{array}{cc} 1-p_{22}+\lambda_2(1-p_{11}) & 1-p_{22}-\lambda_2(1-p_{22}) \\ 1-p_{11}-\lambda_2(1-p_{11}) & 1-p_{11}+\lambda_2(1-p_{22}) \end{array}\right] \times \\
&\quad \left(\begin{array}{c} \sigma_1^2 \\ \sigma_2^2 \end{array}\right).
\end{aligned}$$

Suppose that $\sigma_2 > \sigma_1$ so uncertainty about earnings is greater in the second state than in the first state. Using an example with $p_{11} = p_{22} = 0.8$, $\sigma_2 = 2\sigma_1$, $h = 10$ and

$\theta = 1/4$, Figure 2.2 shows that as the probability of starting from the low volatility state rises from zero to one, the change in the forecast revision predicted to occur between periods $T-h$ and $T-h+1$, $E_{T-h}[\sigma_{e_{T,T-h}} - \sigma_{e_{T,T-h+1}}]F_z^{-1}(1-\theta)$, declines. Note also from the figure that $E_{T-h}[\sigma_{e_{T,T-h}} - \sigma_{e_{T,T-h+1}}] > 0$ which (as we shall see later) is consistent with the bias becoming less negative (i.e. rising) over time as the forecast horizon gets reduced.

The pattern displayed in Figure 2.2 is a general feature that holds not only for this particular example but for any earnings process with mean-reverting volatility as assumed by most models of time-varying volatility. Such mean reversion implies that if the current (conditional) volatility is low, it can be predicted to rise next period. Conversely, if current volatility is unusually high, it can be predicted to decline. In both cases the change in the volatility is partially predictable. It follows from (2.9) that such predictability translates into predictability in earnings revisions.

Furthermore, the future forecast revision $\Delta f_{T,h+1}$ depends on $(\sigma_{e_{T,T-h-1}} - \sigma_{e_{T,T-h-2}})$ and hence on ε^2_{T-h-1} and ε^2_{T-h-2}. Both $\Delta f_{T,h}$ and $\Delta f_{T,h+1}$ depend on ε^2_{T-h-1}. This induces serial dependence (persistence) in the forecast revisions and, by implication, in the forecast errors.[5] The magnitude of this depends on the persistence in the conditional variance of the earnings process. We summarize these results in the following proposition:

Proposition 2.2. *Suppose earnings follow the random walk process (2.10) and that analysts' cost of over- and underpredicting earnings are different, so $\theta \neq \frac{1}{2}$ in (2.1). Then*

1. *Revisions to analysts' earnings forecasts between subsequent periods may be predictable;*

2. *Revisions to analysts' earnings forecasts may be serially correlated.*

[5] See Patton and Timmermann (2006) for a discussion of properties of optimal forecasts under asymmetric loss.

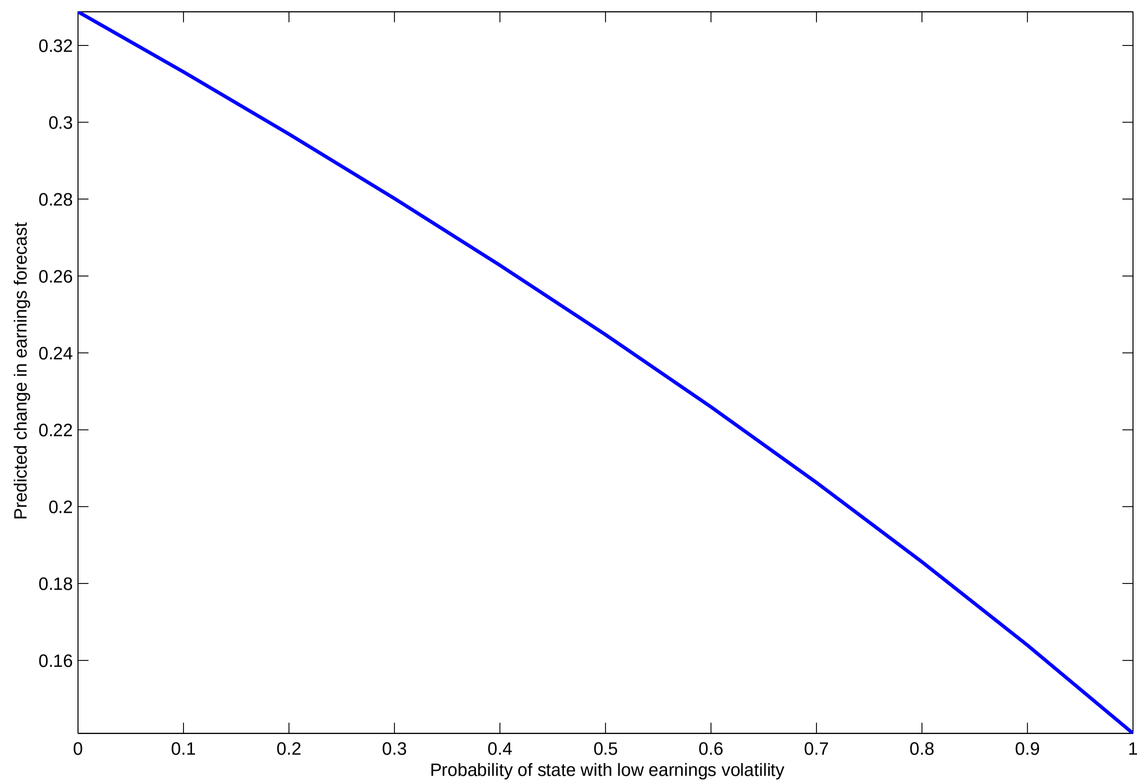

Figure 2.2: The figure reports the expected revisions in the forecasts.

We return to these properties of analysts' forecasts under asymmetric loss in the empirical analysis.

2.3 Empirical Results

How financial analysts form expectations of corporate earnings is a question that has generated significant attention in accounting and finance. This interest is motivated by the importance of earnings expectations to discounted cash flow valuation models−and hence to asset prices−and the key role analysts play in disseminating earnings information to participants in the financial markets. In particular, a number of papers have studied the link between revisions in analysts' earnings estimates and movements in stock prices. Stickel (1991) reports evidence that large earnings revisions affect stock prices, albeit with a delay as stock prices continue to drift in the direction of the revision for about six months after the revision. Park and Stice (2000) find that revisions to the top analysts' earnings estimates explain stock returns in a three day window centered on the revision date. Chen, Francis, and Jiang (2005) argue that the market's reaction to analysts' forecast revisions is consistent with investors' Bayesian learning process. Liu and Thomas (2000) find that analysts' revisions help explain how news on corporate earnings are incorporated in stock prices. Brown, Foster, and Noreen (1985) find that the sign and magnitude of analysts' forecast revisions are positively correlated with the sign and magnitude of the average cumulative abnormal security returns for the twelve-month period preceding revisions. Finally, Klein (1990) finds that firms with large (small) stock returns during the year after an earnings forecast have, on average, negative (positive) forecast errors and positive (negative) revisions one year later.[6]

While the relation between earnings revisions and stock price movements has been studied extensively, there has been little analysis of how the consensus earnings estimate evolves over longer spans of time and as a function of the forecast horizon. Our chapter addresses this issue by studying the properties of monthly revisions to analysts' earnings estimates for the 30 firms included in the Dow Jones index over a 20 year period. Our study adopts a new perspective by considering revisions to the consensus earnings

[6] For additional references to the extensive literature in this area, see the special issue of Journal of Financial and Quantitative Analysis, vol. 41, 2006, dedicated to stock analysts.

estimate every month and by studying how this is linked to the forecast horizon.

2.3.1 Data

Our data source on earnings forecasts is the summary tapes of the Institutional Brokers' Estimate System (I/B/E/S) which spans the period from January 1986 to December 2004, a total of 228 months. The advantage of using such a long time series is that it gives us the ability to better document systematic patterns in revisions to analysts' earnings expectations. We are interested in the behavior of earnings revisions at the individual firm level. To keep the analysis manageable, we focus on the firms included in the Dow Jones 30 Index listed in Table 2.1. Data coverage is excellent for each of these firms which on average is tracked by between 20 and 40 analysts. Moreover, since individual analysts usually do not provide a complete series of forecast updates, we follow Klein (1990), Chaney, Hogan, and Jeter (1999), Easterwood and Nutt (1999) and others in modeling the consensus forecasts (i.e., the cross-sectional average) in order to get the longest contiguous data series. The consensus forecast is generally viewed as highly influential and is the most widely accepted single measure of market expectations of earnings (see, e.g., Brown, Foster, and Noreen (1985)). While individual analysts' forecasts may differ from the consensus expectation, the latter still explains an overwhelming fraction of the time-series variation in individual analysts' views. Moreover, a large literature on forecast combinations suggests that it is very difficult to come up with better forecasts than those from simple averages of individual forecasters. Time-series of monthly revisions to analysts' earnings estimates are constructed as follows. Every third Thursday of the month $(t)-$the so-called statistical date$-$I/B/E/S lists all analysts' earnings estimates entered since the third Thursday of the previous month $(t-1)$. I/B/E/S then computes summary statistics (such as the consensus mean) over this set of individual analysts' estimates. We denote the consensus estimate or forecast of earnings for the current fiscal year recorded during month t for firm j by f_t^j.[7] At the end of the fiscal year, which we denote by T, firm j's actual earnings per share figure, A_T^j, is announced. Our analysis focuses on analysts' forecast of earnings for the current fiscal

[7]To indicate the fiscal year of the earnings figure, T, alternatively the more accurate notation $f_{t,T}^j$ could be used. It is implicit in this notation that $t \leq T$, i.e. the statistical date precedes the announcement date for the earnings figure in question.

Table 2.1: Analyst Coverage by Firm

This table reports company names and ticker legends for the 30 Dow Jones firms in addition to the median, minimum and maximum number of analysts covering the firms during 1986-2004.

Company Name	Ticker	Median	Min	Max
Alcoa Inc.	AA	23	14	33
American Intl. Group Inc.	AIG	28	14	35
Honeywell Intl Inc	ALD	18	5	22
American Express Inc	AXP	21	11	27
Boeing Company	BA	28	13	37
Verizon Communications	BEL	30	15	42
Caterpillar Inc.	CAT	25	14	33
Citigroup, Inc.	CCC2	18	8	33
JP Morgan & Chase & Co	CHL	24	8	30
E.I. du Pont de Nemours and Company	DD	25	10	32
The Walt Disney Company	DIS	28	17	36
General Electric Company	GE	23	12	30
General Motors Corporation	GM	24	11	30
The Home Depot, Inc.	HD	26	15	41
Hewlett-Packard Co.	HWP	30	15	45
International Business Machines	IBM	28	16	45
Intel Corporation	INTC	36	22	43
Johnson & Johnson	JNJ	31	14	38
The Coca-Cola Company	KO	25	13	31
McDonald’s Corporation	MCD	28	15	35
3m Co	MMM	20	11	28
Altria Group Inc	MO	26	6	36
Merck & Co., Inc.	MRK	40	20	51
Microsoft Corporation	MSFT	33	21	44
Pfizer Inc	PFE	38	21	47
The Procter & Gamble Company	PG	21	9	26
AT&T Inc.	SBC	29	19	38
United Technologies Corporation	UTX	24	16	31
Wal-Mart Stores, Inc.	WMT	31	20	47
Exxon Mobil Corp	XON	36	21	47

year. In particular when $t+1 \leq T$, the earnings revision for firm j, Δf^j_{t+1}, for fiscal year T is computed as the difference between the earnings estimates on the statistical dates $t+1$ and t, scaled by the initial stock price on day t. During months where the fiscal year changes from T to $T+1$ we compute the revision to the earnings forecast by comparing the previous month's forecast of earnings for fiscal year $T+1$ to the current month's forecast of this value. This allows us to create a contiguous time-series.

Following studies such as Klein (1990) and Lys and Sohn (1990), we scale the earnings revision by a firm's initial stock price, P^j_t, measured at the close on day t and obtained from the CRSP daily files. The revision to the consensus estimate of firm j's earnings figure between months t and $t+1$ is thus computed as follows:

$$\Delta f^j_{t+1} = 100 \times \left(\frac{f^j_{t+1} - f^j_t}{P^j_t} \right). \tag{2.11}$$

That is, we define the revision to the consensus earnings estimate as the change in the forecast of earnings per share from t to $t+1$, $f^j_{t+1} - f^j_t$, scaled by the initial stock price per share, P^j_t, and multiplied by 100. Revision numbers can therefore be interpreted as a percentage of the stock price.

I/B/E/S reports the consensus estimate of earnings per share rounded to the nearest cent. For many of the months included in our sample, forecast revisions are zero since the arrival of new information between two neighboring months is insufficient to lead to an earnings revision exceeding one cent. While the earnings revision is unlikely to be exactly equal to zero, we record this as a zero observation since such entries are best treated separately and represents 'small revisions'. Such observations could be discarded, but doing so may lead to important biases. A 'no change' forecast may in fact contain valuable information about future revisions, particularly if periods with small revisions tend to be persistent (which we shall subsequently see is indeed the case).

The proportion of 'no change' revisions turns out to be strongly related to the firms' quarterly earnings cycle within each fiscal year. To illustrate this, Figure 2.3 shows the proportion of zero, positive and negative earnings revisions over the course of the fiscal year as a function of the number of months remaining before the announcement date. Numbers are computed as averages across firms and fiscal years. There is a clear quarterly cycle in the proportion of zeros which is 10-15 percentage points lower three, six and nine months prior to the end of the fiscal year than during adjacent months. The corresponding increase in the proportion of positive or negative revisions corresponds to the release of the quarterly earnings numbers and hence signals the arrival of new information. Moreover, there is a clear upward drift in the proportion of zeros which nearly doubles from 35% one year prior to the end of the current fiscal year to 65% one month prior to its end. This suggests that the major part of the revisions in analysts' earnings estimates occur long before the end of the fiscal year.

2.3.2 Revisions to Analysts' Earnings Forecasts

Figure 2.4 plots the time-series of monthly revisions to analysts' earnings expectations for eight representative firms, namely Alcoa, Caterpillar, Citigroup, du Pont, General Motors, IBM, 3M and Exxon Mobil. Most revisions are quite small and lie in a range between -0.2 and 0.2 percent of the stock price. However, occasionally very large revisions occur, as in the case of GM where some revisions exceed one percent of the stock price. When this happens, it is more often than not a downward adjustment in the earnings estimate, i.e. a negative revision. The many months with small (zero) revisions are also apparent from this figure. Furthermore, revisions to earnings expectations appear to be persistent−positive revisions are more likely to follow if the previous revision was also positive. A similar finding holds for the negative revisions.

To capture such characteristics more systematically, Table 2.2 reports descriptive statistics for the time-series of revisions to earnings expectations for each of the 30 firms. Since the starting date varies across firms, the number of observations−shown in the first column−also differs. On average we have 220 monthly observations, just short of 20 years of data. This varies from a minimum of 167 observations for United Technologies to a maximum of 228 observations for a number of firms.

On average 50% of the monthly revisions to the consensus earnings forecasts are smaller than one cent per share and hence get recorded as zero. This grand average conceals substantial variations across firms, however. The proportion of 'no change' revisions to earnings estimates exceeds 80% for firms such as General Electric and Pfizer,

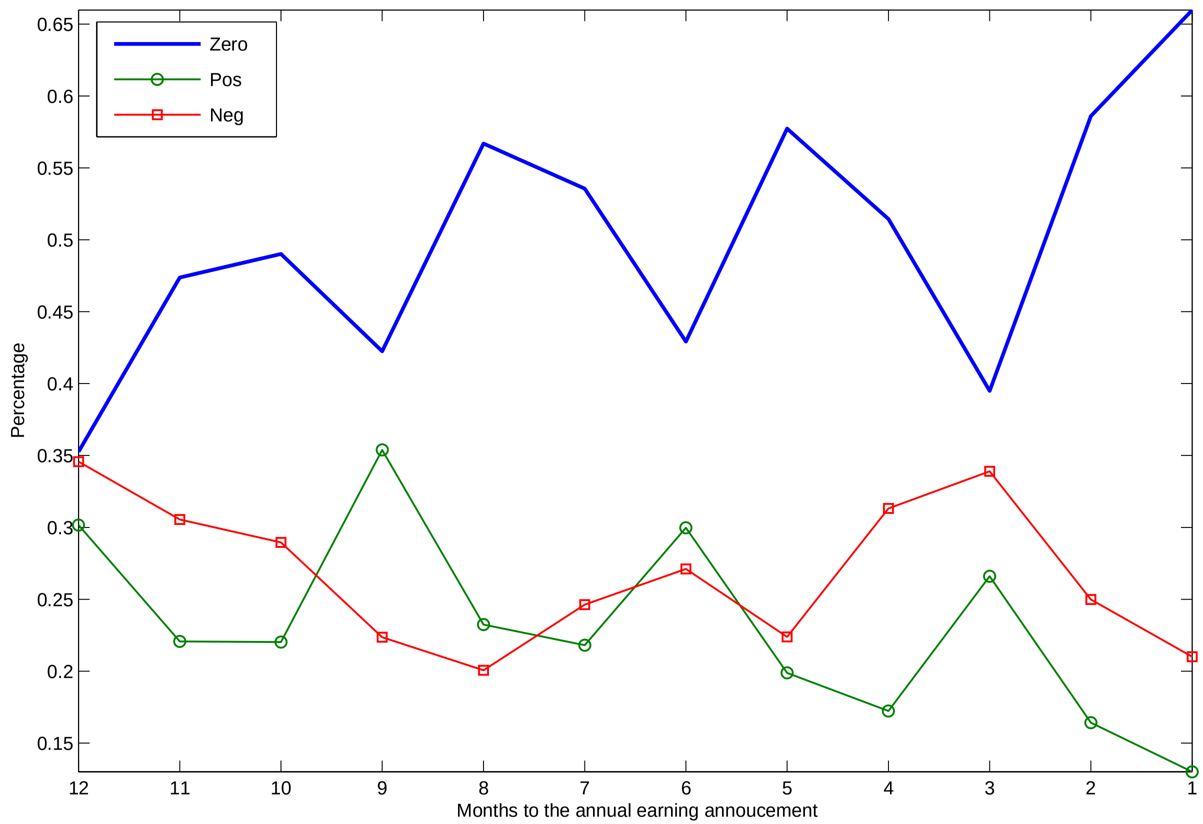

Figure 2.3: The figure reports the average-across firms- percentage of positive, zero, and negative revisions in each forecasting period as a function of the months remaining before the announcement of the annual earnings figure.

while this proportion falls below 20% for Alcoa, JP Morgan & Chase, General Motors and Exxon Mobil. This high proportion of zeros for the median forecast revision is in line with earlier results by Kang, O'Brien, and Sivaramakrishnan (1994) who find that 19-35% of their forecast revisions are zero. Similarly, for their sample of annual earnings forecasts, Brown, Foster, and Noreen (1985) find that the zeros in monthly forecast revisions range up to 80-90%. The average revision to earnings expectations, at -0.025 percent of the stock price, is negative. This pattern gives rise to a greater proportion of downward revisions (27 percent) than upward revisions (23 percent) and occurs consistently across firms. Only five of the 30 firms produce a positive average revision.

Consistent with Proposition 2.2 and earlier findings in the literature, analysts' forecast revisions are serially correlated. The average first order serial correlation coefficient is 0.235.[8] There is considerable variation in the degree of serial correlation across firms. Revisions for General Electric exhibit a very small, positive serial correlation of 0.04 (in line with Brown, Foster, and Noreen (1985) who document relatively weak serial correlation). Revisions are most persistent for Exxon Mobil, with a serial correlation coefficient of 0.68.

Revisions to earnings expectations are negatively skewed with large fat tails as revealed by the estimates of the third and fourth moments. The negative skew suggests that negative revisions are significantly larger than positive ones. Indeed, an interesting feature of the consensus revisions is that their properties differ depending on the sign of the revision. To demonstrate this, Table 2.2 reports the mean and standard deviation separately for positive and negative revisions. For 27 of 30 firms, negative revisions have a larger mean (in absolute terms) and a greater standard deviation than the positive revisions. Moreover, these differences are economically large: The average negative revision is twice as large as the average positive revision, while the average standard deviation of negative revisions is three times as large as the standard deviation of positive ones. Large negative revisions are hence far more common than large positive ones.

[8] Lys and Sohn (1990) estimate the average first order serial correlation for forecast revisions over the 58 companies in their sample to be 0.281. Abarbanell and Bernard (1992) report that the average first-order autocorrelation for their forecast errors is 0.20.

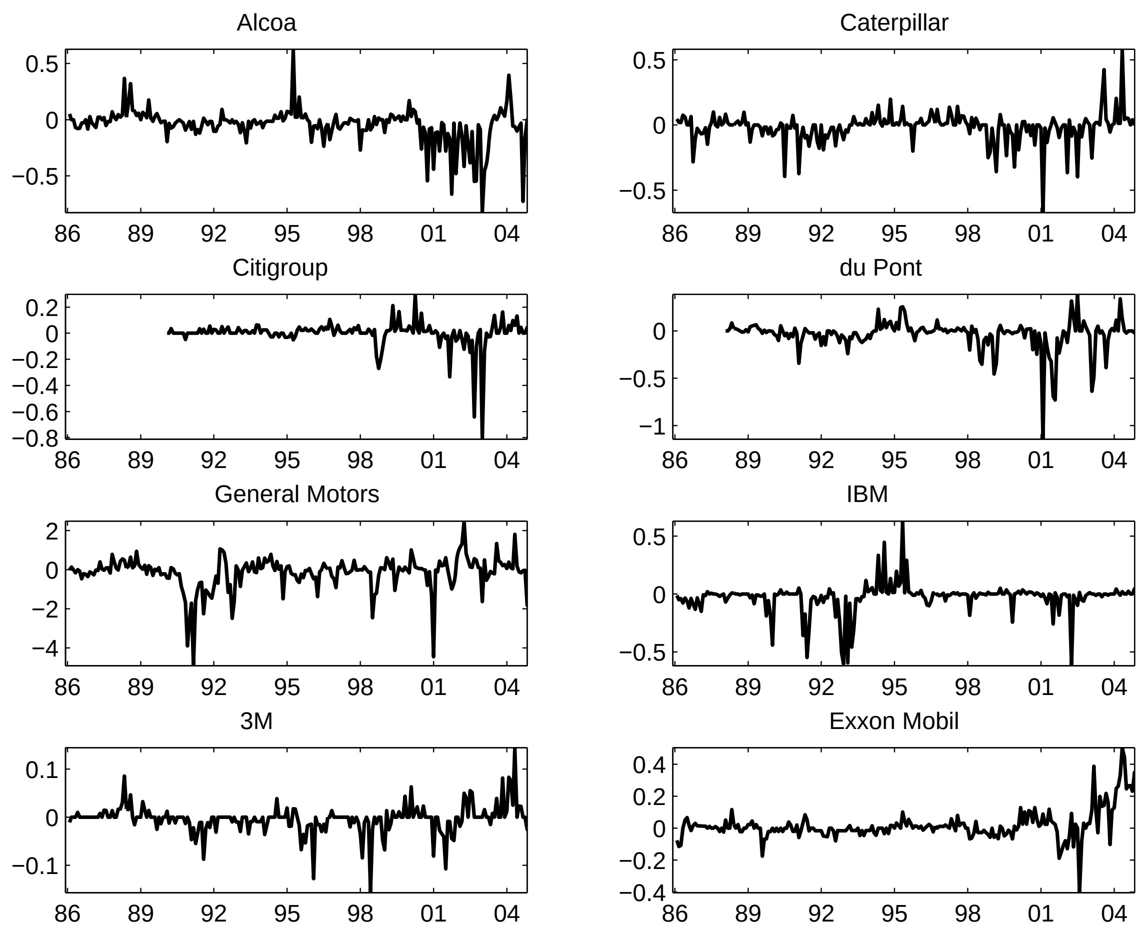

Figure 2.4: Revisions in the consensus earnings forecast between consecutive months (scaled by the initial stock price) for selected DJ30 companies over the sample Jan 1986 - Dec 2004.

Table 2.2: Data Summary

This table reports descriptive statistics for revisions in the IBES consensus earnings forecast for each of the 30 firms in the Dow Jones index over the sample Jan 1986 - Dec 2004. The revision in the consensus earnings forecast between two consecutive months is reported as a percent of the initial month's stock price and is denoted by Δf.

	# Obs	$\Delta f > 0$	$\Delta f = 0$	$\Delta f < 0$	Δf					$\Delta f > 0$		$\Delta f < 0$	
					Mean	Std	AR(1)	Skew	Kurt	Mean	Std	Mean	Std
AA	227	28.2%	11.5%	60.4%	-0.046	0.155	0.461	-1.455	10.714	0.077	0.106	-0.113	0.150
AIG	214	16.8%	69.6%	13.6%	-0.001	0.026	-0.040	-4.433	39.101	0.020	0.021	-0.036	0.052
ALD	227	18.5%	48.9%	32.6%	-0.024	0.126	0.144	-4.422	45.094	0.055	0.103	-0.104	0.180
AXP	227	20.3%	47.6%	32.2%	-0.051	0.198	0.353	-4.851	29.599	0.045	0.052	-0.187	0.304
BA	227	31.7%	27.3%	41.0%	-0.046	0.269	0.227	-5.235	47.156	0.089	0.138	-0.182	0.359
BEL	215	14.9%	59.1%	26.0%	-0.013	0.074	0.058	-7.063	71.543	0.034	0.037	-0.071	0.124
CAT	227	37.0%	26.4%	36.6%	-0.017	0.113	0.210	-0.846	12.816	0.060	0.084	-0.108	0.117
CCC2	179	45.3%	34.6%	20.1%	-0.001	0.101	0.250	-4.510	34.867	0.047	0.048	-0.109	0.170
CHL	228	35.1%	17.5%	47.4%	-0.230	1.242	0.161	-8.587	83.954	0.112	0.183	-0.569	1.740
DD	203	30.0%	21.7%	48.3%	-0.038	0.159	0.454	-2.829	17.853	0.073	0.086	-0.123	0.179
DIS	228	22.4%	47.8%	29.8%	-0.012	0.075	0.380	-1.433	12.224	0.052	0.064	-0.078	0.091
GE	227	7.9%	83.3%	8.8%	-0.003	0.042	0.039	-10.676	149.674	0.029	0.039	-0.060	0.123
GM	227	45.8%	3.5%	50.7%	-0.155	0.832	0.550	-2.122	12.149	0.361	0.397	-0.632	0.870
HD	213	19.2%	74.2%	6.6%	0.006	0.033	0.168	0.700	15.878	0.048	0.043	-0.054	0.052
HWP	228	20.6%	47.4%	32.0%	-0.028	0.174	0.172	-2.694	29.168	0.076	0.159	-0.136	0.243
IBM	227	25.6%	27.8%	46.7%	-0.029	0.128	0.364	-1.305	13.758	0.057	0.112	-0.092	0.141
INTC	227	26.9%	48.5%	24.7%	-0.015	0.132	0.172	-6.616	61.759	0.050	0.062	-0.116	0.227
JNJ	227	15.9%	76.7%	7.5%	0.001	0.013	0.175	0.457	15.351	0.021	0.014	-0.026	0.015
KO	227	13.7%	65.2%	21.1%	-0.006	0.033	0.244	-3.345	24.788	0.028	0.022	-0.044	0.049
MCD	227	11.0%	69.2%	19.8%	-0.005	0.064	0.282	-0.589	22.358	0.085	0.089	-0.072	0.088
MMM	227	19.4%	48.0%	32.6%	-0.003	0.030	0.337	-0.432	10.939	0.032	0.028	-0.029	0.028
MO	227	19.4%	58.1%	22.5%	-0.011	0.070	0.060	-7.781	74.773	0.022	0.018	-0.068	0.132
MRK	227	18.1%	70.9%	11.0%	-0.009	0.096	-0.059	-10.944	134.610	0.021	0.023	-0.112	0.271
MSFT	185	20.5%	71.4%	8.1%	0.005	0.041	0.059	-1.431	46.364	0.041	0.054	-0.048	0.088
PFE	227	9.7%	81.9%	8.4%	-0.001	0.019	0.123	-4.277	37.076	0.023	0.017	-0.037	0.043
PG	228	22.8%	60.5%	16.7%	0.000	0.016	0.064	-1.791	16.378	0.017	0.011	-0.023	0.020
SBC	215	18.6%	54.9%	26.5%	-0.015	0.098	0.080	-6.540	68.908	0.043	0.063	-0.085	0.162
UTX	167	17.4%	61.7%	21.0%	-0.010	0.042	0.569	-5.155	40.421	0.019	0.011	-0.062	0.069
WMT	226	16.8%	66.8%	16.4%	-0.002	0.022	0.325	-0.760	9.380	0.026	0.017	-0.038	0.021
XON	227	45.4%	18.5%	36.1%	0.017	0.094	0.679	1.629	11.055	0.076	0.100	-0.048	0.055
Average	220	23.2%	50.0%	26.8%	-0.025	0.151	0.235	-3.644	39.990	0.058	0.073	-0.115	0.205

2.3.3 Term Structure of Analyst Forecast Errors

To explore the implications of the theoretical analysis in Section 2.2 for how the bias and precision of analysts' forecasts evolve as a function of the forecast horizon, we next study the relationship between the forecast horizon and properties of the forecast error, $e_{T,T-h}$. For each of the 30 firms Table 2.3 shows the mean forecast error, defined as the consensus estimate of earnings for the fiscal year minus the actual earnings figure averaged across time, $\bar{A}_T^j - \bar{f}_{T,T-h}^j$, as a function of the months remaining before the announcement of the actual earnings figure $(h = 12, ..., 2, 1)$. For the vast majority of firms the forecast error starts out being negative at the 12-month horizon (corresponding to overpredictions) but rises systematically and is close to zero−corresponding to largely unbiased forecasts−at the 1-month horizon. For most firms, analysts therefore tend to systematically overpredict earnings with a bias that is greater the longer the forecast horizon.[9] Moreover, this bias is systematically reduced as the end of the current fiscal year draws closer. Few firms display the reverse pattern of initial underpredictions of earnings followed by upwards revisions to the earnings estimates.

These patterns in analyst biases are consistent with Proposition 2.1. As the time to the earnings date draws closer and uncertainty gets reduced, the bias shrinks. Moreover, at the shortest horizon, the bias has practically vanished. Indeed, for two-thirds of the firms the over-prediction bias observed at longer horizons is reversed into a slight under-prediction bias at the shortest horizon. This is consistent with an incentive for firms to meet earnings expectations. Another explanation of this finding is that, as the earnings announcement date approaches and the quality of analysts' information improves, any remaining biases in analyst forecasts become more obvious and hence more costly. We next turn our attention to how the precision of analyst forecasts changes as the time to the end of the fiscal year draws closer. To this end Figure 2.5 plots the root mean squared forecast error, calculated as an average across firms with weights that are inversely proportional to the standard deviation of the forecast errors for the individual firms.[10] At

[9]This is consistent with findings in the literature (e.g., Abarbanell (1991), Jain (1992), Kang, O'Brien, and Sivaramakrishnan (1994) and Lim (2001)) that analyst earnings forecasts on average start out too high and get reduced as the time to the earnings announcement date draws closer.

[10]A plot of the standard deviation (i.e. the square root of the mean squared forecast error minus the squared bias) is nearly identical. The decline in the root mean squared error as the horizon shrinks is therefore not induced by the smaller bias.

Table 2.3: Bias as a function of the forecasting horizon

The table reports the bias for each firm as a function of the forecast horizon (h). The bias is computed as $\frac{1}{\#t=i}\sum_{t=i} A_T^j - f_{t,T}^j$ where A_T^j is the actual earnings figure (published annually), and $f_{t,T}^j$ is the month t consensus estimate of earnings for the fiscal year, T and $i = 1, ..., 12$.

	$h=12$	$h=11$	$h=10$	$h=9$	$h=8$	$h=7$	$h=6$	$h=5$	$h=4$	$h=3$	$h=2$	$h=1$
AA	-0.189	-0.188	-0.174	-0.165	-0.152	-0.135	-0.117	-0.115	-0.088	-0.058	-0.050	-0.035
AIG	-0.045	-0.040	-0.042	-0.048	-0.048	-0.050	-0.053	-0.057	-0.042	-0.035	-0.037	-0.038
ALD	-0.127	-0.108	-0.100	-0.087	-0.083	-0.042	-0.041	-0.019	-0.013	-0.011	0.001	0.004
AXP	-0.169	-0.163	-0.132	-0.116	-0.106	-0.071	-0.075	-0.068	-0.050	-0.025	-0.012	-0.013
BA	-0.256	-0.201	-0.161	-0.153	-0.145	-0.120	-0.099	-0.104	-0.071	0.002	-0.004	0.013
BEL	-0.063	-0.061	-0.057	-0.054	-0.054	-0.048	-0.027	-0.031	-0.006	0.003	0.004	0.006
CAT	-0.032	0.006	0.011	-0.022	-0.036	-0.021	-0.033	-0.033	-0.031	-0.015	0.000	0.006
CCC2	0.001	0.022	0.022	0.007	-0.010	-0.012	-0.026	-0.012	-0.002	-0.010	-0.013	0.006
CHL	-1.183	-0.773	-0.771	-0.747	-0.686	-0.391	-0.331	-0.296	0.020	0.067	0.101	0.148
DD	-0.105	-0.082	-0.102	-0.130	-0.126	-0.104	-0.069	-0.044	-0.033	-0.013	-0.011	0.006
DIS	-0.017	-0.015	-0.019	-0.025	-0.015	-0.008	-0.006	-0.004	0.001	-0.001	0.005	0.009
GE	-0.014	-0.013	-0.012	-0.014	-0.013	-0.012	-0.012	-0.013	-0.010	-0.008	-0.000	-0.001
GM	-0.425	-0.397	-0.475	-0.617	-0.699	-0.555	-0.484	-0.390	-0.319	-0.143	-0.041	0.067
HD	0.030	0.029	0.024	0.017	0.016	0.008	0.005	0.010	0.003	0.003	0.006	0.007
HWP	-0.083	-0.063	-0.070	-0.092	-0.092	-0.073	-0.057	-0.040	-0.036	-0.005	0.001	0.008
IBM	-0.405	-0.299	-0.248	-0.198	-0.196	-0.170	-0.156	-0.156	-0.123	-0.080	-0.033	-0.006
INTC	-0.008	-0.013	0.016	0.014	0.010	0.018	0.010	0.011	0.010	0.013	0.011	0.013
JNJ	0.010	0.012	0.010	0.008	0.007	0.005	0.005	0.005	0.000	-0.001	-0.001	-0.001
KO	-0.030	-0.022	-0.015	-0.015	-0.015	-0.013	-0.015	-0.005	0.003	0.003	0.006	0.003
MCD	-0.015	-0.013	-0.010	-0.010	-0.010	-0.006	-0.007	-0.002	-0.007	-0.007	-0.003	-0.003
MMM	-0.041	-0.020	-0.023	-0.026	-0.020	-0.015	-0.016	-0.014	-0.008	-0.008	-0.003	0.002
MO	-0.053	-0.056	-0.031	-0.031	-0.032	-0.032	-0.027	-0.027	-0.013	-0.005	-0.001	-0.004
MRK	-0.043	-0.043	-0.041	-0.043	-0.043	-0.039	-0.041	-0.034	-0.006	0.003	0.003	0.002
MSFT	0.036	0.042	0.042	0.035	0.029	0.029	0.023	0.016	0.015	0.008	0.005	0.005
PFE	-0.006	-0.001	-0.002	-0.002	0.001	0.004	0.004	0.003	0.004	0.004	0.002	0.003
PG	0.003	0.001	-0.002	-0.002	-0.005	-0.010	-0.011	-0.005	-0.006	-0.008	-0.006	-0.005
SBC	-0.141	-0.126	-0.119	-0.118	-0.118	-0.115	-0.115	-0.116	-0.112	-0.097	-0.096	-0.102
UTX	-0.056	-0.042	-0.038	-0.025	-0.016	-0.012	-0.008	-0.008	-0.003	-0.002	0.000	0.002
WMT	0.001	0.008	0.008	0.005	0.004	0.005	-0.001	0.001	0.003	0.003	0.004	0.005
XON	0.161	0.150	0.146	0.115	0.092	0.079	0.080	0.071	0.065	0.063	0.063	0.046
Average	-0.109	-0.082	-0.079	-0.085	-0.085	-0.064	-0.057	-0.049	-0.028	-0.012	-0.003	0.005

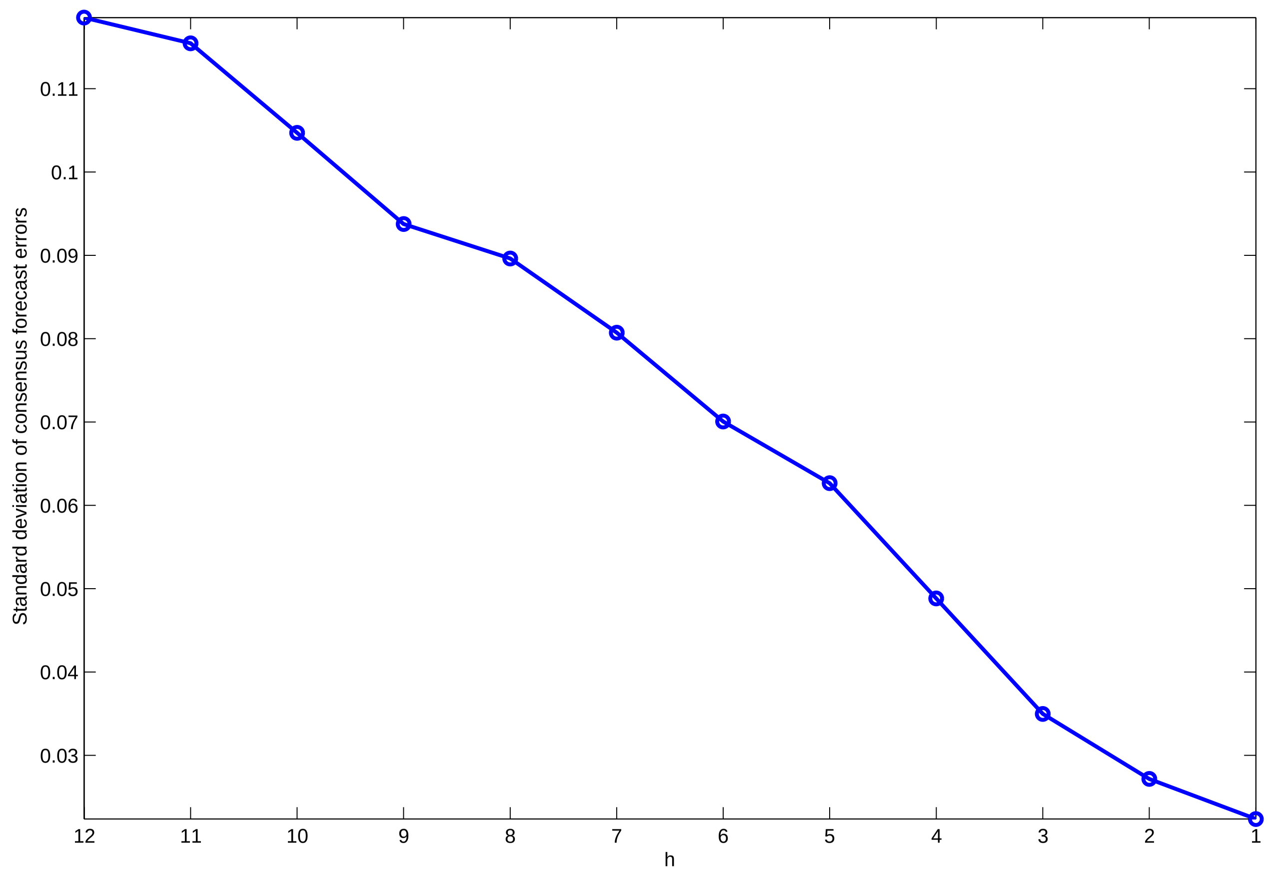

Figure 2.5: The figure reports the weighted average standard deviation of the consensus forecast errors as a function of forecast horizon (h). The weights are inversely proportional to the individual firms standard deviation of the consensus forecasts errors.

the 12-month horizon the average root mean squared forecast error is 12 cents per share. This gets steadily reduced as the fiscal year unfolds and the forecast horizon shrinks. The root mean squared forecast error is nine, seven, four and three cents per share at the nine, six, three and one month horizons, respectively. Hence, the precision of analysts' earnings estimates clearly improves systematically as the fiscal year progresses, although it never tapers off completely.

Figure 2.6 supplements Figure 2.5 by presenting root mean squared error plots for the individual firms. Patterns in the individual firms' plots are a bit noisier since they are not smoothed out by taking cross-sectional averages. Even so, for the vast majority of firms, there is a clear decline in the root mean squared forecast error as the horizon shrinks. Again these findings are consistent with Proposition 2.1 suggesting that the patterns observed in analyst forecast errors could well reflect their economic incentives rather than inefficient use of information.

2.4 Modeling Revisions to Earnings Expectations

One of the implications of the theoretical analysis in Section 2.2 is that we should expect revisions to analysts' earnings forecasts to be predictable whenever the uncertainty surroundings future earnings is itself persistent.[11] In order to model predictability in analysts' earnings forecast we need to account for the different properties and dynamic behavior of negative and positive revisions to earnings expectations and to capture information revealed in draws from the 'no change' state (rather than discarding information on this state). To this end we propose a dynamic three-state model whose state variable tracks revisions to analyst earnings forecasts:[12]

$$s_{t+1} = \begin{cases} +1 & \text{if } \Delta f_{t+1} > 0 \\ 0 & \text{if } \Delta f_{t+1} = 0 \\ -1 & \text{if } \Delta f_{t+1} < 0 \end{cases} . \tag{2.12}$$

[11] This condition seems to be satisfied empirically. For around half of the Dow Jones firms there is evidence of ARCH effects in the quarterly earnings figures.

[12] For simplicity, and without risk of confusion, we have omitted the firm superscript, j. We follow this convention in the subsequent analysis.

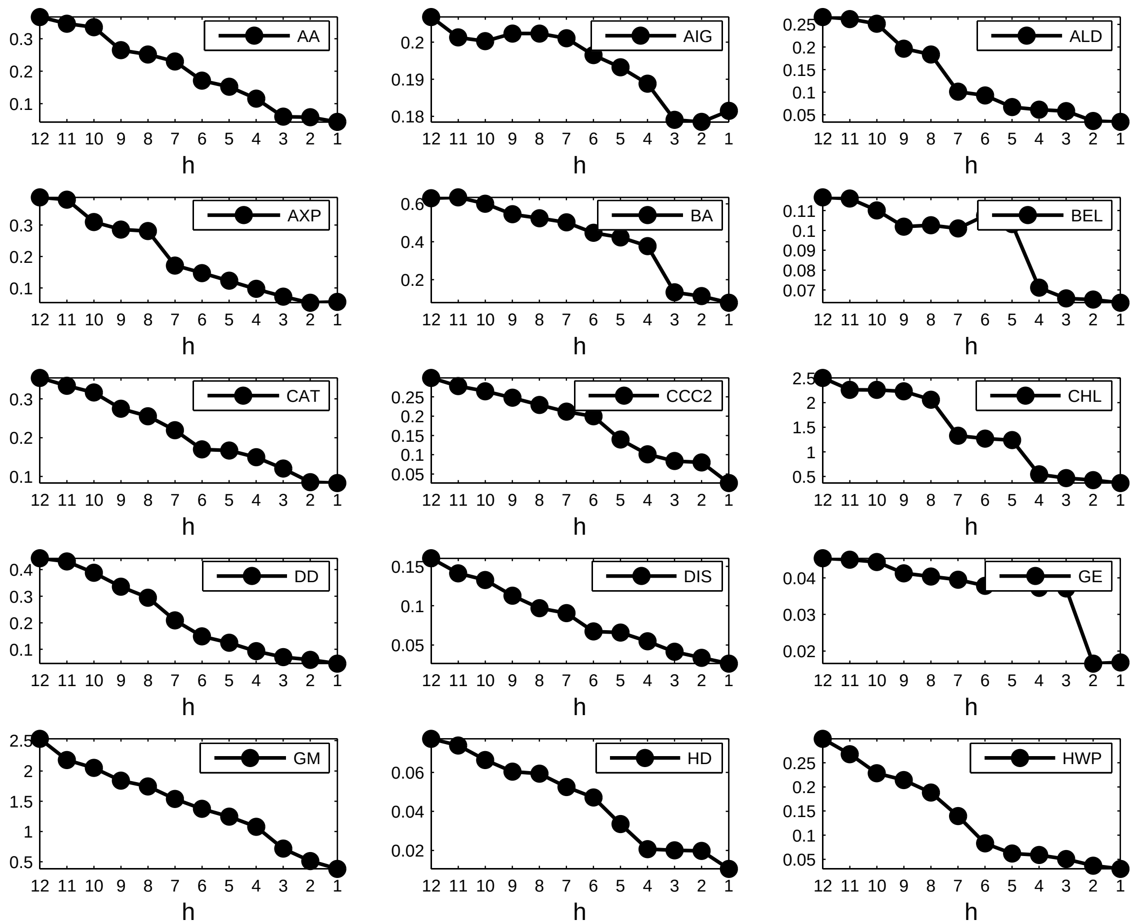

Figure 2.6: This figure reports the standard deviation of the consensus forecast errors as a function of the forecast horizon (h) for the individual firms in the Dow Jones Index.

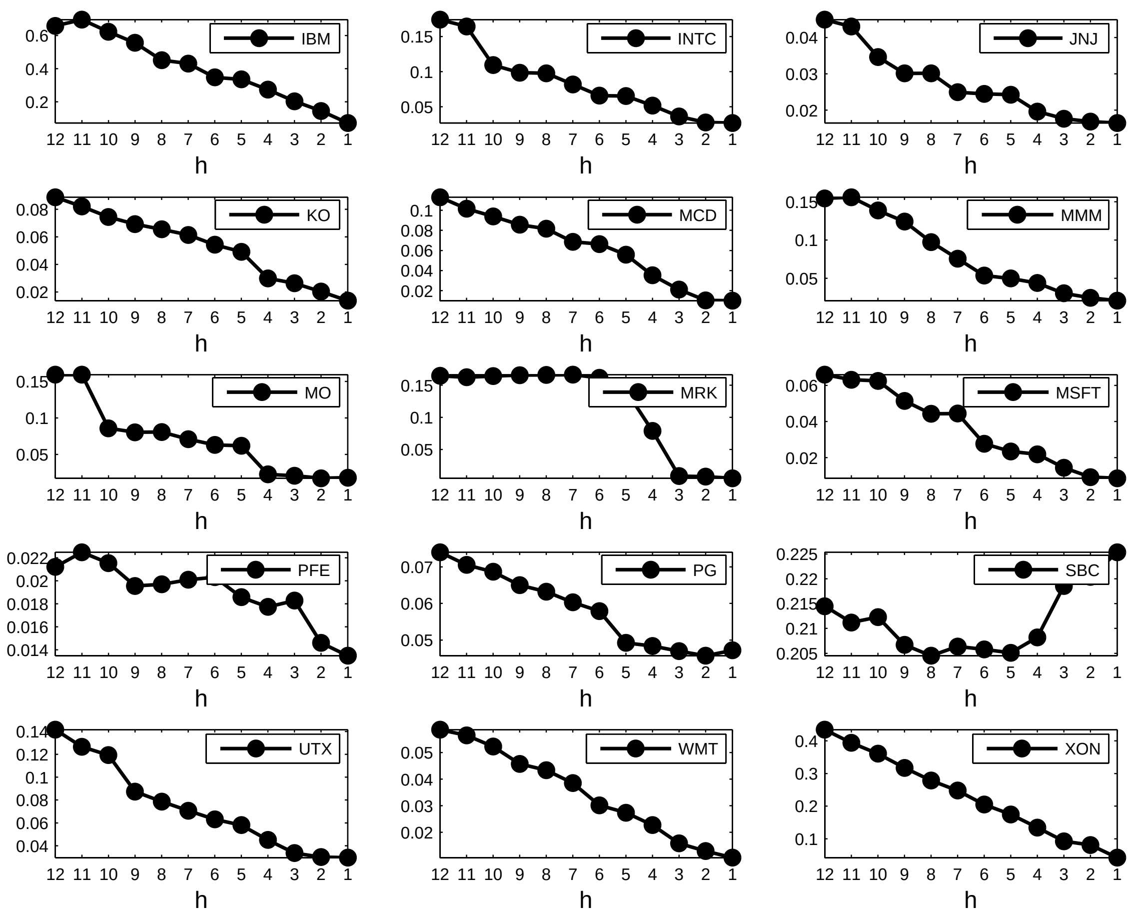

Figure 2.6 (continued)

If $s_{t+1} = 1$, the revision to analyst earnings expectations at time $t+1$ is positive, while conditional on $s_{t+1} = -1$, it is negative. Finally, if $s_{t+1} = 0$, the revision to the earnings expectation is zero (i.e. less than one cent per share). Treating small changes (zeros) separately is important if there are months where essentially no news arrives. A model that pools 'no news' and news months is likely to be misspecified as it cannot fully capture the dynamics of earnings expectations in months with news which may be quite different from the dynamics in months with no news.

Table 2.2 showed that revisions to earnings expectations are not normally distributed. In particular, there are far more large negative revisions than we should expect under a normal distribution. To capture this aspect of the data, we assume that both positive and negative revisions to earnings expectations$-$the latter with their sign reversed so they become positive$-$are log-normally distributed. The log-normal distribution accommodates fatter tails than the normal distribution and is thus better suited for our data which clearly is affected by outliers.

Table 2.2 also indicated that the mean and variance of positive and negative revisions to earnings expectations are very different. Such differences in the behavior of positive and negative revisions to the consensus estimate could reflect asymmetric incentives encountered by analysts so we do not want to impose that the parameters for positive and negative revisions are identical. In this spirit we assume that positive revisions to earnings expectations are log-normally distributed with mean μ_1 and variance σ_1^2, while negative revisions (with their sign reversed) are log-normally distributed with mean μ_{-1} and variance σ_{-1}^2.The distribution of Δf_{t+1}, $g(\Delta f_{t+1})$ conditional on s_{t+1} is thus described by the following mixture of log-normals:

$$\begin{aligned} g(\Delta f_{t+1}|s_{t+1} = 1) &= \frac{1}{\Delta f_{t+1}\sqrt{2\pi\sigma_1^2}} \exp\left(\frac{-(\ln(\Delta f_{t+1}) - \mu_1)^2}{2\sigma_1^2}\right) \\ g(\Delta f_{t+1}|s_{t+1} = 0) &= \mathbf{1}_{\Delta f_{t+1}=0} \\ g(-\Delta f_{t+1}|s_{t+1} = -1) &= \frac{1}{|\Delta f_{t+1}|\sqrt{2\pi\sigma_{-1}^2}} \exp\left(\frac{-(\ln(|\Delta f_{t+1}|) - \mu_{-1})^2}{2\sigma_{-1}^2}\right). \end{aligned} \tag{2.13}$$

Here $\mathbf{1}_{\Delta f_{t+1}=0}$ is an indicator function that equals one if $\Delta f_{t+1} = 0$ and otherwise is zero. Notice that we use the absolute value, $|\Delta f_{t+1}|$, when $\Delta f_{t+1} < 0$ so the logarithm of the revision to earnings expectations is well-defined.

To complete the model, we need to characterize the factors that determine transitions between the three states, representing positive, zero and negative revisions to earnings expectations. We consider two models. The first assumes that state transition probabilities $p_{i,k} = P(s_{t+1} = k|s_t = i)$ are constant:

$$\mathbf{P} = \begin{pmatrix} p_{1,1} & 1 - p_{1,-1} - p_{1,1} & p_{1,-1} \\ 0.5(1 - p_{0,0}) & p_{0,0} & 0.5(1 - p_{0,0}) \\ p_{-1,1} & 1 - p_{-1,1} - p_{-1,-1} & p_{-1,-1} \end{pmatrix}. \qquad (2.14)$$

State transitions are parameterized so there are only five free parameters, namely $p_{1,1}$ (the probability of continued positive revisions), $p_{0,0}$ (the probability of continued zero revisions), $p_{-1,-1}$ (the probability of continued negative revisions), $p_{1,-1}$ (the probability of switching from a positive to a negative revision) and $p_{-1,1}$ (the probability of switching from a negative to a positive revision). Since probabilities must add to one across the columns, the only constraint is that the probabilities of moving from the 'zero' state to the negative and positive earnings revision states are identical. This constraint is supported by our data and reduces the number of parameters to be estimated.

This model captures the possibility that positive revisions to earnings expectations are more likely to occur if the previous revision was also positive than if it was negative, i.e. $p_{1,1} > p_{-1,1}$, as Figure 2.4 suggests holds for our data. Similarly, negative revisions to earnings expectations are more likely to follow from past negative revisions than from past positive revisions whenever $p_{-1,-1} > p_{1,-1}$. The model is also consistent with persistence in the mean of the revisions$-$this occurs when $\mu_1 \neq \mu_{-1}$ and $p_{1,1} \neq p_{-1,1}$ or $p_{-1,-1} \neq p_{1,-1}$$-$or in their volatility (assuming $\sigma_1^2 \neq \sigma_{-1}^2$). One of the advantages of treating small ('zero') changes separately is that it can accommodate the clustering of periods with low (high) volatility in revisions. Finally, the model can capture skew and kurtosis in revisions to earnings expectations, properties that the estimates in Table 2.2 show are present in the data.

Our second model lets the diagonal elements of the state transitions (the 'stayer' probabilities) vary as a function of a set of state or predictor variables, x_{t-1}, which along with data on revisions to earnings expectations are collected in the information set $\Omega_{t-1} = \{\Delta f_{t-1}, x_{t-1}\}$. Transitions between states s_{t-1} and s_t are thus allowed to depend on x_{t-1}: $p_{i,k,t-1} = P(s_t = k|s_{t-1} = i, x_{t-1})$. More specifically, we use a logit specification,

$$p_{i,k,t-1} \equiv \frac{\exp(\delta_{1,i,k} + \delta_{2,i,k}x_{t-1})}{1 + \exp(\delta_{1,i,k} + \delta_{2,i,k}x_{t-1})}. \tag{2.15}$$

This gives rise to a time-varying transition probability matrix

$$\mathbf{P}_{t-1} = \begin{pmatrix} p_{1,1,t-1} & 1 - p_{1,-1,t-1} - p_{1,1,t-1} & p_{1,-1,t-1} \\ 0.5(1 - p_{0,0,t-1}) & p_{0,0,t-1} & 0.5(1 - p_{0,0,t-1}) \\ p_{-1,1,t-1} & 1 - p_{-1,1,t-1} - p_{-1,-1,t-1} & p_{-1,-1,t-1} \end{pmatrix}. \tag{2.16}$$

2.4.1 Estimation

Since all variables, including the underlying state, are observed, we can estimate our model by maximum likelihood methods. The chief complication is the highly non-linear functional form, which means that numerical optimization methods have to be used.

At time t, the log-likelihood function conditional on the information set Ω_{t-1} (which includes knowledge of s_{t-1}), takes the form of a mixture distribution:

$$\begin{aligned} LL_{t|\Omega_{t-1}} &= ln[p_{s_{t-1},1} \times g\left(\Delta f_t | s_t = 1, \Omega_{t-1}\right) + p_{s_{t-1},0} \times \mathbf{1}_{\Delta f_t = 0} + \\ &\quad p_{s_{t-1},-1} \times g\left(-\Delta f_t | s_t = -1, \Omega_{t-1}\right)], \end{aligned} \tag{2.17}$$

where the densities $g(\cdot)$ are given by (2.13). Summing across all time periods, $t = 1, .., T$ and integrating across the three values that s_{t-1} can take $\{-1, 0, 1\}$, we get the full sample log-likelihood function:

$$LL = \sum_{t=1}^{T} \sum_{s_{t-1}=-1}^{1} \mathbf{1}_{s_{t-1}} LL_{t|\Omega_{t-1}}. \tag{2.18}$$

Estimation proceeds by choosing parameters $\{\mu_1,\ \mu_{-1},\ \sigma_1^2,\ \sigma_{-1}^2,\ p_{1,1},\ p_{0,0},\ p_{-1,-1},\ p_{1,-1}, p_{-1,1}\}$ to maximize the log-likelihood function LL in (2.18), using Quasi-Newton methods with a mixed quadratic and cubic line search procedure.

2.4.2 Persistence and Asymmetries in Revisions to Earnings Expectations

Table 2.4 reports parameter estimates for the simplest sign mixture model (2.13) with constant transition probabilities (2.14).[13] First consider the mean and volatility parameters in the positive and negative revision states, respectively. The first column of Table 2.4 reports the ratio of the mean parameter in the negative revision state to its value in the positive revision state (exponentiated since we use a log-normal distribution). Consistent with the evidence in Table 2.2, most of the ratios reveal that the magnitude of negative revisions to earnings expectations are 30% to 170% higher than their positive counterparts. Similarly, the ratios of the volatility estimates in the negative and positive revision states, shown in column 2, exceed unity for all but two firms and are significantly above unity for around half of the firms. This corresponds to significantly higher volatility of negative than of positive revisions to analysts' earnings expectations.

Turning to the persistence in revisions to earnings expectations, far from following a random process, the sign of revisions to earnings expectations is highly persistent for many of the firms. This is clear from columns 3-7 in Table 2.4 which show parameter estimates and standard errors for the five transition probability parameters. For example, for Alcoa (AA) the probability that a positive revision follows a previous positive revision is 66% while the probability that a negative revision follows a previous negative revision is 82%. Whenever the probability of observing a positive revision is greater if the previous revision was positive ($p_{1,1}$) than if it was negative ($p_{-1,1}$), positive revisions to earnings expectations are persistent. This condition holds empirically for 24 of the 28 firms for which a sufficient number of sign changes was observed. A similar argument for persistence among negative revisions to earnings expectations can be made when $p_{-1,-1}$ exceeds $p_{1,-1}$. This holds empirically for every single firm for which estimates of both parameters are available.

To investigate the presence of persistence in revisions to earnings expectations more formally, we adopted a likelihood ratio test of the null hypothesis of no sign persistence ($p_{1,1} = p_{1,-1}$ and $p_{-1,1} = p_{-1,-1}$). Column eight in Table 2.4 shows that this hypothesis

[13] Estimates were found to be robust to different starting values. Standard errors are obtained using the delta method.

Table 2.4: Parameter Estimates for the Sign Mixture Model with Constant Transition Probabilities

For each of the DJ30 firms a sign mixture model was estimated to revisions in the consensus earnings forecast between consecutive months, Δf_t: $log(|\Delta f_t|) = \beta_{1,s_t} + \epsilon_t, \quad \epsilon_t \sim N\left(0, \sigma^2_{s_t}\right)$ with state transition probabilities $p_{i,k} = P\left(s_t = i | s_{t-1} = k\right)$ and $i, k = 1, 0, -1$. States 1,0,-1 capture positive, zero and negative earnings revisions, respectively. Standard errors appear in parentheses to the right of the parameter estimates. The first and second columns report the ratio of the means and volatilities in the negative versus the positive states. The column *No Persist.* reports the Wald test for the joint restrictions $p_{1,1} = p_{-1,1}$ and $p_{1,-1} = p_{-1,-1}$ and is a test of no persistence in the signs of the earnings revisions. (-) indicates too few transitions between states to allow estimation of the parameters. ** and * denote significance at 5% and 10% level respectively. $\overline{p}_{i,k}$ are the steady state probabilities.

	$\frac{exp(\beta_{1,-1})}{exp(\beta_{1,1})}$	$\frac{\sigma_{-1}}{\sigma_1}$	$p_{1,1}$	$p_{1,-1}$	$p_{0,0}$	$p_{-1,1}$	$p_{-1,-1}$	*No Persist.*	$\overline{p}_{1,1}$	$\overline{p}_{0,0}$	$\overline{p}_{-1,-1}$
AA	1.35**	1.13	0.66 (0.06)	0.17 (0.05)	0.15 (0.07)	0.10 (0.03)	0.82 (0.03)	60.95**	0.32	0.12	0.57
AIG	1.41*	1.27	0.11 (0.05)	0.11 (0.05)	0.72 (0.04)	0.14 (0.06)	0.41 (0.09)	7.96**	0.13	0.68	0.19
ALD	1.67**	1.22	0.19 (0.06)	0.29 (0.07)	0.60 (0.05)	0.18 (0.04)	0.51 (0.06)	6.18**	0.19	0.49	0.32
AXP	2.70**	1.53**	0.35 (0.07)	0.20 (0.06)	0.66 (0.05)	0.15 (0.04)	0.64 (0.06)	20.42**	0.20	0.47	0.33
BA	1.60**	1.19	0.44 (0.06)	0.18 (0.05)	0.31 (0.06)	0.18 (0.04)	0.65 (0.05)	31.08**	0.30	0.27	0.42
BEL	1.56**	1.39**	0.19 (0.07)	0.16 (0.07)	0.68 (0.04)	0.16 (0.05)	0.48 (0.07)	8.45**	0.17	0.60	0.24
CAT	1.83**	1.06	0.61 (0.05)	0.12 (0.04)	0.30 (0.06)	0.12 (0.04)	0.64 (0.05)	45.62**	0.36	0.27	0.38
CCC2	1.59**	1.54**	0.62 (0.05)	0.09 (0.03)	0.47 (0.06)	0.20 (0.07)	0.57 (0.08)	26.60**	0.38	0.33	0.28
CHL	2.43**	1.39**	0.60 (0.05)	0.26 (0.05)	0.26 (0.07)	0.19 (0.04)	0.65 (0.05)	32.78**	0.37	0.17	0.46
DD	1.41**	1.17	0.54 (0.06)	0.18 (0.05)	0.21 (0.06)	0.10 (0.03)	0.72 (0.05)	41.18**	0.29	0.21	0.50
DIS	1.39*	1.17	0.30 (0.06)	0.16 (0.05)	0.58 (0.05)	0.10 (0.04)	0.65 (0.06)	24.23**	0.19	0.46	0.35
GE	1.48	1.48*	— (–)	0.28 (0.53)	0.87 (0.21)	— (–)	— (–)	92.01**	0.06	0.87	0.07
GM	1.24	1.37**	0.70 (0.43)	0.24 (0.46)	— (–)	0.25 (0.73)	0.73 (0.71)	47.03**	0.48	0.04	0.49
HD	1.08	1.01	0.22 (0.38)	— (–)	0.75 (0.18)	— (–)	0.29 (0.59)	232.37**	0.12	0.75	0.13
HWP	1.65**	1.08	0.26 (0.06)	0.09 (0.04)	0.44 (0.05)	0.10 (0.03)	0.52 (0.06)	18.97**	0.22	0.47	0.31
IBM	1.65**	1.15	0.43 (0.07)	0.21 (0.05)	0.42 (0.06)	0.13 (0.03)	0.72 (0.04)	34.38**	0.25	0.28	0.47
INTC	1.60**	1.31**	0.42 (0.06)	0.15 (0.05)	0.58 (0.05)	0.11 (0.04)	0.55 (0.07)	21.45**	0.23	0.48	0.30
JNJ	1.25	1.06	0.17 (0.45)	— (–)	0.77 (0.18)	— (–)	0.35 (0.51)	209.75**	0.11	0.76	0.14
KO	1.31**	1.67**	0.13 (0.06)	0.19 (0.07)	0.74 (0.04)	0.13 (0.05)	0.52 (0.07)	8.58**	0.13	0.65	0.23
MCD	0.78	1.06	0.54 (0.41)	— (–)	0.76 (0.19)	— (–)	0.31 (0.32)	488.32**	0.18	0.69	0.12
MMM	0.91	1.02	0.39 (0.31)	— (–)	0.57 (0.19)	0.07 (0.49)	0.62 (0.26)	773.12**	0.21	0.51	0.29
MO	1.85**	1.66**	0.18 (0.06)	0.14 (0.05)	0.56 (0.04)	0.08 (0.04)	0.37 (0.07)	7.14**	0.18	0.58	0.24
MRK	2.22**	1.87**	0.15 (0.06)	0.10 (0.05)	0.71 (0.04)	0.20 (0.08)	0.16 (0.07)	0.93	0.15	0.70	0.14
MSFT	1.03	0.95	0.14 (0.48)	— (–)	0.70 (0.19)	0.20 (0.69)	0.33 (0.59)	96.51**	0.15	0.69	0.15
PFE	1.21	1.49*	— (–)	— (–)	0.83 (0.19)	— (–)	— (–)	NaN	0.07	0.85	0.07
PG	1.23*	1.29*	0.29 (0.06)	0.16 (0.05)	0.68 (0.04)	0.21 (0.07)	0.37 (0.08)	5.01*	0.20	0.60	0.20
SBC	1.52**	1.35**	0.17 (0.06)	0.15 (0.06)	0.55 (0.05)	0.14 (0.05)	0.41 (0.07)	7.12**	0.19	0.55	0.26
UTX	2.68**	1.61**	0.14 (0.06)	0.10 (0.06)	0.66 (0.05)	0.06 (0.04)	0.54 (0.08)	11.22**	0.14	0.60	0.26
WMT	1.45**	1.04	0.39 (0.33)	— (–)	0.75 (0.19)	0.16 (0.47)	0.35 (0.36)	273.02**	0.18	0.69	0.13
XON	0.83	0.72**	0.68 (0.05)	0.16 (0.04)	0.24 (0.07)	0.21 (0.04)	0.61 (0.05)	41.45**	0.45	0.19	0.36

is soundly rejected for all but one of the firms (Microsoft) in our sample. Interestingly, for 24 of the 30 firms in Table 2.4, the probability of continued negative revisions to the earnings estimates exceeds that of continued positive revisions. Although negative revisions to earnings expectations tend to be larger than positive ones, 'bad news' seems to travel slower than good news in the sense that it takes longer for it to get fully incorporated into analysts' earnings expectations.

The finding that the current earnings state predicts future revisions to analysts' expectations is consistent with our model in Section 2.2. This follows since volatility is very different in the positive and negative earnings states. Persistence in the underlying state means that earnings volatility is also predictable (persistent) and so, as stated in Proposition 2.2, revisions to earnings forecasts should also themselves be predictable.

The last three columns in Table 2.4 report the steady-state or average probabilities of positive, negative and zero revisions. These estimates can be compared to the observed frequencies reported in columns 2-4 in Table 2.2, all of which are closely matched by the model. For example, for Alcoa and AIG, the probability of a zero state is 11% and 69% compared with probabilities of 12% and 68% implied by our estimates.

Comparing our estimates of $p_{0,0}$ to the steady-state probability $\bar{p}_0$, it is clear that the probability of observing a small (less than one cent per share) change in earnings expectations is typically raised by 5-10% if the previous change was small. This suggests some degree of volatility clustering (with clusters of small revisions) in the revision process for analysts' earnings expectations.

2.4.3 Sources of Persistence in Revisions to Earnings Expectations

We conclude from Table 2.4 that revisions to analysts' earnings expectations are highly persistent with an asymmetric distribution for positive and negative values. To better understand what gives rise to such effects it is interesting to link revisions in earnings expectations to economic state variables. To this end we consider several variables, some of which have previously been studied in the accounting and finance literature, while others are new.

As our first measure we consider past returns, a variable that has been considered in many previous studies of analyst expectations, i.e., Brown, Foster, and Noreen (1985), Klein (1990), Stickel (1991), Liu and Thomas (2000) and Park and Stice (2000). Asymmetric information models such as Diamond and Verrecchia (1981) and Glosten and Milgrom (1985) imply that price changes may contain information about fundamentals and so forecast revisions could be positively correlated with prior price changes when private information gradually gets incorporated into stock prices. Comparing the proportion of positive revisions given past positive returns to the proportion of positive revisions given past negative returns, Abarbanell (1991) finds a positive correlation between price changes and forecast revisions.[14]

Building on these results, we compute the lagged cumulated return using daily data on stock returns obtained from CRSP. More precisely, let P_τ^j be the closing stock price for firm j on day τ, so the continuously compounded single-day stock return is given by $r_\tau^j = log(P_\tau^j / P_{\tau-1}^j)$. The cumulated return between time t and $t+1$ (typically one month apart) is then given by

$$r_{t:t+1}^j = \sum_{t<\tau\leq t+1} r_\tau^j = log(P_{t+1}^j) - log(P_t^j). \tag{2.19}$$

To see whether revisions to analysts' earnings expectations adjust gradually to news, we also use past revisions as a predictor variable for future revisions. If revisions to earnings expectations between two adjacent periods occur gradually, we should expect to find that past revisions predict future values.

Another strand of the literature studies the effect that uncertainty has on earnings forecasts. Indeed, Proposition 2.2 suggests that variables tracking uncertainty in the earnings process should help predict revisions to analysts' earnings expectations. In an experimental setting, Miller and Sedor (2006) vary the levels of uncertainty about future earnings and report that revisions to analysts' earnings forecasts are influenced by observed price changes only when uncertainty about future earnings is high. When uncertainty about future earnings is low, price changes are not reflected in forecast revi-

[14] Cooper, Day, and Lewis (2001) rank individual analysts' performance based on the time of the earnings forecast, the abnormal trading volume associated with the forecast, and forecast accuracy. Distinguishing between high-tech and low-tech firms, they find that, independent of analyst performance, revisions of high-tech firms are significantly related to past excess returns. However, for low-tech firms, only low-performing analysts tend to revise expectations based on information contained in past excess returns.

sions.

Measuring such uncertainty directly is complicated since earnings data is not available at high frequency. To the extent that greater earnings uncertainty is reflected in more volatile stock returns, it is reasonable to consider a measure of return volatility. We adopt a measure that builds on the literature on realized volatility (see Andersen, Bollerslev, and Diebold (2006)) and compute the realized volatility in stock returns between time t and $t+1$ as

$$vol_{t:t+1}^{j} = \sum_{t<\tau\leq t+1} \left(r_{\tau}^{j}\right)^{2}. \tag{2.20}$$

Finally, to explore the possibility that macroeconomic information predicts revisions to earnings expectations, we use the lagged 3-month T-bill rate ($Tbill$) as a predictor variable. This variable has been used extensively in the finance literature as a way of capturing variations in the state of the economy and tracking risk premia. Most studies find a negative association between stock returns and past interest rates (see, e.g. Fama and French (1988), Fama and Schwert (1977) and Ferson and Harvey (1991)).[15] Interestingly, the nominal T-bill rate has also been linked to return volatility (see, e.g. Glosten, Jagannathan, and Runkle (1993)) which, in light of our discussion in Section 2.2, suggests that this variable should be correlated with any predictable components in revisions to analysts' earnings forecasts.

In summary, at a given point in time, t, we use the revision, Δf_t^j, the cumulated stock return $r_{t-1:t}^j$, its volatility, $vol_{t-1:t}^j$, and the 3-month T-bill rate, $Tbill_t$, to predict the subsequent earnings revision between time t and $t+1$. Our timing convention ensures that all variables are known to the analysts when the consensus forecast at time t was computed.

Table 2.5 reports estimates of the slope coefficients, $\delta_{2,i,k}$, for the logit specification (2.15) with different explanatory variables. For 11 of the 30 firms, the short interest rate is significant in predicting future revisions, while the corresponding numbers are 4, 17 and 7 for past returns, past revisions and past volatility, respectively. Past revisions to the consensus earnings estimate are thus best at forecasting the direction of future

[15] Aiolfi and Rodriguez (2006) investigate how financial analysts incorporate macroeconomic information into their earnings forecasts for different industries and find that macroeconomic information helps explain revisions to the average analyst's forecast.

revisions to earnings expectations. For 23 of the 30 firms at least one covariate appears to be significant in predicting the sign of future revisions, with the interest rate and past returns clearly the best predictor variables.

2.4.4 Interest Rates and Revisions to Earnings Expectations

To better see how the probability of a negative, zero or positive state is affected by the lagged interest rate, Figure 2.7 plots the time-series of the estimated transition probabilities for Caterpillar and Coca Cola. Variations in the short interest rate give rise to substantial changes in the probability of staying in the current state. Lower interest rates (observed towards the end of our sample) increase the probability of continued positive revisions to earnings expectations (shown in the top left window) or a shift from a negative towards a positive earnings revision (bottom right window). Conversely, higher interest rates increase the probability of continued negative revisions (shown in the bottom left window) and also increase the probability of a shift from a positive to a negative revision in earnings expectations. Judged by their magnitude changes in these probabilities can be very large and variations of up to 40% in $p_{1,1}$ or $p_{0,0}$ are observed during our sample.

Figure 2.8 shows predicted along with 'actual' values, i.e. revisions to the consensus earnings expectations. These plots are based on the three-state model that has the T-bill rate as a predictor variable. The individual plots highlight a number of interesting points. First, even though no model can reasonably be expected to capture the large outliers in revisions to earnings expectations and we should expect the correlation between the predicted and actual value to be quite low, it is clearly feasible to predict the direction of the earnings forecasts, i.e. the sign of the revision. Secondly, the predictor variables manage to identify important components of the variation in revisions to earnings expectations, a point that is particularly visible in the plots for Alcoa, du Pont, 3M and Exxon Mobil.

So far we have assumed that the predictor variables, x_t, only affect the state transitions but not the mean of the revision in earnings expectations within each state. It is natural, however, to consider whether the predictor variables directly affect the mean earnings revision within each state. We do so by modifying the earlier model as follows

$$\Delta f_{t+1} = \beta_{1,s_{t+1}} + \beta_{2,s_{t+1}} x_t + \epsilon_{t+1}, \quad \epsilon_{t+1} \sim N\left(0, \sigma^2_{s_{t+1}}\right). \tag{2.21}$$

Since the evidence in support of time-varying mean revisions to earnings expectations is strongest for the short interest rate ($Tbill$), we confine our discussion to this model, results from which are displayed in Table 2.6. When entering this model, the T-bill rate is significant for almost every single firm.

Higher interest rates appear to have two separate effects on how analysts revise their earnings expectations: A magnitude (level) and a persistence effect. Higher interest rates lead to smaller expected future revisions in the consensus estimate: In both the positive and negative revision state all significant estimates of $\beta_{2,1}$ are negative. Higher interest rates also reduce the probability of staying in the positive revision state or, equivalently, increase the probability of transitioning to the negative revision state or staying there if past revisions were negative. Both effects suggest an adverse effect of interest rates on future revisions to earnings expectations.

2.5 Predictability of Actual and Expected Earnings

A concern that is often raised in the forecasting literature is that evidence of in-sample predictability−i.e. predictability of some variable established on the same sample used to fit the model−is too weak a threshold to pass since the model could overfit the data. A remedy used to deal with this problem is to evaluate the model's forecasting performance in a subsequent out-of-sample period. Models that have overfitted the data can expect to see their forecasting performance deteriorate significantly when evaluated on fresh (out-of-sample) data.

Table 2.5: Parameter Estimates for the Sign Mixture Model with Time-Varying Transition Probabilities

A sign mixture model was estimated for revisions to the consensus earnings forecasts of each of the DJ30 firms: $log(|\Delta f_t|) = \beta_{1,s_t} + \epsilon_t, \quad \epsilon_t \sim N\left(0, \sigma^2_{s_t}\right)$ with state transition probabilities $p_{i,k} = P\left(s_t = i | s_{t-1} = k\right) = \Phi\left(\delta_{1,ik} + \delta_{2,ik} x_{t-1}\right)$ where x_{t-1} is the lagged value of the 3-month T-bill rate (Tb), past revisions (Rev), past returns (Ret) and past volatility (Vol) and $i, k = 1, 0, -1$. ** and * denote significance at the 5% and 10% level respectively.

	Tb			Ret			Rev		Vol		
	$\delta_{2;1,1}$	$\delta_{2;0,0}$	$\delta_{2;-1,-1}$	$\delta_{2;1,1}$	$\delta_{2;0,0}$	$\delta_{2;-1,-1}$	$\delta_{2;1,1}$	$\delta_{2;-1,-1}$	$\delta_{2;1,1}$	$\delta_{2;0,0}$	$\delta_{2;-1,-1}$
AA	0.12	0.23	-0.31	-0.06	0.07	0.07	1.66*	1.91**	0.52	-0.57	-0.53
AIG	-0.42	0.25	-0.63	1.37	-0.13	0.24	-2.69	-0.52	-0.23	1.64	1.70
ALD	-1.45**	-0.03	-0.22	0.42	0.07	-0.21	0.43	2.28**	-0.24	-0.16	-0.38
AXP	-0.28	0.27	0.43*	-0.24	0.02	-0.07	1.33	1.93**	-12.13*	0.01	1.17
BA	0.27	-0.21	0.19	0.28	-0.36	-0.24	1.40*	2.67**	-0.25	-0.46	-0.29
BEL	-0.43	0.36*	-0.33	-0.32	0.37	0.14	-0.66	1.35	-0.09	-0.41	-0.30
CAT	-0.40*	0.37	0.22	-0.11	-0.05	0.13	0.59	0.08	0.13	-0.94	-1.74
CCC2	0.36	0.57**	-0.26	0.13	-0.05	-0.53	-0.25	4.01**	1.17	-0.70	0.41
CHL	-0.55**	-0.33	0.11	0.78**	-0.33	-0.15	0.08	2.09**	1.91	0.36	0.11
DD	0.30	-0.04	-0.17	-0.48	-0.25	-0.76*	2.05**	1.18*	-3.30	3.58	0.85
DIS	-0.21	0.38*	0.22	0.97	-0.26	-0.58	-0.03	0.33	-2.76	0.40	0.77
GE	—	0.40*	—	—	0.07	—	—	—	—	0.16	—
GM	-0.38*	-7.15	-0.01	0.17	-5.24**	-0.20**	0.98**	0.88**	0.51	-3.40	-0.23
HD	-0.86*	0.67**	-0.45	0.28	0.31*	0.07	3.55*	2.14	-2.82	-0.30*	0.22
HWP	-0.00	0.20	-0.12	0.31	-0.14	-0.56*	-2.38	1.97**	-0.84	-0.48	1.89*
IBM	0.31	-0.05	0.21	0.07	-0.10	-0.70**	3.79**	2.66**	-0.25	0.34	-0.61
INTC	-0.09	0.41**	0.01	0.84*	0.08	-1.01**	1.15	-0.37	-0.49	-0.01	1.86*
JNJ	-0.30	0.38**	-0.28	0.57	0.19	0.17	6.37*	1.57	-12.17	-0.06	-5.75
KO	-1.37*	0.67**	0.18	-0.81	-0.12	-0.37	-9.84	2.98*	-18.24	3.10	4.47
MCD	-3.44	0.62**	0.23	-0.18	-0.14	-0.19	6.52**	2.40*	-7.17	0.39	0.93
MMM	-0.12	0.08	-0.32	-0.34	0.17	-0.51	1.64	1.42	-1.69	-0.13	3.87
MO	-0.32	0.18	-0.15	1.54*	-0.24	-0.19	3.76	0.06	-4.52	11.19**	-0.31
MRK	-0.31	0.45**	-0.64	0.53	-0.08	-0.58	3.78	1.40	2.73	0.93	-7.19
MSFT	-0.74	0.47**	-0.09	-0.15	0.00	-1.77*	1.26	-4.55	-0.94	-0.06	-1.00
PFE	—	0.05	—	—	-0.03	—	—	—	—	0.33	—
PG	0.03	0.51**	0.59	0.76	0.59*	0.12	-5.07	-1.69	-4.50	-0.70	-0.70
SBC	-0.70	0.73**	-0.35	2.28**	0.06	0.12	2.43	2.30*	1.36	0.06	-0.16
UTX	0.68	0.62**	-0.27	0.97	-0.46	0.42	-1.45	6.86**	-10.36	1.90	15.69*
WMT	0.41	0.75**	0.78*	-0.15	0.38**	-0.44	-7.76**	6.23**	4.93*	-0.24	7.89**
XON	-0.33	0.18	-0.29	0.08	0.18	-0.25	3.04**	1.08	-0.10	-12.79	-1.23

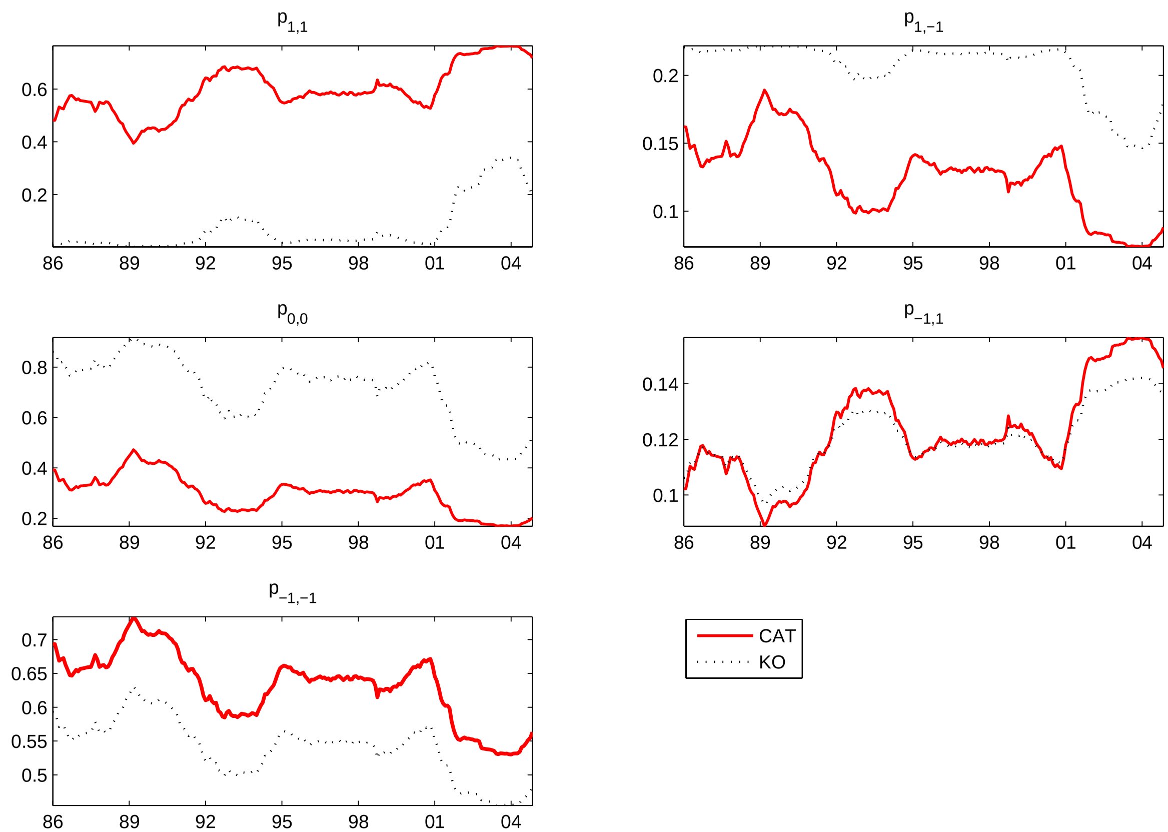

Figure 2.7: Time-varying transition probabilities for Caterpillar (CAT) and The Coca-Cola Company (KO) during 1986-2004 obtained from a three-state model with time-varying transition probabilities that depend on the lagged T-bill rate. State 1 captures positive earnings revisions, state 0 zero revisions and state -1 negative earnings revisions.

In the context of predictability of stock returns, Goyal and Welch (2003) make this point forcefully. They find that while the dividend yield appear to forecast stock returns in-sample, predictability breaks down out-of-sample. Studies such as Kang, O'Brien, and Sivaramakrishnan (1994) and Lys and Sohn (1990) find evidence of in-sample predictability of revisions to analysts' earnings forecasts.[16] However, no previous study has documented whether this carries over to out-of-sample predictability, so this evidence must be viewed with some skepticism since it is not clear if it could genuinely have been exploited to improve real-time forecasts. Our contiguous time series of earnings revisions is ideally suited to test for out-of-sample predictability.

In our forecasting experiment, we keep the last 10 years of data for each firm (or, equivalently, 120 monthly observations) as the out-of-sample period. The preceding data is used to estimate the parameters of the forecasting model. This is all done in 'real time' in order to simulate the forecasts that an analyst with access only to the historical data could have computed (i.e., without the benefit of hindsight). Suppose, for example, that in December 1995 we are interested in forecasting the revision to earnings expectations for January 1996. Then we only use data up to December 1995 to estimate the parameters of our model. These estimates are in turn used to predict the revision to the earnings estimate for January 1996. The subsequent month we add one monthly observation and re-estimate the model so that the forecast for February 1996 only incorporates data up to January 1996 and so forth.

This approach has the advantage that it takes the effect of parameter estimation error and model misspecification into account: If our model fits the revisions data poorly or its parameters are imprecisely estimated, this will be reflected in imprecise forecasts. Of course, the forecasting performance evaluated in this way sets the bar very high, since the more traditional in-sample or 'fitted' values almost always perform better since they have the benefit of using data from the full sample.

As a benchmark for evaluating our model's forecasting performance, we also report predictions generated by more traditional time-series models that take revisions to earn-

[16] Lys and Sohn (1990) regress forecast revisions on the cumulative returns of the market portfolio as well as the firm's cumulative stock returns since the forecast and surrounding the revision. They find that revisions to analysts' earnings forecasts reflect some, but not all, of the new information available between the dates of the forecasts and revisions.

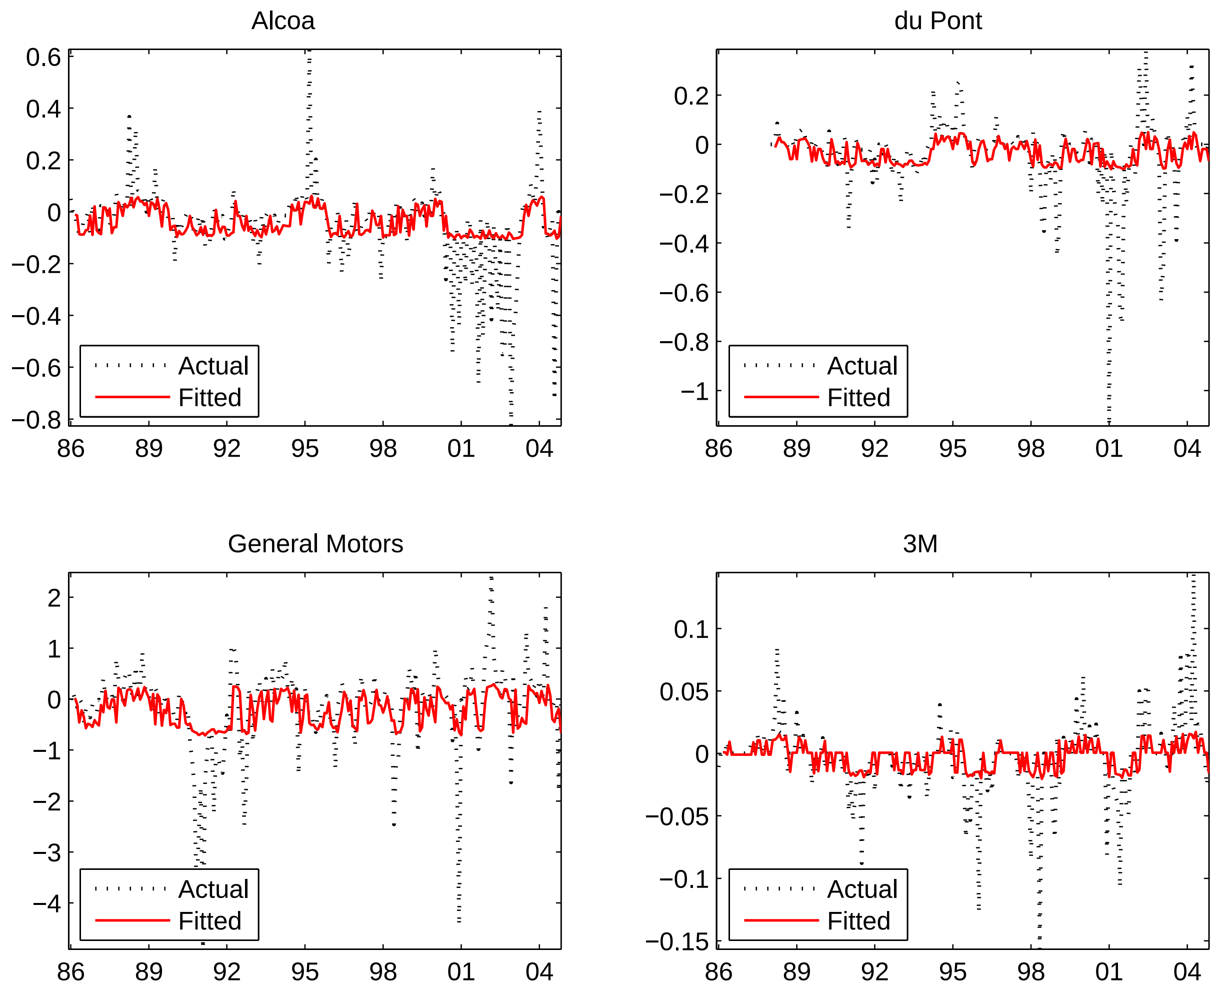

Figure 2.8: Revisions in the consensus earnings forecast between consecutive months (scaled by the initial stock price) for selected DJ30 companies over the sample Jan 1986 - Dec 2004. Fitted values are obtained from a three-state model with time-varying transition probabilities based on past earnings revisions.

Table 2.6: Estimates for the Mixture Model with Time-Varying Mean and Transition Probabilities

A sign mixture model was estimated for revisions in the consensus earnings forecast for each of the DJ30 firms: $log(|\Delta f_t|) = \beta_{1,s_t} + \beta_{2,s_t} Tb_{t-1} + \epsilon_t, \quad \epsilon_t \sim N\left(0, \sigma^2_{s_t}\right)$ where Tb_{t-1} is the lagged value of the 3-month T-bill rate and $p_{i,k} = P\left(s_t = i | s_{t-1} = k\right) = \Phi\left(\delta_{1,ik} + \delta_{2,ik} Tb_{t-1}\right)$ are the state transition probabilities between states $i, k = 1, 0, -1$. State 1,0,-1 capture positive, zero and negative earnings revisions, respectively. (-) indicates too few transitions between states to allow estimation of the parameters. ** and * denote significance at 5% and 10% level, respectively.

	$\beta_{2,1}$	$\beta_{2,-1}$	$p_{1,1}$		$p_{1,-1}$	$p_{0,0}$		$p_{-1,1}$	$p_{-1,-1}$	
			δ_1	δ_2	δ_1	δ_1	δ_2	δ_1	δ_1	δ_2
AA	-0.13	-0.35**	1.01	0.12	-0.00	-2.26	0.23	0.24	3.02**	-0.31
AIG	-0.24**	-0.40**	-1.18	-0.42	-1.95**	0.35	0.25	-1.18**	1.07	-0.63
ALD	-0.20*	-0.12	1.53*	-1.45**	-0.61*	0.47	-0.03	-0.57	1.03*	-0.22
AXP	-0.29**	-0.10	0.22	-0.28	-0.85**	0.05	0.27	-0.31	0.05	0.43*
BA	-0.31**	-0.46**	-0.55	0.27	-0.73**	-0.37	-0.21	0.06	0.90	0.19
BEL	-0.36**	-0.28**	-0.36	-0.43	-1.39**	-0.12	0.36*	-0.80**	0.99	-0.33
CAT	-0.07	-0.20*	1.73**	-0.40*	-0.79**	-1.76**	0.37	-0.69*	0.43	0.22
CCC2	-0.17**	-0.30*	-0.09	0.36	-1.23**	-1.54**	0.57**	-0.13	1.41*	-0.26
CHL	-0.21	0.08	2.72**	-0.55**	0.65*	-0.34	-0.33	0.16	1.15**	0.11
DD	-0.41**	-0.35**	-0.08	0.30	-0.44	-1.25	-0.04	-0.53	1.80**	-0.17
DIS	-0.53**	-0.69**	-0.12	-0.21	-1.22**	-0.62	0.38*	-0.89**	0.50	0.22
GE	-0.32**	-0.61**	—	0.84**	-0.93**	0.95**	0.40**	—	—	1.30**
GM	-0.30**	0.08	3.37**	-0.38*	1.43**	—	-5.53**	2.67**	3.75**	-0.01
HD	-0.44**	-0.19	0.10	-0.86*	—	-0.41	0.67**	—	0.03	-0.45
HWP	-0.44**	-0.38**	-0.94	-0.00	-2.05**	-0.73	0.20	-1.39**	0.58	-0.12
IBM	-0.05	-0.12	-0.55	0.31	-0.56	-0.21	-0.05	-0.13	1.08*	0.21
INTC	-0.50**	-0.40**	0.16	-0.09	-1.06**	-0.70	0.41**	-1.15**	0.46	0.01
JNJ	-0.23**	-0.02	-1.00	-0.30	—	0.31	0.38**	—	-0.02	-0.28
KO	-0.11*	-0.14	0.20	-1.37*	-1.25**	-0.58	0.67**	-1.04**	0.04	0.18
MCD	-0.51**	-0.61**	3.49**	-3.44	—	-0.38	0.62**	—	-1.32	0.23
MMM	-0.24**	0.08	-0.19	-0.12	—	0.07	0.08	-1.53**	1.42*	-0.32
MO	-0.28**	-0.25	-0.49	-0.32	-1.61**	-0.21	0.18	-1.95**	-0.14	-0.15
MRK	-0.26**	-0.80**	-0.92	-0.31	-2.01**	-0.19	0.45**	-1.16**	-0.47	-0.64
MSFT	-0.67**	-0.19	-0.54	-0.74	—	-0.19	0.47**	-0.85	-0.19	-0.09
PFE	-0.15	-0.36**	—	1.57**	—	1.45**	0.05	—	—	1.18**
PG	-0.07	0.06	-0.68	0.03	-1.25**	-0.46	0.51**	-0.69	-1.62	0.59
SBC	-0.55**	-0.49**	-0.14	-0.70	-1.50**	-1.66**	0.73**	-1.14**	0.60	-0.35
UTX	-0.15	-0.05	-4.09	0.68	-1.99**	-1.71**	0.62**	-1.95**	1.23	-0.27
WMT	-0.04	0.27**	-1.24	0.41	—	-0.69	0.75**	-1.10**	-2.03*	0.78*
XON	-0.61**	-0.17*	2.17**	-0.33	-0.06	-1.67	0.18	0.13	1.87**	-0.29

ings expectations as the dependent variable and always includes the lagged revision in addition to other explanatory variables (x):

$$\Delta f_{t+1} = \beta_0 + \beta_1 \Delta f_t + \beta_2 x_t + \varepsilon_{t+1}. \tag{2.22}$$

Here x_t is selected from the same set of covariates used in the three-state model, i.e. the lagged T-bill rate, past stock returns or past volatility.

2.5.1 Out-of-Sample Forecasting Results

Before proceeding to the empirical results, it is worth pointing out that even the simple model with constant transition probabilities (2.14) and a constant mean and variance parameter within each state can capture time-variations in revisions to analysts' earnings expectations. This seemingly surprising property arises because the state probabilities vary over time. To see this, notice that, using properties of the log-normal distribution, the predicted value of the earnings revision next period as a function of the current state, s_t, is

$$\Delta \hat{f}_{t+1} = \begin{cases} p_{1,1,t}\exp(\mu_1 + 0.5\sigma_1^2) - p_{1,-1,t}\exp(\mu_{-1} + 0.5\sigma_{-1}^2) & \text{if } s_t = 1 \\ (1 - \frac{p_{0,0,t}}{2})(\exp(\mu_1 + 0.5\sigma_1^2) - \exp(\mu_{-1} + 0.5\sigma_{-1}^2)) & \text{if } s_t = 0 \\ p_{-1,1,t}\exp(\mu_1 + 0.5\sigma_1^2) - p_{-1,-1,t}\exp(\mu_{-1} + 0.5\sigma_{-1}^2) & \text{if } s_t = -1 \end{cases}. \tag{2.23}$$

Clearly the predicted earnings revision will vary as a function of the current state, s_t, whenever $\mu_1 \neq \mu_{-1}$ or $\sigma_1 \neq \sigma_{-1}$. If state transitions are allowed to vary through time, this gives rise to further variation in the predicted consensus revisions.

The most common measure of forecasting performance is the correlation between the actual revision to earnings expectations and the predicted value. Values of this measure, computed out-of-sample, are reported in Table 2.7. We show results both for the three-state model with constant transition probabilities ('constant') in addition to the models that allow for time-varying state transitions and predictions from a first-order autoregressive (AR(1)) model. The three-state models generate average correlations around 0.20, with a slightly higher median value (0.22) for the model that uses the past interest rate as a predictor variable. The three-state models generate significantly higher correlations with actual revisions than the linear autoregressive models accomplish and generates R^2-values nearly twice as high as those observed for the linear model.

Table 2.8 reports the proportion of correctly predicted positive, zero and negative values of the revisions to earnings expectations−a statistic commonly referred to as the hit rate.[17] Viewed across firms the average hit rate for the three-state models is 58%.[18] The autoregressive model generates hit rates around 48 percent−almost 10% below those recorded for the three-state models. This indicates a clear advantage from incorporating the separate information in the sign of past revisions and is perhaps to be expected in view of our findings in Section 2.3 that the persistence of negative, 'no change', and positive revisions to earnings expectations as well as the magnitude of positive and negative consensus revisions are so different.

2.5.2 Predicted Revisions and Actual Earnings Per Share

So far our analysis focused on modeling the process whereby analysts revise their earnings expectations. This allowed us to understand analyst behavior and the factors determining the updates to analysts' beliefs. A question of separate interest is whether our findings of predictability in revisions to analysts' earnings expectations translate into an ability to forecast the 'actuals', i.e. the realized earnings figure announced by the firms once a year.

To see if this is the case, we estimate the following model for the actual annual

[17]The likelihood that the predicted value from the AR(1) model is exactly equal to zero is of course very low. To be consistent with our treatment of the consensus earnings estimate in the three-state model, we therefore categorize forecast revisions from the autoregressive model as zero whenever the predicted earnings revision is smaller than one cent per share.

[18]We also computed predictions for the three-state models extended to incorporate time-varying mean earnings revisions within each state (using (2.21)). The forecasts implied by these models are given by

$$\Delta\hat{f}_{t+1} = \begin{cases} p_{11,t}\exp(\beta_{11}+\beta_{21}x_t+0.5\sigma_1^2) - p_{1,-1,t}\exp(\beta_{1,-1}+\beta_{2,-1}x_t+0.5\sigma_{-1}^2) & \text{if } s_t=1 \\ (1-\frac{p_{00,t}}{2})(\exp(\beta_{11}+\beta_{21}x_t+0.5\sigma_1^2) - \exp(\beta_{1,-1}+\beta_{2,-1}x_t+0.5\sigma_{-1}^2)) & \text{if } s_t=0 \\ p_{-1,1,t}\exp(\beta_{11}+\beta_{21}x_t+0.5\sigma_1^2) - p_{-1,-1,t}\exp(\beta_{1,-1}+\beta_{2,-1}x_t+0.5\sigma_{-1}^2) & \text{if } s_t=-1 \end{cases}.$$

We found that the forecasting performance produced by these models was very similar to that of the simpler models in Tables 7 and 8 and therefore do not report these results separately. Results are available on request.

Table 2.7: Correlation between Revisions to Consensus Earnings Estimates and Predicted Values

This table reports the correlation coefficient between revisions to consensus earnings estimates and the predicted values from the three-state model using a constant (Const), a constant and the lagged 3-month T-bill (Tb), a constant and the lagged holding period return (Ret), a constant and the past revision (Rev), and a constant and the lagged return volatility (Vol). We also report results based on a first-order autoregressive (AR(1)) model. All forecasts are computed recursively out-of-sample using 120 monthly observations. ** and * denote significance at the 5% and 10% level, respectively.

	Three-State Model					Linear AR(1) Model			
	Const	Tb	Ret	Rev	Vol	Tb	Ret	Rev	Vol
AA	0.53**	0.52**	0.53**	0.53**	0.53**	0.36**	0.25**	0.36**	0.30**
AIG	0.30	0.31	0.30	0.24	0.18	-0.05	-0.06	-0.06	-0.06
ALD	0.22**	0.24**	0.20*	0.21**	0.06	0.01	-0.00	0.05	0.05
AXP	0.46**	0.45**	0.44**	0.40**	0.38**	0.10	0.18**	0.16*	0.16*
BA	0.23**	0.22**	0.25**	0.26**	0.24**	0.04	0.05	0.05	0.05
BEL	0.12	0.09	-0.06	0.15	0.01	-0.04	-0.04	-0.04	-0.04
CAT	0.34**	0.35**	0.33**	0.32**	0.32**	0.13**	0.12**	0.11**	0.12**
CCC2	0.24**	0.25**	0.25**	0.24**	0.24**	0.22**	0.24**	0.23**	0.22**
CHL	0.29**	0.24**	0.31**	0.36**	0.25**	-0.15	0.18**	0.01	-0.00
DD	0.43**	0.44**	0.44**	0.46**	0.44**	0.35**	0.34**	0.34**	0.34**
DIS	0.40**	0.37**	0.41**	0.40**	0.40**	0.26**	0.32**	0.32**	0.32**
GE	0.01	-0.00	0.01	0.01	0.01	0.01	-0.00	-0.01	-0.01
GM	0.35**	0.36**	0.36**	0.38**	0.34**	0.37**	0.37**	0.35**	0.36**
HD	-0.05	-0.05	-0.04	-0.01	-0.03	0.21**	0.09*	0.10*	0.10*
HWP	0.14	0.11	0.16	0.17	0.18	0.07	0.08	0.07	0.07
IBM	0.35**	0.34**	0.36**	0.32**	0.33**	0.12*	0.11*	0.07	0.07
INTC	0.13*	0.11*	0.18**	0.12*	0.17**	-0.00	0.00	-0.00	-0.00
JNJ	0.14	0.08	0.13	0.12	0.12	0.15**	0.17**	0.16**	0.15**
KO	0.14	0.12	0.08	0.13*	0.18*	0.12**	0.12**	0.16**	0.16**
MCD	0.34**	0.36**	0.33**	0.35**	0.32**	0.18**	0.20**	0.20**	0.20**
MMM	0.42**	0.40**	0.44**	0.41**	0.43**	0.26**	0.24**	0.24**	0.24**
MO	-0.05	-0.05	-0.05	-0.05	-0.06	0.12	-0.01	-0.03	-0.03
MRK	0.04	0.09	-0.08*	-0.06	0.11	0.15	-0.08**	-0.08*	-0.08*
MSFT	0.04	-0.02	0.01	0.04	-0.01	-0.03	0.04	0.03	0.04
PFE	-0.14	-0.14	-0.13	-0.14	-0.14	0.05	-0.01	-0.01	-0.00
PG	0.08	0.02	0.01	0.09	0.07	-0.02	-0.03	-0.02	-0.02
SBC	0.19	0.20	0.19	0.11	0.29**	0.04	0.05	0.05	0.05
UTX	-0.25**	-0.26**	-0.31**	-0.25**	-0.29**	-0.06	0.06	-0.07	0.07**
WMT	0.19**	0.22**	0.17**	0.21**	0.23**	0.27**	0.27**	0.28**	0.28**
XON	0.56**	0.55**	0.55**	0.62**	0.54**	0.67**	0.66**	0.66**	0.66**
Average	0.21	0.20	0.19	0.20	0.20	0.13	0.13	0.12	0.12

Table 2.8: Percentage of Correctly Predicted States

This table reports the percentage of correctly predicted states representing positive, zero or negative signs of the revisions to the consensus earnings forecast. Transition probabilities are modeled using a constant (Const), a constant and the lagged 3-month T-bill rate (Tb), a constant and the lagged holding period return (Ret), a constant and the lagged revision to the earnings forecast (Rev), and a constant and the lagged return volatility (Vol). Forecasts were also computed using a linear autoregressive model with a constant, the lagged earnings revision and each of the covariates. All forecasts are computed recursively out-of-sample, using 120 monthly observations.

	Three-State Model					Linear AR(1) Model			
	Const	Tb	Ret	Rev	Vol	Tb	Ret	Rev	Vol
AA	0.71	0.70	0.71	0.71	0.71	0.63	0.51	0.60	0.59
AIG	0.69	0.73	0.68	0.69	0.72	0.61	0.63	0.62	0.63
ALD	0.58	0.58	0.58	0.57	0.57	0.49	0.47	0.53	0.50
AXP	0.69	0.66	0.68	0.67	0.68	0.13	0.12	0.13	0.13
BA	0.49	0.48	0.47	0.47	0.45	0.40	0.41	0.40	0.41
BEL	0.61	0.57	0.59	0.62	0.54	0.50	0.48	0.49	0.49
CAT	0.53	0.52	0.52	0.49	0.52	0.42	0.36	0.36	0.37
CCC2	0.52	0.52	0.49	0.53	0.51	0.45	0.53	0.48	0.49
CHL	0.61	0.60	0.63	0.62	0.61	0.38	0.38	0.37	0.38
DD	0.56	0.55	0.53	0.56	0.56	0.47	0.41	0.45	0.43
DIS	0.50	0.50	0.49	0.49	0.50	0.47	0.51	0.49	0.49
GE	0.79	0.79	0.79	0.79	0.79	0.63	0.72	0.72	0.72
GM	0.70	0.70	0.68	0.68	0.70	0.59	0.64	0.60	0.61
HD	0.62	0.58	0.62	0.61	0.63	0.50	0.47	0.46	0.47
HWP	0.50	0.46	0.47	0.51	0.53	0.41	0.43	0.39	0.39
IBM	0.44	0.41	0.43	0.45	0.42	0.38	0.40	0.37	0.38
INTC	0.35	0.38	0.41	0.34	0.33	0.45	0.44	0.48	0.46
JNJ	0.69	0.68	0.69	0.68	0.69	0.62	0.63	0.62	0.62
KO	0.56	0.55	0.53	0.54	0.55	0.47	0.45	0.47	0.47
MCD	0.53	0.56	0.52	0.57	0.53	0.45	0.47	0.47	0.47
MMM	0.57	0.55	0.53	0.53	0.57	0.45	0.44	0.45	0.44
MO	0.54	0.52	0.53	0.53	0.47	0.44	0.41	0.45	0.45
MRK	0.64	0.64	0.63	0.64	0.64	0.46	0.52	0.51	0.51
MSFT	0.60	0.58	0.60	0.60	0.59	0.43	0.41	0.42	0.42
PFE	0.74	0.74	0.74	0.74	0.74	0.61	0.63	0.62	0.61
PG	0.54	0.54	0.54	0.55	0.54	0.50	0.49	0.50	0.49
SBC	0.48	0.53	0.48	0.53	0.51	0.40	0.41	0.41	0.41
UTX	0.67	0.67	0.65	0.67	0.67	0.45	0.50	0.45	0.47
WMT	0.46	0.46	0.41	0.47	0.48	0.42	0.44	0.44	0.44
XON	0.69	0.69	0.68	0.68	0.65	0.58	0.58	0.59	0.59
Average	0.59	0.58	0.58	0.58	0.58	0.47	0.48	0.48	0.48

earnings announced at time T, A_T :

$$A_T = \beta_0 + \beta_1 f_{T,T-1} + \beta_2 \Delta \hat{f}_{T,T-1} + \varepsilon_T, \tag{2.24}$$

where $f_{T,T-1}$ is the consensus earnings forecast produced the month prior to the earnings announcement and $\Delta \hat{f}_{T,T-1}$ is the revision to the consensus earnings estimate that our three-state model predicts will occur between period $T-1$ and T. To conduct the test, for each firm we use data on the realized earnings per share for the fiscal year since this is the variable targeted by financial analysts. For each fiscal year, I/B/E/S reports actual earnings as soon as they are released to the market. These earnings figures are then adjusted for comparability with analysts' forecasts.[19]

Panel A in Table 2.8 reports results from applying regression (2.24) to our data. For each of the four predictor variables we get a different value of $\Delta \hat{f}_{T,T-1}$ so the table shows four sets of results. As expected, analysts' estimates of the earnings figure that gets announced the following month are very precise with slope coefficients, β_1, near one and R^2-values (not shown to preserve space) that typically exceed 0.97.

As argued by Fama and French (2000), time-series estimation of earnings equations lack power due to the small number of annual earnings observations for most firms. Here, since we only have 19 annual observations on the actual earnings for each firm, we do not expect to have much power in detecting predictability from $\Delta \hat{f}_{T,T-1}$. Interestingly, however, for around half of the firms$-$most often the past T-bill rate or past revisions$-$at least one of the explanatory variables leads to a coefficient on $\Delta \hat{f}_{T,T-1}$ that is statistically significant.

The forecast error in analysts' earnings estimate, $A_T - f_{T,T-1}$, can alternatively be viewed as the predicted revision to the consensus earnings estimate, $\Delta \hat{f}_{T,T-1}$, plus a forecast error. This follows since the earnings estimate will be identical to the actual earnings figure once this has been released and thus becomes known. It is therefore natural to conjecture that the predicted revision in the consensus earnings forecast may be correlated with the forecast error, $A_T - f_{T,T-1}$. This suggests a regression closely

[19] I/B/E/S does not require analysts to forecast earnings per share in the basic or diluted format, but instead lets the majority rule. In cases where an analyst follows a firm on a basis that is different from the consensus, I/B/E/S adjusts the corresponding estimates to conform with the majority.

related to (2.24):

$$A_T - f_{T,T-1} = \beta_0 + \beta_1 \Delta \hat{f}_{T,T-1} + \varepsilon_T. \tag{2.25}$$

This regression is interesting in view of the empirical findings that analysts' earnings forecasts are superior to those generated by traditional time series models, see e.g. Elton and Gruber (1972) and Brown and Rozeff (1978).

Panel B in Table 2.8 reports the outcome of applying this regression to our data. Errors in analysts' forecasts, $A_T - f_{T,T-1}$, are often explained by the predicted revision to the earning estimate generated by the three-state models based on the short interest rate or the past revisions. Furthermore, the coefficients mostly have the sign one would expect. Of 28 significant estimates of the β_1 coefficient, 23 are positive. A forecast of a positive earnings revision one month prior to the earnings announcement thus systematically translates into the actual earnings figure coming in higher than the consensus estimate.

The consensus forecast is widely regarded to be difficult to beat and the high quality of this forecast is reflected in the high values of the R^2 in the regressions reported in Panel A of Table 2.8. We should therefore not expect to find much predictability in the consensus forecast error, $A_T - f_{T,T-1}$. Indeed, the R^2-value reported by Abarbanell and Bernard (1992) and Easterwood and Nutt (1999) from regressions of forecast errors on prior-year earnings changes lie in the range of 0.01-0.02. While the explanatory power of our prediction of the consensus revision, $\Delta \hat{f}_{T,T-1}$, is quite low in many cases, for more than half of the firms the three-state model which includes the lagged T-bill rate generates an R^2 that exceeds 10%.

2.6 Conclusion

A key to understanding financial analysts' forecasts of corporate earnings is how analysts incorporate new information in their forecasts as the time to the earnings announcement date draws closer. As part of this process analysts must balance the need for forecast accuracy versus the benefits from issuing biased forecasts. We presented a simple theoretical framework for understanding how this trade-off evolves as a function of the forecast horizon and found that the implications of this theory$-$that the magni-

Table 2.9: Forecasts of Actual Earnings

Panel A reports the results of the regression $\left(\frac{A_t}{p_{t-1}}\right) = \beta_0 + \beta_1 f_{t-1,t} + \beta_2 \Delta\widehat{f}_{t-1,t} + \epsilon_t$, where A_t is the actual earnings figure (published annually), $f_{t-1,t}$ is the previous month's consensus estimate of earnings and $\Delta\widehat{f}_{t-1,t}$ is the predicted consensus revision scaled by the stock price obtained from the three-state models. These models are based on covariates such as the lagged 3-month T-bill rate (Tb), the lagged holding period return (Ret), the lagged revision to the earnings forecast (Rev), and the lagged return volatility (Vol). Panel B reports the results of the regression $\left(\frac{A_t - f_{t-1,t}}{p_{t-1}}\right) = \beta_0 + \beta_2 \Delta\widehat{f}_{t-1,t} + \epsilon_t$. ** and * denote significance at the 5% and 10% level, respectively, for tests of $\beta_1 = 1$ and $\beta_2 = 0$.

Panel A

	Tb		Ret		Rev		Vol	
	β_1	β_2	β_1	β_2	β_1	β_2	$\beta1$	β_2
AA	0.98	0.02	0.98	-0.13	0.98	-0.23	0.98	-0.18
AIG	0.96	25.93*	0.94**	22.92	0.93**	20.71	0.93**	7.68
ALD	1.01	1.45	1.00	1.97	1.01	3.24**	1.01	3.02
AXP	1.01	1.67	1.02	1.14	1.03	0.24	1.01	1.97*
BA	1.00	0.02	1.00	-0.09	1.00	-0.09	1.00	-0.34
BEL	1.01	1.08	1.01	-0.20	1.01	-0.16	1.01	0.21
CAT	1.05**	2.06**	1.06**	1.86**	1.06**	1.91**	1.05**	1.81**
CCC2	1.00	1.30**	1.00	1.29**	1.00	1.12**	1.00	0.79
CHL	0.97	0.51	0.97	1.45	0.97	0.43	0.97	0.97
DD	1.01	0.65	1.01	0.63	1.01	0.67	1.01	0.71
DIS	1.00	0.75	1.01	0.82	1.01	0.73	1.00	0.70
GE	1.01	6.78	1.01	14.36	1.01	8.68	1.01	2.59
GM	1.02	0.09	1.03	-0.09	1.02	0.00	1.02	0.05
HD	1.02**	-0.99**	1.02**	-0.89**	1.02**	-1.22**	1.02**	-0.86**
HWP	1.05**	3.05**	1.05**	1.93	1.06**	3.56**	1.05**	3.20**
IBM	1.02	0.75	1.02	0.64	1.02*	0.84**	1.02*	0.80
INTC	1.03**	1.86**	1.03**	1.52*	1.03**	1.72**	1.04**	1.72**
JNJ	1.01	4.49	1.01	4.85	1.01	4.62	1.01	4.25
KO	1.02**	1.56*	1.02**	2.82**	1.02**	3.00**	1.02**	2.43**
MCD	1.00	0.29	1.00	0.32	1.00	0.33	1.00	0.36
MMM	1.00	0.91*	1.00	0.91	1.00	1.00*	1.00	1.10**
MO	0.99**	-4.37**	0.99**	-3.77**	0.99**	-3.74**	0.99**	-3.71**
MRK	1.00	0.95	1.00	2.11	1.00	0.32	1.00	0.89
MSFT	1.00	-2.52*	1.00	-0.11	1.00	-2.55	1.00	-0.73
PFE	1.00	1.12	1.00	3.02	1.00	2.94	1.00	2.69
PG	1.03	-1.57	1.03*	-2.49	1.03	-1.54	1.03	-1.88
SBC	1.18**	14.62	1.15**	14.93	1.11*	-0.64	1.15**	15.50
UTX	1.03	3.08**	1.02	4.15**	1.03	2.75**	1.02	3.63**
WMT	1.00	0.80	1.00	1.63	1.00	0.92	1.00	-0.25
XON	1.04**	1.59**	1.04**	1.34**	1.04**	2.07**	1.04**	1.52**

Panel B

	Tb		Ret		Rev		Vol	
	β_2	R^2	β_2	R^2	β_2	R^2	β_2	R^2
AA	-0.06	0.00	-0.21	0.01	-0.25	0.01	-0.26	0.01
AIG	32.40**	0.31	24.44	0.13	17.97	0.07	6.95	0.02
ALD	1.22	0.04	1.80	0.07	2.99**	0.19	2.56	0.11
AXP	1.78	0.12	1.30	0.06	0.41	0.01	2.05*	0.18
BA	0.03	0.00	-0.09	0.00	-0.08	0.00	-0.34	0.01
BEL	1.31	0.01	0.06	0.00	0.07	0.00	0.47	0.00
CAT	2.32**	0.22	2.04**	0.19	2.06**	0.20	2.05**	0.20
CCC2	1.31**	0.30	1.26**	0.27	1.11**	0.38	0.81	0.16
CHL	0.11	0.00	0.93	0.01	-0.01	0.00	0.65	0.00
DD	0.69	0.10	0.67	0.10	0.68	0.11	0.73	0.11
DIS	0.74	0.05	0.75	0.06	0.69	0.04	0.67	0.04
GE	0.23	0.00	17.79	0.05	12.47	0.03	6.40	0.01
GM	0.20	0.00	0.00	0.00	0.14	0.00	0.13	0.00
HD	-4.30**	0.47	-3.29**	0.23	-4.70**	0.29	-3.09**	0.24
HWP	2.21	0.11	1.50	0.05	1.47	0.04	1.97	0.06
IBM	0.89	0.12	0.82	0.11	0.91**	0.20	0.90	0.12
INTC	1.60*	0.14	1.20	0.08	1.55*	0.14	1.36	0.11
JNJ	6.21*	0.17	6.66*	0.18	8.03*	0.16	6.01**	0.19
KO	1.05	0.05	1.61	0.07	1.18	0.03	1.53	0.08
MCD	0.18	0.02	0.31	0.08	0.29	0.09	0.32	0.09
MMM	0.92*	0.16	0.92	0.13	1.00**	0.19	1.10**	0.20
MO	-2.93	0.12	-3.23	0.13	-2.84	0.10	-3.55**	0.22
MRK	0.58	0.03	1.95	0.12	0.34	0.03	0.82	0.03
MSFT	-2.11*	0.19	-0.24	0.00	-2.02	0.14	-0.80	0.02
PFE	6.88	0.01	9.22	0.01	9.22	0.01	8.75	0.01
PG	-0.80	0.00	-1.17	0.00	-0.05	0.00	-0.33	0.00
SBC	-11.24	0.10	1.18	0.00	-9.84	0.12	-0.53	0.00
UTX	3.37**	0.34	4.40**	0.52	2.94**	0.45	3.81**	0.64
WMT	0.32	0.01	0.90	0.01	0.52	0.01	-0.33	0.01
XON	1.55	0.12	1.27	0.09	2.48**	0.32	1.46	0.11

tude of the bias shrinks while the forecast accuracy increases as the forecast horizon is reduced−could be confirmed empirically. Forecast accuracy appears to become more important as the earnings announcement date is approached, more information is available and biases become both more easily detectable and (as a result) more costly.

Our analysis also uncovered strong asymmetries and persistence in revisions to earnings expectations, in part related to the magnitude of past revisions−small revisions are more likely to follow small revisions−in part related to the sign of the revision in earnings expectations: Continuation of revisions of the same sign are more likely than sign reversals. Variables such as past revisions and the lagged interest rate were found to predict future revisions in analyst earnings expectations. These variables contain information that is relevant not only for the revisions to analysts' earnings forecasts but also for the actual earnings figures which is what ultimately matters to investors and financial analysts. Higher interest rates tend to reduce both the magnitude and the probability of a positive future revision. If earnings expectations are lowered following higher interest rates, ceteris paribus stock prices should also decline. This holds in addition to the increased discount rate channel that investors are likely to apply to future earnings when interest rates increase and can thus help explain the significant negative relation between stock returns and interest rates.

Our results on asymmetries and persistence in revisions to the consensus earnings estimate bear an interesting relation to recent findings by Conrad, Cornell, Landsman, and Rountree (2006) that whereas analysts are equally likely to upgrade or downgrade a stock following a large increase in the stock price, they are more likely to issue a downgrade following a large negative movement in the stock price. Consistent with our findings of persistence in the direction of earnings forecasts, they also find that analysts recommendations are "sticky" and that analysts appear reluctant to issue a downgrade−a finding confirmed by Clarke, Ferris, Jayaraman, and Lee (2006).

Time-series data on revisions to analysts' earnings expectations could in principle allow us to relate the well-documented properties of stock returns−such as skewness, fat tails and volatility clustering (autoregressive conditional heteroskedasticity (ARCH))−to the revision process for earnings expectations. As pointed out by Kothari (2001, p. 145), analysts forecasts "... affect the information environment and influence the level and variability of security prices." If analysts' earnings expectations were all that mattered to stock prices, then there should be a one-to-one mapping between the properties of stock returns and properties of the revisions to earnings expectations. This tight link is broken when expected returns can vary over time and are linked to earnings news. Nevertheless, further analysis of the time-series of analyst earnings expectations offers the potential of gaining additional insights into the properties of stock returns and the role of financial analysts in transmitting earnings news.

This chapter is based on *Managing Earnings Expectations: Persistence, Asymmetry and Predictability in Analysts' Earnings Forecasts*, joint with Marco Aiolfi and Allan Timmermann.

3

Financial Analysts' Forecast Revisions and Macroeconomic Information

3.1 Introduction

Financial analysts periodically provide earnings per share forecasts for publicly-traded firms. Market participants pay close attention to such forecasts as a summary of the market's expectations of a firm's financial health. Moreover, as analysts update their expectations about a firm's earnings per share potential, more information is revealed to the market about the firm's condition. As such, analysts' forecasts and, more importantly, analysts' revisions of these forecasts play an important role in many portfolio decisions. In fact, Stickel (1991), Park and Stice (2000), Liu and Thomas (2000), and Gleason and Lee (2002) show that analysts' forecast revisions are correlated with movements in stock prices.

It is a well-known fact in the financial economics literature that corporate earnings tend to follow macroeconomic movements and industry-wide events (e.g., O'Brien (1994), Brown and Ball (1967), Gonedes (1973), and Magee (1974)). What is not clearly understood in the literature is how financial analysts absorb new macroeconomic information into their predictions of earnings per share. For any given firm, earnings per share is defined to be net income (revenues less expenses) divided by the average number of common shares in a given period. Therefore, we should expect analysts' prediction updates to depend on the firm's operating characteristics since different macroeconomic scenarios would affect the revenue and expense structure in different ways. In this chapter, we decompose financial analysts' corporate earnings prediction updates into the contributions from firm-specific attributes and the contributions made by new macroeconomic information. With the exception of Abarnanell and Bushee (1997), who consider whether or not fundamental variables have an impact on earnings revisions during periods of high and low GDP growth and inflation, we are not aware of any study that disentangles the effects of firm characteristics versus the economic environment on financial analysts' forecast revisions.

We also investigate how financial analysts incorporate macroeconomic information into their earnings per share forecast revisions for different industries. It is clear that new macroeconomic information does not affect corporate earnings equally across industries. For instance, earnings for firms in highly leveraged or capital-intensive industries may be more sensitive to interest rate fluctuations through a cost-of-capital channel than earnings for firms in labor-intensive or nondurables industries. Similarly, highly consumer demand-driven industries are likely to have earnings profiles more responsive to real output and employment measures. Financial analysts following firms in these industries are likely to rely more heavily on macroeconomic information when updating their predictions.

To address the question of whether or not analysts' revisions can be explained by the behavior of macroeconomic variables, we compile data from three different sources. First, we exploit the long history of earnings per share forecasts by analyst and firm, available in the Institutional Brokers' Estimate System to create a continuous forecast revision for the average analyst. Next, we collect stock price information from the Center for Research in Security Prices to compute firm-specific stock returns. Finally, we follow Ludvigson and Ng (2006) and use the many macroeconomic time series available from Stock and Watson (2005) as our indicators of the macroeconomic environment.

Since the cross-section dimension of the many macroeconomic variables available from Stock and Watson (2005) is extremely large, we follow the methodology proposed in Stock and Watson (2002a, 2002b, 2004) in order to achieve dimension reduction. These papers show that consistent estimates of the space spanned by the common factors may be constructed using principal components analysis. This procedure has also been used to forecast measures of macroeconomic activity (e.g., Stock and Watson (2002a, 2002b, 2004)). We follow the literature and estimate a time series of common factors from each set of macroeconomic variables using principal components analysis. By summarizing the information from a large number of series in a few estimated factors, we eliminate the arbitrary reliance on a small number of imperfectly measured indicators to proxy for macroeconomic fundamentals, and make feasible the use of a vast set of economic variables that are more likely to span the unobservable information sets of financial market participants.

Our results indicate that accounting for macroeconomic information enhances the understanding of the average financial analyst's forecasting pattern. The economic environment in which a firm operates is most often positively correlated with financial analysts' updates. Our financial condition measure as well as our metric for inflationary pressure are usually negatively associated with forecast revisions, the main exceptions being the chemical and telecommunication sectors. By including these macroeconomic factors as regressors to explain the dynamics of forecast revisions by financial analysts, the average adjusted R^2 is double the average adjusted R^2 for regressions including only firm-level components, such as a firm's cumulative past excess stock return. Moreover, not surprisingly, the impact of macroeconomic factors has different implications for different industries. Our estimates suggest that analysts are more likely to update predictions for more capital-intensive, consumer demand-driven, or highly-leveraged industries. Furthermore, the magnitude and direction of the coefficients on the macroeconomic factors vary significantly across sectors, offering further evidence to support our study at the industry-level.

We argue these results offer a contribution to the literature on financial analysts' forecast revisions. In particular, we present evidence of the explanatory power of macroeconomic factors in determining prediction updates. Previous work neglecting these macroeconomic indicators may suffer from an omitted variable bias. In addition, it is imperative that the literature consider the effects of different factors separately for different industries, as the economic environment is likely to differentially affect operating conditions in different industries.

The rest of this chapter is organized as follows. In the next section, we offer a brief discussion of the previous literature on the correlation between forecast revisions and stock price movements. In Section 3.3, we describe the data used for estimation and offer a descriptive analysis of the data in Section 3.4. Section 3.5 lays out the econometric framework, adapted from Stock and Watson (2002a, 2002b, 2004). The empirical results from our estimation are presented in Section 3.6 and we derive the main conclusions in the final section.

3.2 Extant Literature

The process of forecast updating should reflect the arrival of new information either through firm-specific or through other factors that characterize the economic environment in which traded firms operate. The existing literature has focused on whether or not analyst forecast revisions are linked to firm-specific stock price movements. The main papers, discussed in detail in Section 3.2.1, vary by the data source used as well as the measurement and definition of the variables in question. We also offer support from the literature on the influence of macroeconomic indicators on corporate earnings, as a motivation for the current study.

3.2.1 Stock Prices and Forecast Revisions

There is little argument among financial economists that stock prices reflect information about firms' fundamentals. Since internal information on a firm's conditions affects managers' investment decisions, in markets with asymmetric information, changes in stock prices may signal private information about real investment opportunities. It is not surprising, therefore, that the literature has found that changes in stock prices have a positive impact on financial analysts' earnings per share forecasts and revisions.

Brown, Foster, and Noreen (1985) find that the sign and magnitude of analysts' forecast revisions are positively correlated with the sign and magnitude of the average cumulative abnormal security returns for the twelve-month period preceding revisions. Building on this work, in order to test the cognitive bias theory of share price reversals, Klein (1990) inspects the behavior of forecast errors and forecast revisions one year later. The author finds that firms with large (small) stock returns during the year after the forecast have, on average, negative (positive) forecast errors and positive (negative) revisions one year later.

Based on the asymmetric information models presented in Diamond and Verrecchia (1981) and Glosten and Milgrom (1985), Abarbanell (1991) examines whether analysts' earnings predictions incorporate the information present in stock price changes. More precisely, he argues if price changes enhance analysts' information about fundamentals, then forecast revisions will be positively associated with prior price changes. Using a new test criterion based on the difference between the frequency of positive revisions given past positive returns and the frequency of positive revisions given past negative returns, Abarbanell (1991) finds a positive correlation between stock price changes and forecast revisions.

Lys and Sohn (1990) regress forecast revisions on the cumulative returns of the market portfolio as well as the cumulative stock returns of the firm, since the forecast and surrounding the revision. They find that analysts' earnings forecast revisions reflect some, but not all, of the new information available between the forecasts and revisions. Furthermore, even after controlling for corporate accounting disclosures, analysts' earnings revisions incorporate the information in security prices.

Cooper, Day, and Lewis (2001) construct a binary ranking of an individual analyst's performance based on the time of the earnings forecast, the abnormal trading volume associated with the forecast, and the accuracy of the forecast. Distinguishing between high-tech and low-tech firms, they find that, independent of analyst performance, revisions of high-tech firms are significantly related to past excess returns. However, for the case of low-tech firms, only low-performing analysts tend to revise expectations based on information contained in past excess returns.

In an experimental setting, Miller and Sedor (2006) vary the levels of uncertainty about future earnings. They find that analysts' earnings forecast revisions are influenced by observed price changes only when uncertainty about future earnings is high. However, when uncertainty about future earnings is low, price changes are not reflected in forecast revisions.

In contrast to the literature focusing on stock price changes, Chaney, Hogan, and

Jeter (1999) argue that the frequency and magnitude of restructuring charges may affect the average analyst's forecast revision behavior. In fact, they find that subsequent to a restructuring charge announcement, on average, analysts revise downward one- and two-year ahead forecasts.[1]

3.2.2 Macroeconomic Information

Arguably, economic conditions influence the number and types of real investment opportunities available to a firm, and hence a firm's earnings performance (e.g., Brown and Ball (1967), Gonedes (1973), and Magee (1974)). However, the hypothesis that macroeconomic information wields important effects on financial analysts' corporate earnings forecasts has little empirical support, despite the strong intuitive appeal. By ignoring these important macro variables, the literature on forecast revisions may suffer from an omitted variable bias. In fact, Brown (1993), Schipper (1991), and Ramannath, Rock, and Shane (2006) remind us of the need for research concerning the relationship between the decision process of financial analysts and the roles of macroeconomic and industry factors.

We motivate our study of the effects of macroeconomic fundamentals on analysts' earnings prediction updates with evidence from the literature on the effects of macroeconomic factors on firms' earnings potential. Veronesi (1999) suggests investor uncertainty regarding the economic environment affects the volatility of stock prices. Furthermore, the effect is larger in recessionary times; that is, investors expectations about future cash flows are more sensitive to new macroeconomic information during times of greater investor uncertainty.

Moreover, macroeconomic factors may have different implications for different industries. For instance, earnings in highly consumer demand-driven sectors, like consumer durables, may be far more sensitive to economic conditions. Likewise, firms in highly regulated sectors, like utilities and financial services, may have operating char-

[1] These studies differ not only by the methodological approach but also by the data used to measure analysts' forecasts and revisions. Klein (1990) and Chaney, Hogan, and Jeter (1999) use consensus forecast data provided by I/B/E/S, while Cooper, Day, and Lewis (2001) use individual analyst data from I/B/E/S. Abarbanell (1991) and Lys and Sohn (1990) also use individual analyst data from the Value Line Investment Survey and Zacks Investment Research, respectively.

acteristics and earnings profiles markedly different from firms in other sectors. O'Brien (1994) finds that macroeconomic shocks to industrial production, interest rates, and stock returns help to explain industry-average forecast errors. Furthermore, the author finds that the sensitivity of forecast errors to shocks in interest rates and stock returns are industry-specific, while the sensitivity of forecast errors to industrial production growth appears constant and positive economy-wide.

We argue the current literature on forecast revisions neglects these crucial macroeconomic elements. We breakdown the effects of firm-specific characteristics like stock price movements from the macroeconomic information available to analysts at the time of the forecast revisions, in order to cope with a potential omitted variable bias. Finally, we augment the current literature by examining these effects for the economy as a whole, and for specific industries.

3.3 Data

We are interested in evaluating financial analysts' use of firm-specific versus systematic macroeconomic information in updating earnings forecasts across different industries. For this purpose, we obtain average analyst forecasting data from the summary tapes of the Institutional Brokers' Estimate System (I/B/E/S), daily stock price and return data from the Center for Research in Security Prices (CRSP), and, following Ludvigson and Ng (2006), we use the key macroeconomic data compiled by Stock and Watson (2005) from the Global Insights Basic Economics database.

3.3.1 Forecast Revisions Data

The data on financial analysts' forecast revisions, by firm, are collected from the Institutional Brokers' Estimate System (I/B/E/S) provided by the company I/B/E/S International Inc. I/B/E/S gathers and compiles consensus and individual forecasts from security analysts at the firm level, including earnings per share, revenue, cash flow, long-term growth projections, and stock recommendations. The main variable of interest, for the purpose of this study, is financial analysts' earnings per share forecasts for a given fiscal year.

Every third Thursday of the month, I/B/E/S computes several statistical measures for the distribution of forecasts provided by financial analysts working in different sectors of the economy. Following I/B/E/S, we refer to each third Thursday as a statistical date. We hereafter denote each third Thursday as statistical date s.

Like Klein (1990) and Chaney, Hogan, and Jeter (1999), we use the consensus forecast from I/B/E/S, as individual analysts often do not provide a complete series of forecast updates. Moreover, even if individual analyst forecast predictions were available for the full time period, the series may be subject to considerable error. The consensus mean forecast is hereafter referred to as the forecast of the "average analyst".

We represent the average analyst's forecast for firm j for the fiscal year ending in $T+1$, entered on statistical date s in the fiscal year $T+1$ by $f^{j,T+1}_{s,T+1}$. We refer to the stock price for firm j on statistical date s of the fiscal year $T+1$ to be $p^{j}_{s,T+1}$. At the end of the fiscal year $T+1$, firm j's earnings per share results are announced. We represent firm j's realized earnings per share by $y^{j,T+1}$.

For the fiscal year $T+1$, we compute the average analyst's forecast revision of firm j on statistical date $s+1$ as follows:[2]

$$rev^{j,T+1}_{s+1,T+1} = \frac{f^{j,T+1}_{s+1,T+1} - f^{j,T+1}_{s,T+1}}{p^{j}_{s,T+1}}. \tag{3.1}$$

That is, we define the average analyst's forecast revision at statistical date $s+1$ to be the change in the average analyst's forecast for fiscal year $T+1$ between statistical periods s and $s+1$. The difference in the forecasts is then scaled by the stock price at statistical date s. In this way, we can interpret each revision as a percent value of the stock price. Similarly, we compute the revision at statistical date $s+2$ by the change between the average prediction at $s+2$ and $s+1$, scaled by the stock price in $s+1$. Figure 3.1 demonstrates our construction of the average analyst's forecast revisions over time.

We note one small caveat to the construction of average forecast revisions pictured in Figure 3.1. As analysts continually update forecasts, it may be that we observe statis-

[2] As statistical dates occur with higher frequency (monthly) than fiscal years, a one-period ahead statistical date ($s+1$) need not occur in the same time period as a one-period ahead fiscal year ($T+1$).

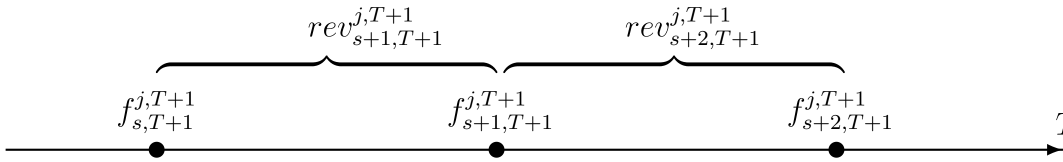

Figure 3.1: Construction of Forecast Revisions

tical dates with multiple forecast revision observations if analysts provide forecasts for different fiscal year targets. For instance, at statistical date s, we may observe forecast revisions for the fiscal year ending in T and the fiscal year ending in $T+1$. As in Figure 3.1, the general case for revisions of forecasts for fiscal year $T+1$ are entered during fiscal year $T+1$. However, to create a contiguous time series of forecast revisions, the first revision of fiscal year $T+1$ depends on a forecast for fiscal year $T+1$ made prior to the end of fiscal year T.

Figure 3.2 demonstrates our treatment of forecast revisions at the beginning of each fiscal year in order to achieve a continuum of revisions over the sample period. The forecast revision at statistical date $s+2$ corresponds to the fiscal year ending in $T+1$. However, to maintain a sequence of revisions over time, we use the forecast for fiscal year $T+1$ issued before the end of fiscal year T. Please note that the forecast revision for statistical date $s+1$ corresponding to the fiscal year ending in year T is defined exactly as pictured in Figure 3.1.

Our sample consists of 5,004 firms with consensus forecast revision data between January 1986 and December 2003. From this sample, we exclude firms with a discontinuity in the sequence of forecast revisions. Finally, we only include firms with a minimum of 60 observations, that is, a minimum of 60 continuous statistical periods of forecast revisions. The final sample consists of 2,856 firms, representing all sectors of the economy.

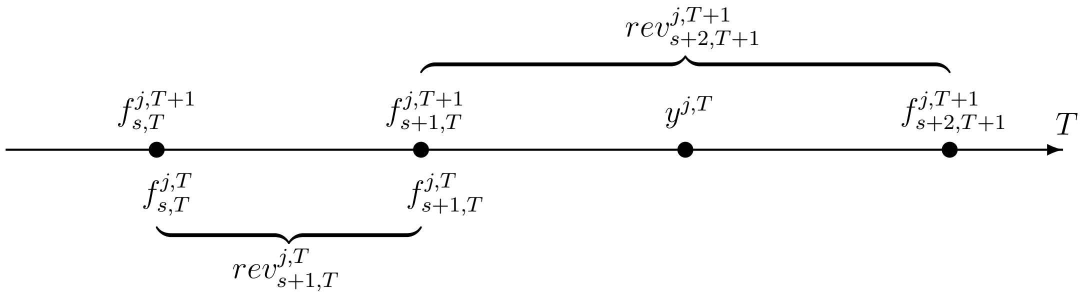

Figure 3.2: Treatment of the First Revision

3.3.2 Stock Price Data

The stock price and stock return data are collected from the Center for Research in Security Prices (CRSP). We use the daily stock price information in CRSP in order to scale the average analyst's change of opinion as in (3.1). We also make use of the CRSP data on each firm's excess return over New York Stock Exchange (NYSE) Equally-Weighted index, in order to compute the cumulative past excess return for each firm.

We match the daily stock price data from CRSP by statistical date to our revision database. Using the matched data, we compute the average analyst's forecast revisions by firm for each statistical date, following the definition in (3.1). We also use the daily excess stock return information to generate a proxy for firm-specific information that may affect an analyst's prediction update. As is suggested in Section 3.2.1, if analysts believe that traders are better informed about a firm's "true" profit possibilities, they may use the information contained in stock returns when they update earnings per share predictions.

More formally, we define the excess stock return on day t for firm j to be r^j_t, and the cumulative excess return for firm j between statistical dates s and $s+1$ is defined as follows:

$$\widetilde{r}^j_{s+1} = \sum_{s<t\leq s+1} r^j_t. \tag{3.2}$$

Along with daily stock price data, another advantage of the CRSP data is the availability of a firm's industry classification. We expect the relationship between a firm's earnings potential and macroeconomic conditions to vary across segments of the economy due to the variation in firms' real operating characteristics. For this reason, we classify firms into industry groups following the Fama-French industry classification. Fama and French categorize firms into 12 broad industry groups based on their Standard Industrial Classification (SIC) codes, available in CRSP.[3] The 12 industries are as follows: Consumer Durables; Consumer Nondurables; Manufacturing; Oil, Gas, and Coal Extraction and Products (Energy); Chemicals and Allied Products; Business Equipment; Telephone and Television Transmission; Utilities; Wholesale, Retail, and Some Services (Shops); Health Care, Medical Equipment, and Drugs (Health); Finance and Insurance Companies (Money); and Other including mines, construction, transportation, hotels, business services, and entertainment.[4]

3.3.3 Macroeconomic Data

Our chapter expands on the current literature to include macroeconomic variables in explaining analyst forecast revisions. We follow Ludvigson and Ng (2006) and use the macroeconomic data compiled by Stock and Watson (2005) from the Global Insights Basic Economics database.

The panel of 132 macro variables are sorted into three broad groups: 1) economic activity, 2) financial conditions, and 3) inflationary measures. The first group consists of 71 variables that track economic activity including different real output and income measures, employment and hours, real retail, manufacturing and trade sales, consumer spending, housing starts, inventories and inventory sales ratios, orders and unfilled orders. The second group contains 37 variables that capture financial conditions, such as monetary aggregates and credit measures, interest rates and interest rate spreads, stock market indicators, and foreign exchange measures. Finally, the third group consists of 24 inflation measures based on compensation and labor costs, capacity utilization measures, and price indexes. The data are monthly and span the period from January 1960 to December 2003.

[3]Our industry classification is based on each firm's most recent SIC code.

[4]Please see the data library at Kenneth French's website for the SIC codes in each broad industry group.

3.4 Descriptive Statistics

Before we move to the econometric analysis designed to test whether macroeconomic factors hold explanatory power for financial analysts' earnings forecast revisions, we offer a brief examination of the various databases used for our tests. Our data are shown to exhibit the stylized characteristics of the earnings forecast revisions' literature. Furthermore, we go beyond the previous work to illustrate the major differences across industries.

3.4.1 Forecast Revisions

For each of the 2,856 firms in our sample, we generate earnings forecast revisions for the average analyst over time as defined in (3.1). Then, for each firm, we compute a single mean revision over the sample period. Figure 3.3 plots the distribution of these mean revisions across all firms in each of the 12 Fama-French industrial groups.

Across most sectors, the distribution of mean revisions are skewed to the left, demonstrating that our data are consistent with the literature's findings that negative revisions are relatively more frequent. The telecommunications industry is an exception, though it appears the positive skew may be a result of a select few positive outliers. In the business equipment and utilities sectors, we also observe significant negative outliers that may alter the distribution. The health services and consumer nondurables industries appear to best represent the standard distribution, centered around zero revisions.

For each statistical date, we compute the average revision across firms in a given industry. The resulting time series, displayed in Figure 3.4, demonstrates roughly the same patterns as in Figure 3.3. Across most industries, average forecast revisions are more often negative. Once again, we point to the consumer durables and health services sectors, where average revisions over time are generally centered around zero revisions.

We remark that the pattern of average forecast revisions is particular to the industry being considered. For instance, mean forecast revisions for the business equipment sector display significant volatility through time, as is evidenced by the broad range of values (from a positive revision of around 20 percent of the stock price to to a negative revision of over 80 percent of the stock price); though, we note the existence of a large

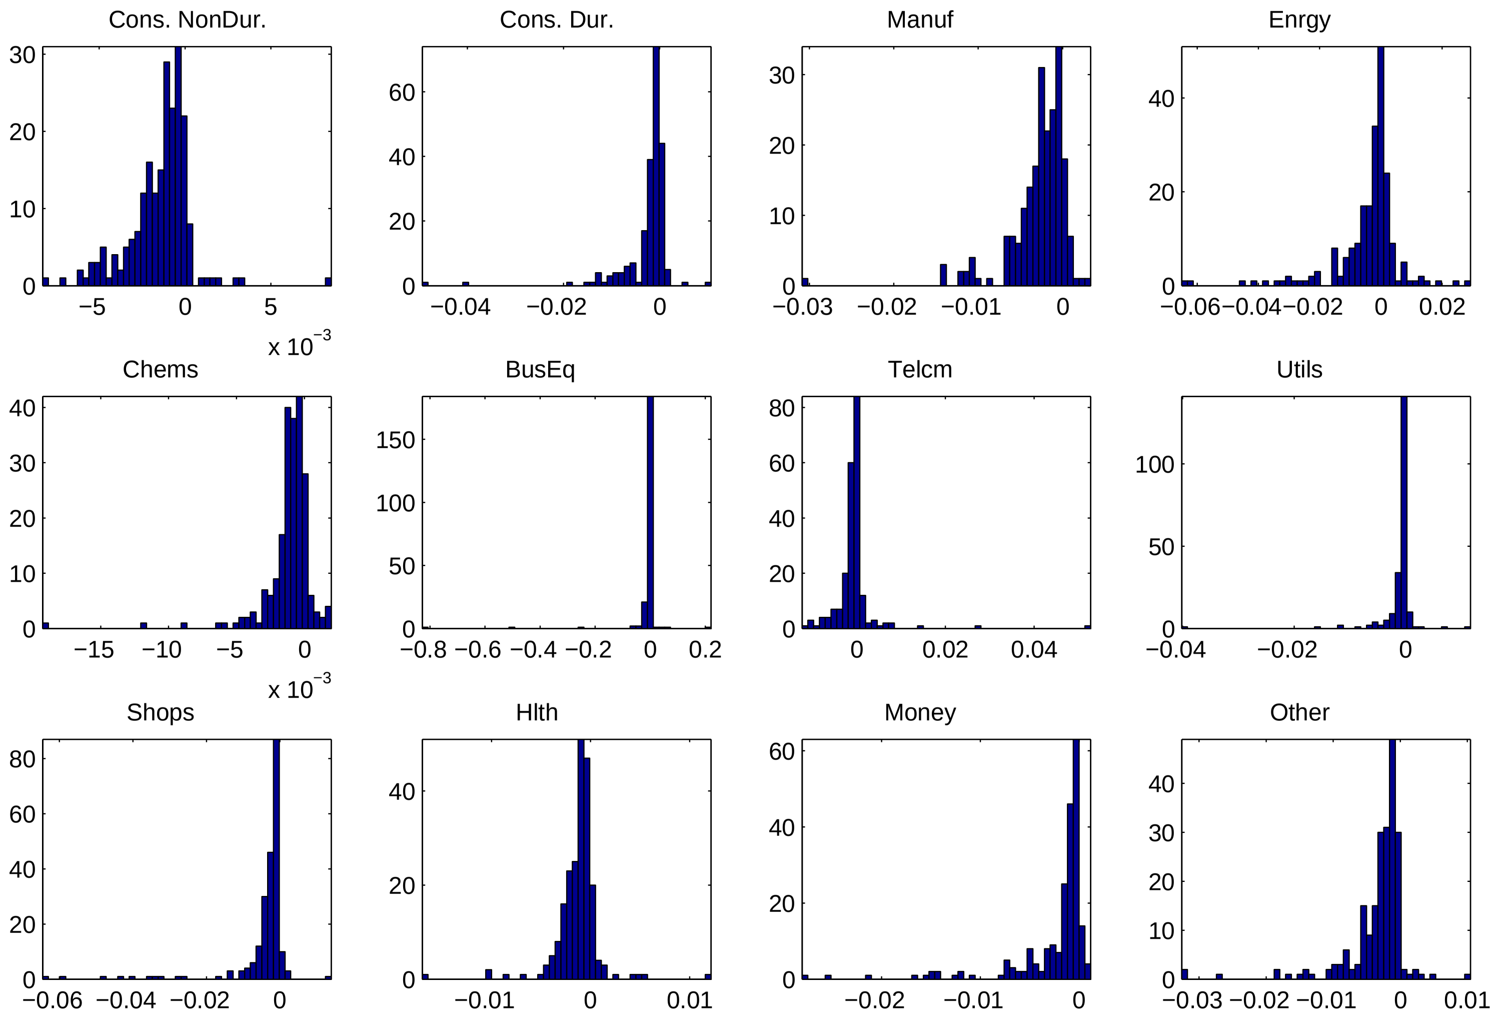

Figure 3.3: Distribution of Mean Revisions by Industry

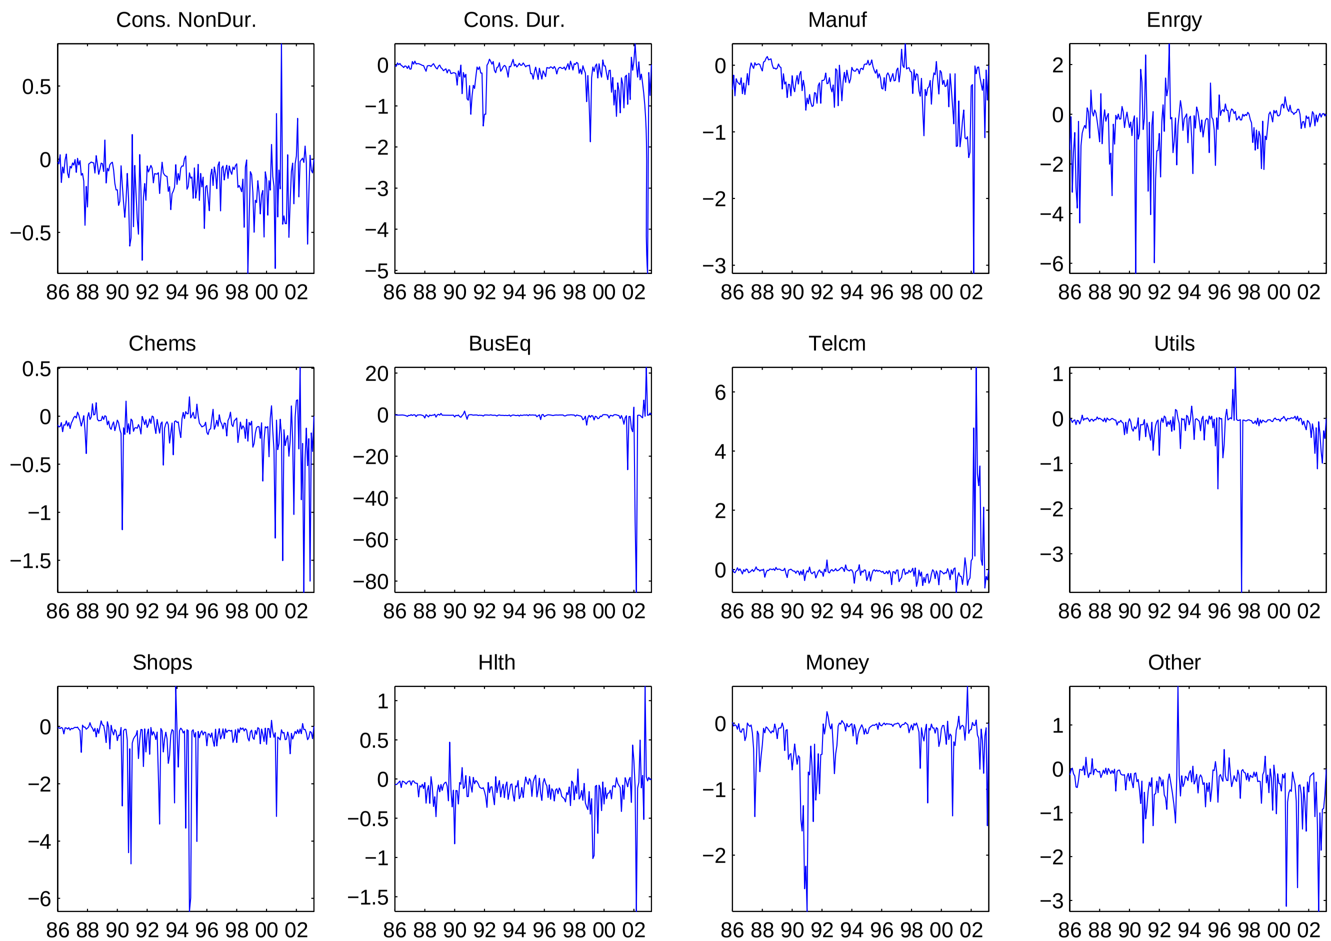

Figure 3.4: Time Series of the Mean Revisions by Industry

negative outlier. Average forecast revisions for the energy, shops, telecommunications, and consumer durables sectors also display a high level of volatility over time, while revisions for the consumer nondurables, financial services, and health sectors appear most stable over time. A likely reason for these patterns is that different industries react in different ways to short run economic fluctuations.

In Table 3.1, we provide descriptive statistics for the distributions of mean revisions by industry illustrated in Figure 3.3. First, we note that our data are consistent with findings in the literature (e.g., Abarbanell (1991) and Jain (1992)). The average revision over all firms in all industries in the sample is negative, at -.383 percent of the stock price (see column (5)). This suggests that average earnings forecasts start out too high and are lowered as the time to the earnings announcement date draws closer. This gives rise to a greater number of downward revisions than upward revisions. As is displayed in columns (2) and (4) of the table, across all industries, 34.1 percent of revisions are negative, while 23.1 percent of revisions are positive. The data also show a high percentage (42.8 percent) of zero revisions (see column (3)).[5]

We now turn our analysis to the differences across sectors. Within the 2,856 I/B/E/S firms in our data, there is large variation in the number of firms characterizing each of the 12 industry groups (see column (1)). There are, on average, 238 firms per industry; however, the consumer durables sector is represented by 60 firms, while the business equipment and financial services sectors offer the largest representation with 459 and 543 firms, respectively.

In terms of the distribution of the revisions by sector, the business equipment sector exhibits the largest (in absolute terms) mean revision, while the telecommunications sector is the only industry with a positive mean revision. Recall, however, our comments with respect to the distribution of firm mean revisions for the business equipment and telecommunications sectors in Figure 3.3. There appears to be significant outliers in the data that may alter these distributional characteristics, deserving special attention. We note further that the large standard deviation for the business equipment sector may be a result of outlier observations. Furthermore, the large degree of negative skewness for

[5]Aiolfi, Rodriguez, and Timmermann (2006) focus on firms in the Dow Jones 30 Index and propose a nonlinear econometric model in order to cope with the asymmetric behavior of positive and negative earnings forecast revisions.

the business equipment sector and the relatively small degree of negative skewness for the telecommunications sector reinforce this point.

With this in mind, we note a pattern in the data in sectors without significant outliers. The chemicals, utilities, health services, and consumer nondurables sectors are characterized by small mean revisions and low variance, while the retail, energy, consumer durables, and manufacturing sectors display a high variance across revisions and larger mean revisions. Taken together, these points suggest analysts' sensitivity to new information is not uniform across sectors. In fact, if macroeconomic information is of any importance we should expect analysts to rely more on macroeconomic environment factors for capital-intensive, highly regulated, highly leveraged or demand-driven industries, like the retail, energy, consumer durables, and manufacturing sectors.

The energy sector is characterized by the most revisions (both positive and negative), while analysts are least likely to revise forecasts for the health sector. One interpretation of these figures could rely on the relative capital-intensities of these industries. The energy sector is arguably more capital-intensive than the health services sector. Because of this, analysts may weight more heavily the conditions in the economic environment for the energy sector, while they may weight more heavily firm-specific characteristics for relatively less capital-intensive sectors like the health services sector. In fact, a casual glance at the distribution of zero revisions by sectors offers some support for this interpretation.

Table 3.1 also presents information on the average, minimum, and maximum number of analysts covering the average firm in the industry. First, we note that, on average, seven analysts provide revisions for their initial estimates across all firms in all industries. At least three and at most twelve analysts provide forecast revisions. The energy and utilities sectors have the highest average analyst coverage.

3.4.2 Stock Returns

We now turn our focus to the stock price data from CRSP. We calculate a firm-specific cumulative stock return and excess stock return as defined in (3.2). We then calculate the equally-weighted average return over all firms in an industry.

In Table 3.2, we report descriptive statistics for the equally-weighted monthly industry-

Table 3.1: Forecast Revisions

This table reports descriptive statistics for the average analyst's earnings forecast revisions across 2,856 I/B/E/S companies clustered by industry. The sample spans the period January 1986 to December 2003. For the fiscal year $T+1$, we compute the average analyst's forecast revision of firm j on statistical date $s+1$ as in (3.1). Mean and standard deviation statistics are reported as percentages. The last three columns of the table report the average, minimum, and maximum number of analysts covering the average firm in the industry.

	$\#firms$	$rev > 0$	$rev = 0$	$rev < 0$	Mean	Std	AR(1)	Skew	Kurt	Analysts Coverage		
										Mean	Min	Max
Cons. NonDur.	154	21.5%	45.7%	32.8%	-0.195	0.924	0.177	-2.65	28.94	9	3	13
Cons. Dur.	60	25.3%	39.3%	35.4%	-0.345	1.645	0.221	-3.37	29.83	8	3	12
Manuf	301	22.0%	39.7%	38.3%	-0.325	1.359	0.194	-3.34	28.18	8	3	12
Enrgy	120	30.4%	25.0%	44.7%	-0.397	2.014	0.268	-1.73	20.55	12	5	19
Chems	68	21.3%	42.1%	36.5%	-0.141	0.601	0.226	-2.86	30.55	10	4	15
BusEq	459	21.7%	47.9%	30.4%	-1.665	9.197	0.163	-3.54	32.42	9	3	15
Telcm	89	24.2%	40.3%	35.5%	0.027	2.017	0.113	-1.54	31.39	10	4	16
Utils	131	24.0%	39.9%	36.1%	-0.170	0.910	0.243	-3.19	34.08	13	6	19
Shops	324	21.2%	49.3%	29.5%	-0.520	2.069	0.173	-3.53	29.89	8	3	13
Hlth	246	19.9%	51.6%	28.5%	-0.176	1.220	0.136	-2.83	29.73	8	3	14
Money	543	23.9%	47.4%	28.7%	-0.286	1.514	0.163	-3.31	34.27	8	4	12
Other	361	21.6%	45.1%	33.3%	-0.400	2.125	0.163	-2.79	26.43	7	3	12
Average	238	23.1%	42.8%	34.1%	-0.383	2.133	0.187	-2.89	29.69	7	3	12

average returns and the equally-weighted monthly industry-average excess stock returns over the NYSE Equally-Weighted index. We note that the average returns over this period are positive and unusually high when compared to returns computed from the Fama-French database.[6] A reason for this finding may be the sample selection bias involved in our study. By selecting firms with reasonable coverage over the sample period, we place more weight on firms with higher returns since, as pointed by McNichols and O'Brien (1997), analysts expend less effort in their coverage of bad news stocks.

Firms in the business equipment and health services sectors display the largest mean stock returns, as well as the largest mean excess stock returns, over the sample period. Meanwhile, firms in the utilities sector realized the largest losses over the sample period when compared to the equally-weighted NYSE portfolio return. The volatility of excess returns over this sample period appears to be highly industry-specific. The business equipment and energy sectors report the greatest heterogeneity of returns, while excess returns for the consumer nondurables sector present relatively stable behavior.

Table 3.2: Equally-Weighted Monthly Industry Returns

This table reports summary statistics for equally-weighted monthly industry returns from statistical date s to $s + 1$ over the sample January 1986 to December 2003. "Return" refers to holding-period returns computed as in (3.2) and "Excess Return" refers to the excess returns with respect to equally-weighted NYSE stocks.

	Return				Excess Return			
	Mean	Std	Min	Max	Mean	Std	Min	Max
Cons. NonDur.	1.45%	4.8%	-23.2%	17.9%	0.22%	1.9%	-5.2%	6.4%
Cons. Dur.	1.44%	6.8%	-32.6%	21.6%	0.15%	3.1%	-14.2%	11.0%
Manuf	1.44%	6.1%	-26.4%	19.9%	0.23%	2.3%	-7.0%	6.6%
Enrgy	1.46%	6.6%	-25.1%	19.3%	0.24%	5.1%	-17.5%	17.4%
Chems	1.36%	5.0%	-24.4%	15.1%	0.16%	2.2%	-7.3%	7.1%
BusEq	2.40%	8.7%	-32.0%	32.6%	1.19%	5.8%	-19.8%	20.9%
Telcm	1.72%	6.5%	-26.7%	20.7%	0.44%	4.4%	-15.0%	14.4%
Utils	1.09%	3.6%	-14.3%	9.6%	-0.11%	4.0%	-15.1%	15.8%
Shops	1.75%	6.0%	-25.3%	19.5%	0.52%	2.5%	-6.3%	9.0%
Hlth	2.19%	7.0%	-26.2%	27.4%	0.99%	4.4%	-13.4%	21.7%
Money	1.67%	4.9%	-20.6%	23.9%	0.46%	2.2%	-8.1%	6.0%
Other	1.65%	5.7%	-26.6%	18.1%	0.44%	2.0%	-7.4%	6.0%

[6] These returns are not presented here, but are available from the authors by request.

3.5 Econometric Framework

Our objective is to describe the dynamics of the average analyst's earnings forecast revisions for a given firm j. In order to avoid cumbersome notation we define revisions, rev_{s+1}^{j}, by omitting the fiscal year index as follows:

$$rev_{s+1}^{j} = rev_{s+1,T+1}^{j,T+1}.$$

In this chapter, we go beyond the current literature to represent forecast revisions in terms of firm-specific information, X_s^j, and economic environment information, Z_s. The standard econometric approach is to select a set of pre-determined conditioning variables corresponding to statistical date s and then to estimate

$$rev_{s+1}^{j} = c^j + \alpha^j X_s^j + \beta^j Z_s + \epsilon_{s+1}^{j} \tag{3.3}$$

by ordinary least squares.

In a market with asymmetric information, changes in firms' stock prices may signal private information about a firm's operating environment not available to individual analysts. Therefore, we follow the literature and adopt the cumulative past excess return for each firm as a proxy for firm-specific information in X_s^j. To be precise, the cumulative excess return for firm j between statistical dates $s-1$ and s, $\widetilde{r}_s^j$, is defined as in (3.2).

In order to define the vector Z_s of macroeconomic information, we follow Ludvigson and Ng (2006) and use a panel of 132 macroeconomic variables from Stock and Watson (2005) that represent three broad categories: 1) economic activity, 2) financial conditions, and 3) inflationary measures. However, each category contains too many individual series for our analysis; that is, the number of predictor variables is large relative to the time series dimension. Any estimation of (3.3) that makes use of individual variables would suffer from a significant reduction in the degrees of freedom. Stock and Watson (2002a, 2002b, 2004) show that predictions of real economic activity are greatly improved relative to low-dimensional forecasting regressions when forecasts are based on the estimated factors of large data sets.

3.5.1 Common Factors: Model Specification

The common factors methodology is fundamentally motivated by economic theory. Information about unobservable shocks (e.g., aggregate demand or supply shocks) can often be extracted from a set of economic variables driven by the exogenous shocks. Moreover, our analysis of financial analyst forecast behavior relies on monthly revisions data, but monthly data for macroeconomic indicators are rare. This approach proves to be quite useful in situations with scarce data. Even when a full monthly time series is available, we may expect considerable measurement error. With a large, representative cross-section of information, and provided that measurement errors are mostly idiosyncratic, we can construct common factors that are robust to measurement error from the cross-sectional variables.

Stationarity of the economic variables is an important condition for the common factors analysis, so the series are transformed by taking logarithms and/or first differencing. We first normalize each series to have zero mean and unit variance. We then stack the variables related to each group i into a vector Z_s^i. With the assumption that the vector, Z_s^i, of time series observations over N^i economic variables for the time period $s = 1, ..., S$ exhibits a common factor, it can be represented as

$$Z_s^i = \lambda^i f_s^i + e_s^i \tag{3.4}$$

where f_s^i is the $r \times 1$ vector of latent common factors, λ^i is the corresponding $N^i \times r$ vector of latent factor loadings, and e_s^i is a $N^i \times 1$ vector of idiosyncratic errors. The researcher generally does not observe the vector of factors, f_s^i. In order to extract the factors from the observable economic variables, Z_s^i, we must make a key identifying assumption. As is typical in the literature, we assume that the errors, e_s^i, are mutually orthogonal with respect to f_s^i, although we allow that they are correlated across series and through time.

Under the assumed orthogonality between the common factors and the idiosyncratic disturbances, we can consider a decomposition of the covariance matrix of Z_s^i. The common component can be approximated by projecting on the first r static principal components of the covariance matrix. In Stock and Watson (2002a), the principal component estimators of the factors emerge as the solution to the following least squares problem:

$$\min_{f_s^i, \lambda^i} S^{-1} \sum_{s=1}^{S} \left(Z_s^i - \lambda^i f_s^i\right)' \left(Z_s^i - \lambda^i f_s^i\right)$$

subject to the restriction $\lambda^{i\prime}\lambda^i = I$. The resulting estimator of the factors, $\widehat{f_s^i}$, is the first r static principal components of Z_s^i.

We point out two additional benefits to this type of analysis. First, by summarizing the information from a large number of series in a few estimated factors, we eliminate the arbitrary reliance on a small number of imperfectly measured indicators to proxy for macroeconomic fundamentals, and make feasible the use of a vast set of economic variables that are more likely to span the unobservable information sets of financial market participants. Also, Stock and Watson (2002a) provide both theoretical arguments and empirical evidence that the principal components factor estimates are consistent even in the face of temporal instability in the individual time series used to construct the factors. The reason is that such instabilities may average out in the construction of common factors if the instability is sufficiently dissimilar from one series to the next.

The macroeconomic data from Stock and Watson (2005) used for this analysis spans the time period January 1960 to December 2003, while our data on the average analyst's forecast revisions covers only the period January 1986 to December 2003. With this in mind, we face a decision about the length of the macroeconomic time series to include in constructing the common factors. We opt to use the full sample for each of the macroeconomic time series in estimating the common factors, as this approach can be thought of as providing smoothed estimates of the latent factors, as shown in Brandt and Kang (2004).

3.5.2 Common Factors: Descriptive Statistics

We collect the common factor estimated from each group into the vector

$$\widehat{F}_s = (\widehat{f}_s^{ACT}, \widehat{f}_s^{FIN}, \widehat{f}_s^{INFL}).$$

Figure 3.5 plots the time series for each of the three estimated factors: economic ac-

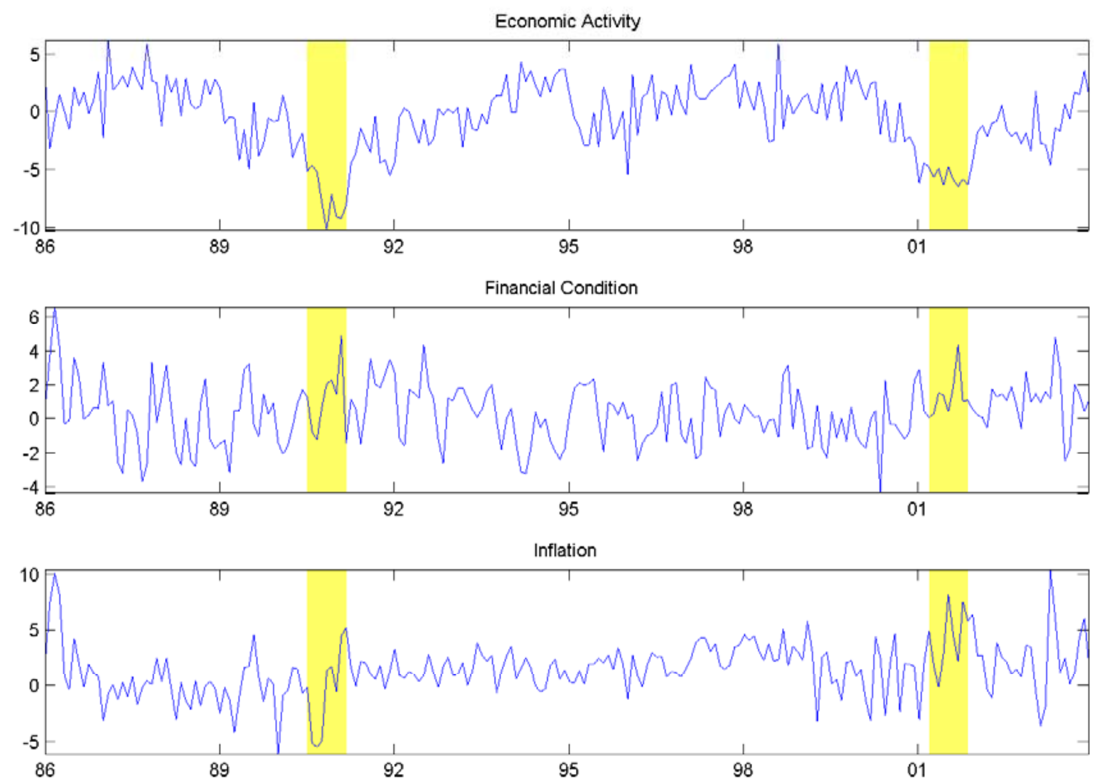

Figure 3.5: Plot of Estimated Macroeconomic Factors and NBER Recessionary Periods

tivity ($\widehat{f}_s^{ACT}$), financial conditions ($\widehat{f}_s^{FIN}$), and inflationary measures ($\widehat{f}_s^{INFL}$) over the full sample period. It is clear the estimated factor for economic activity is a good representation of the underlying variables, as the shaded areas corresponding to recessionary periods defined by the Business Cycle Dating Committee of the National Bureau of Economic Research (NBER), fall in line with downturns in the factor. The factors estimated from the financial conditions and inflation variables are considerably more volatile than the economic activity factor. We argue this reflects the more volatile nature of the underlying variables in these categories.

Table 3.3 reports descriptive statistics for the estimated macroeconomic factors. Panel A reports the first order autocorrelation, standard deviation, and percentage of the variance of the panel explained by each factor. We see that the common factor extracted from the inflation variables is the most highly persistent and explains roughly 46 percent of the variation in the data while the factor extracted from the financial measures is the least persistent and explains roughly 25 percent of the variation in the data.

Panel B reports the correlation coefficients among the factors. We note the negative correlation between the common factors for economic activity and financial conditions, possibly due to the interest rate component embedded within the set of financial conditions variables. We also point out that the inflationary measure common factor is weakly correlated with the economic activity factor. We hypothesize that this poor correlation is due to the well-known instability of the Phillips curve. Since interest rates are part of the financial variables, the negative correlation between inflation and financial common factors is no surprise by the Fisher equation.

3.5.3 Estimation with Common Factors

Returning to the beginning of the section, we now have all of the appropriate elements to estimate the regression in (3.3). Bai and Ng (2006) show that the ordinary least squares estimates from factor-augmented forecasting regressions are $\sqrt{T}$ consistent and asymptotically normal, and that pre-estimation of the factors does not affect the consistency of the second-stage parameter estimates. We substitute the estimated common factors, $\widehat{F}_s$, for the full set of economic environment variables, Z_s, in order to reduce the dimensionality of the problem. Using ordinary least squares with White (1980) robust

Table 3.3: Estimated Macroeconomic Factors

This table reports descriptive statistics for common factors estimated over the sample period from January 1960 to December 2003. The factors include: economic activity ($\widehat{f}_s^{ACT}$), financial conditions ($\widehat{f}_s^{FIN}$), and inflationary measures ($\widehat{f}_s^{INFL}$). Panel A reports the first-order autocorrelation, standard deviation, and percentage of the variance of the panel explained by each factor. Panel B reports the correlation coefficients among the factors.

Panel A

	AR(1)	Std	R^2
$\widehat{f}_s^{ACT}$	0.740	4.595	0.302
$\widehat{f}_s^{FIN}$	0.435	2.806	0.254
$\widehat{f}_s^{INFL}$	0.759	3.344	0.466

Panel B

	$\widehat{f}_s^{ACT}$	$\widehat{f}_s^{FIN}$	$\widehat{f}_s^{INFL}$
$\widehat{f}_s^{ACT}$	1.000	-0.260	-0.044
$\widehat{f}_s^{FIN}$	-0.260	1.000	-0.281
$\widehat{f}_s^{INFL}$	-0.044	-0.281	1.000

standard errors, we estimate

$$rev^j_{s+1} = c^j + \alpha^j \widetilde{r}^j_s + \beta^j \widehat{F}_s + \epsilon^j_{s+1} \tag{3.5}$$

where c^j is a firm-specific constant, $\widetilde{r}^j_s$ is the cumulative past return for each firm, $\widehat{F}_s$ is the vector of estimated common factors, and ϵ^j_{s+1} is a vector of white-noise error terms.

3.6 Empirical Results

The objective of this chapter is to investigate how financial analysts incorporate macroeconomic fundamentals into their earnings per share forecast revisions for different industries. We first present results on the direction and sensitivity of the average analyst's revisions to each explanatory variable. Next, we demonstrate that the model that best explains the average analyst's forecast revisions, indeed, incorporates macroeconomic factors. In the analysis that follows, we restrict our presentation of the results to the case of excess returns as firm-specific information.[7]

3.6.1 Estimation Results

For each of the 2,856 firms in our sample, we estimate three versions of the regression in (3.5). First, we estimate the model common in the literature, which includes only firm-specific excess returns as a determinant of analysts' forecast revisions. Next, we estimate a regression using our three estimated macroeconomic factors. Finally, we combine the analyses and estimate the model in (3.5), which includes both excess returns and the macroeconomic factors.

Table 3.4 reports for each of the three models the percentage of cases, computed across all firms within an industry, for which the coefficients, α^j and β^j, in

$$rev^j_{s+1} = c^j + \alpha^j \widetilde{r}^j_s + \beta^j \widehat{F}_s + \epsilon^j_{s+1}$$

are significantly positive or negative at 10 percent level. The coefficients were computed

[7]Our estimations suggest that using stock returns as the firm-specific component leads to the same set of conclusions.

with ordinary least squares estimation using White (1980) robust standard errors.

To further support our hypothesis that macroeconomic indicators have additional explanatory power in determining financial analysts' earnings forecast revisions when considering firms' excess stock returns, Table 3.4 also reports the average adjusted R^2, $\overline{R}^2$, from (3.5) across all firms in each industry.

Focusing on the first row of results for each industry, we note that analysts' revisions are more often positively correlated with a firm's past cumulative excess stock return in all industries we consider, consistent with the current literature on analyst forecast revisions. In fact, the coefficient on firms' excess stock returns are very rarely negative and significant across all industries. In 48.3 percent of the cases, the α coefficient on the industry-average excess stock return for the energy sector is a positive and significant determinant of financial analysts' forecast revisions. Excess stock returns in the telecommunications and financial services sectors are positively and significantly associated with the average analyst's revision far less frequently than in other sectors.

For the model with the three macroeconomic factors as regressors (second row for each industry), there is little heterogeneity among industries with respect to the direction of the parameters of interest. Concerning the economic activity measure, we observe, across all industries, that this factor correlates positively with the average analyst's revision more often than it is negatively correlated. The energy and health services sectors are the rare cases for which the factor representing economic activity is more often negatively correlated with financial analysts' revisions. In the energy sector, the coefficient on the economic activity factor is negative and significant 9.2 percent of the time, while it is positive and significant only 8.3 percent of the time. Similarly, in the health services sector, the coefficient on the economic activity factor is negative and significant 16.3 percent of the time, while it is positive and significant only 9.3 percent of the time. On the contrary, in the highly demand-driven sector of consumer durables, the coefficient on the economic activity factor is positive and significant 53.3 percent of the time.

The financial condition measure correlates negatively with analysts' revisions more often than positively in all industries, most notably, in the energy and retail sectors. As these sectors tend to be highly leveraged, it is understandable that analysts are highly sensitive to interest rate fluctuations when updating predictions for these sectors. Finally,

Table 3.4: Summary of Regression Coefficients

This table reports the percentage of cases in which the coefficients, α^j and β^j, in the OLS regression $rev^j_{s+1} = c^j + \alpha^j \widetilde{r}^j_s + \beta^j \widehat{F}_s + \epsilon^j_{s+1}$ are significantly positive or negative at 10 percent level, computed across all firms within an industry. $\widetilde{r}^j_s$ is the holding-period excess return with respect to the equally-weighted CRSP index for each firm, $\widehat{f}^{ACT}_s$ is the first factor extracted from economic activity variables, $\widehat{f}^{FIN}_s$ is the first factor extracted from the financial variables, and $\widehat{f}^{INFL}_s$ is the first factor extracted from different inflation measures. $\overline{R}^2$ is the adjusted-R^2 of the regressions. t-statistics are computed using White (1980) robust standard errors.

	$\widetilde{r}^j_s$		$\widehat{f}^{ACT}_s$		$\widehat{f}^{FIN}_s$		$\widehat{f}^{INFL}_s$		$\overline{R}^2$
	$\alpha > 0$	$\alpha < 0$	$\beta > 0$	$\beta < 0$	$\beta > 0$	$\beta < 0$	$\beta > 0$	$\beta < 0$	
Cons. NonDur.	46.1%	0.0%	-	-	-	-	-	-	0.03
	-	-	22.1%	7.8%	3.2%	5.8%	7.1%	8.4%	0.02
	46.8%	0.0%	22.1%	7.1%	1.9%	5.8%	6.5%	8.4%	0.04
Cons. Dur.	38.3%	0.0%	-	-	-	-	-	-	0.02
	-	-	53.3%	1.7%	3.3%	6.7%	10.0%	5.0%	0.05
	38.3%	1.7%	56.7%	1.7%	3.3%	6.7%	10.0%	3.3%	0.07
Manuf	45.2%	0.0%	-	-	-	-	-	-	0.02
	-	-	38.5%	2.7%	4.3%	8.3%	6.6%	15.0%	0.03
	45.2%	0.0%	39.2%	3.0%	5.3%	7.0%	5.0%	13.6%	0.05
Enrgy	48.3%	0.0%	-	-	-	-	-	-	0.03
	-	-	8.3%	9.2%	2.5%	18.3%	1.7%	36.7%	0.04
	42.5%	0.0%	9.2%	8.3%	1.7%	15.8%	1.7%	33.3%	0.07
Chems	38.2%	0.0%	-	-	-	-	-	-	0.03
	-	-	33.8%	5.9%	2.9%	2.9%	10.3%	13.2%	0.03
	41.2%	1.5%	36.8%	5.9%	2.9%	4.4%	10.3%	10.3%	0.05
BusEq	40.1%	0.7%	-	-	-	-	-	-	0.03
	-	-	25.9%	4.6%	3.1%	6.5%	6.1%	9.4%	0.02
	39.0%	0.4%	23.3%	4.4%	2.8%	7.0%	5.4%	9.4%	0.04

Table 3.4 (continued)

	$\widetilde{r}_s^j$		$\widehat{f}_s^{ACT}$		$\widehat{f}_s^{FIN}$		$\widehat{f}_s^{INFL}$		$\overline{R}^2$
	$\alpha > 0$	$\alpha < 0$	$\beta > 0$	$\beta < 0$	$\beta > 0$	$\beta < 0$	$\beta > 0$	$\beta < 0$	
Telcm	15.7%	0.0%	-	-	-	-	-	-	0.01
	-	-	22.5%	5.6%	6.7%	4.5%	5.6%	6.7%	0.02
	13.5%	0.0%	21.3%	5.6%	5.6%	4.5%	6.7%	5.6%	0.03
Utils	34.4%	0.8%	-	-	-	-	-	-	0.02
	-	-	25.2%	3.8%	3.1%	4.6%	6.1%	16.0%	0.02
	29.0%	0.8%	24.4%	3.8%	3.1%	9.2%	6.1%	16.0%	0.04
Shops	47.2%	0.3%	-	-	-	-	-	-	0.03
	-	-	25.3%	6.2%	3.7%	9.9%	6.5%	7.7%	0.02
	46.3%	0.6%	25.0%	6.2%	4.3%	7.4%	4.9%	9.0%	0.05
Hlth	31.7%	0.8%	-	-	-	-	-	-	0.02
	-	-	9.3%	16.3%	4.1%	4.9%	8.1%	6.1%	0.01
	31.3%	1.6%	8.1%	15.0%	4.9%	3.7%	8.9%	5.7%	0.03
Money	26.9%	0.6%	-	-	-	-	-	-	0.02
	-	-	21.2%	8.7%	4.4%	5.2%	11.4%	5.7%	0.02
	26.2%	0.4%	19.9%	7.0%	4.1%	5.7%	9.4%	5.7%	0.04
Other	39.1%	0.6%	-	-	-	-	-	-	0.03
	-	-	24.1%	6.9%	3.0%	7.2%	9.1%	6.9%	0.02
	37.7%	0.8%	25.2%	5.5%	2.8%	6.4%	7.2%	5.8%	0.05

the inflation factor appears to have no distinct pattern in the data across industries. In the energy sector 36.7 percent of the individual firm regressions reported a negative and significant effect while in the money and consumer durables the frequency of positive and significant coefficients seems to be the double of negative ones.

Finally, in the third row of each industry, we report the results of the regression in (3.5) including firm-specific excess returns and macroeconomic factors. For all industries, the $\overline{R}^2$ is increased from the model common in the literature including only firm excess returns. On average, the explanatory power is more than doubled when we estimate the full model. In trying to explain what drives financial analysts' forecast revisions, it is clear we should not exclude the factors representing the economic environment. Furthermore, we again point out the variability in the explanatory power of the macroeconomic factors across industries. The $\overline{R}^2$ increases the most for the consumer nondurables and energy sectors, the sectors for which the economic activity factor and inflation factor, respectively, are most often correlated with the financial analysts' revisions.

3.6.2 Model Selection

We provide further support for the argument that the model including firm-specific past excess returns and macroeconomic factors best explains the dynamics of financial analysts' forecast revisions. We begin by estimating the regression in (3.5), where X_s^j is the cumulative past excess return for each firm, $\widetilde{r}_s^j$, as defined in (3.2), designed to represent the firm-specific component and $\widehat{F}_s$ is the vector of estimated common factors. The vector includes three elements: the common factor for the economic activity variables ($\widehat{f}_s^{ACT}$), the common factor for the financial condition variables ($\widehat{f}_s^{FIN}$), and the common factor for the inflationary measures ($\widehat{f}_s^{INFL}$). For each 64 models, derived from the combination of the four available explanatory variables, we evaluate the corresponding Bayesian Information Criterion (BIC) and R^2, following the statistical model selection literature. Ultimately, the preferred set of regressors will minimize the BIC, as in Stock and Watson (2002b).

In Table 3.5, we present results from the model selection. Each column reports the percentage of cases for which each variable or factor was selected by BIC to be included in the regression analysis; that is, how often each variable or factor helps to explain analysts' prediction updates. We first note that across all firms in all industries in the sample, the firm-specific component proxied by firms' cumulative past excess returns holds explanatory power 45.4 percent of the time. This finding confirms results in the literature that stock returns are associated with analysts' earnings forecast revisions.

In addition, all three of the macroeconomic factors are selected as explanatory regressors alongside the firm-specific component. The economic activity factor helps to explain forecast revisions in 32.2 percent of cases. Though the financial condition and inflation factors are chosen to explain analysts' forecast revisions less frequently (14.1 percent and 22.1 percent of cases), we argue they, nevertheless, may represent a significant share of the variation of revisions. We consider this as evidence that the current literature, which excludes information on the macroeconomic environment as determinants of financial analysts' revisions, may suffer from an omitted variable bias.

As we shift focus to individual industries, we notice considerable heterogeneity in the model selection across the variables and factors. This is not surprising, as we expect the economic environment to have a larger impact on operating conditions for some sectors. Recall from Table 3.1 that some sectors experienced small mean revisions and a low variance of revisions, while other sectors displayed large mean revisions and a high variance of revisions. Based on these findings, we argued analysts' sensitivity to new information is not uniform across sectors. Table 3.5 supports this argument.

Though the factors extracted from the financial condition variables and inflationary measures are selected less often than firms' excess stock returns, the economic activity factor often explains revisions more frequently than returns. Most notably is the consumer durables sector, for which the economic activity factor has explanatory power more often than firm-level excess stock returns. Perhaps analysts following firms in this industry rely more on economic activity variables because of the sector's highly consumer demand-driven nature. Economic activity variables also explain revisions more often than returns for the telecommunications sector. Returns have explanatory power in a low 27.0 percent of cases for the telecommunications sector, while the economic activity and inflation factors have explanatory power in 28.1 and 31.5 percent of cases, respectively. In fact, the telecommunications sector is the only sector for which no sin-

Table 3.5: Model Selection

This table reports the average percentage of cases for which each variable or factor is included in the model selected by BIC, computed across firms in the same industry. The last row reports the average computed across all 2,856 firms. $\widetilde{r}_s^j$ is the holding-period excess return with respect to the equally-weighted CRSP index for each firm, $\widehat{f}_s^{ACT}$ is the first factor extracted from economic activity variables, $\widehat{f}_s^{FIN}$ is the first factor extracted from the financial variables, and $\widehat{f}_s^{INFL}$ is the first factor extracted from different inflation measures.

	$\widetilde{r}_s^j$	$\widehat{f}_s^{ACT}$	$\widehat{f}_s^{FIN}$	$\widehat{f}_s^{INFL}$
Cons. NonDur.	53.9%	28.6%	13.6%	14.3%
Cons. Dur.	40.0%	66.7%	13.3%	15.0%
Manuf	46.8%	40.5%	13.3%	19.9%
Enrgy	52.5%	15.8%	19.2%	37.5%
Chems	39.7%	38.2%	7.4%	25.0%
BusEq	47.7%	30.1%	13.1%	21.6%
Telcm	27.0%	28.1%	22.5%	31.5%
Utils	40.5%	31.3%	14.5%	21.4%
Shops	50.0%	31.8%	11.1%	19.8%
Hlth	48.4%	19.1%	18.7%	21.1%
Money	40.3%	36.5%	14.5%	21.7%
Other	45.2%	32.7%	12.7%	24.7%
Avg	45.4%	32.2%	14.1%	22.1%

gle factor or variable appears to significantly dominate the results. Perhaps this reflects financial analysts' insensitivity to new information.

On the contrary, returns are chosen to explain financial analysts' revisions almost four times as often as is the economic activity factor for the energy sector (52.5 percent for returns compared to 15.8 percent for economic activity). We also remark that the financial condition factor is selected in 37.5 percent of cases for this sector. Energy is the only sector for which financial conditions variables help to explain revisions more frequently than does the economic activity factor. We interpret this to reflect the sensitivity of interest rates and the stock market to worldwide oil markets. This result could also be a sign of the high debt-to-equity structure of the energy sector. Returns are also selected more frequently for the health services and consumer nondurables sectors. These results fall in line with our analysis about the sensitivity of revisions to new information in Section 3.4.1.

In all, these results offer some support for the hypothesis that the current literature focusing on the analysis of analyst forecast revisions based on firm-specific characteristics like stock price changes and excluding key macroeconomic fundamentals, indeed, have a potential omitted variable problem. In trying to explain what drives financial analysts' forecast revisions, it is clear we should not exclude the factors representing the economic environment. Moreover, the individual sensitivity displayed by each sector with respect to each macroeconomic class of variables could depend on the durability of the industry output. For example, it is well-known the demand for durable products, such as investment goods, is more responsive to interest rate changes through a cost-of-capital channel than is the demand for nondurable products, like food.

3.7 Conclusion

The main objective of this chapter is to address the issue of whether or not the average analyst incorporates more general economic performance measures when updating corporate earnings predictions. We contribute to the literature by adapting a framework that allows us to summarize groups of macroeconomic variables into a smaller set of factors, and we show that these macroeconomic factors help to explain the average analyst's opinion. Moreover, the explanatory power and direction of such factors strongly depend on the industry in question.

Analysts' forecasts often serve as a measure of market expectations and are consistently used in investment decisions. Given the importance of analysts' forecasts to market participants, analysts often update their forecasts as they process new information. This updating process may depend on firm-specific information such as past stock returns or more general economic conditions. However, since the set of variables that constitute the full set of economic conditions may be large, any study that makes use of the individual variables would suffer from a significant reduction in the degrees of freedom. We adapt an econometric methodology that helps us to investigate the impact of three large classes of macroeconomic variables: economic activity, financial conditions, and inflation. For each group, we extract a factor that summarizes the information content within that class. We then study how each factor performs in explaining the updating process of the average financial analyst in conjunction with firm-specific information such as past excess stock returns.

Our main results suggest that information about the economic environment does, indeed, help to explain the average analyst's forecast revision. Moreover, the magnitude and direction of the effect of macroeconomic factors depend on the industry considered. As firms in the same industry are likely to share common operating characteristics, these results have a strong intuitive appeal.

Given that measures of economic environment do help to explain the updating process of financial analysts, an interesting question is whether these macroeconomic variables can help to predict the revisions of the average analyst. Though this question is beyond the scope of this chapter, we believe it deserves consideration in future research.

This chapter is based on *Financial Analysts' Forecast Revisions and Macroeconomic Information* joint with Marco Aiolfi.

Appendix A

Assumptions and Proofs - Chapter 1

A.1 Assumptions

(A.1.1) Θ is a compact subset of $\mathbb{R}^h$ and $\theta^* \in \dot{\Theta}$.

(A.1.2) The h-vector Z_t (with first component 1) is such that, given $(\alpha, \theta) \in (0,1) \times \mathbb{R}_+$, $\forall \theta \in \dot{\Theta}$, $E[Z_t e^{-\rho w_{t+1}(\theta)}] \neq 0$ element by element, $E[Z_t 1_{\{e_{t+1}(\theta)<0\}} e^{-\rho w_{t+1}(\theta)}] \neq 0$, $E[Z_t Z_t']$ exists and is positive definite, and $E[\rho^2 (1_{\{e_{t+1}(\theta)<0\}} - \alpha)^2 - \rho] \neq 0$.

A.2 Proofs

Proof of Proposition 1.2.

Clearly the loss function satisfies $L(0) = 0$. We need to show that to the right of zero the function is nondecreasing and to the left of zero the function is nonincreasing. In order to do this, we invoke a result that states that if L has a right derivative at zero with $L'_+(0) > 0$, then $\exists \delta > 0$ such that $0 < x < \delta \Longrightarrow L(0) < L(x)$ and if $L'_-(0) < 0$ we have $L(X) > L(a) \forall x < a$ in a neighborhood of a Lima (1992, p. 208).

For the loss function defined above, in order to have

$$
\begin{aligned}
L'_+(0) &= \alpha U'(\cdot) > 0 \\
L'_-(0) &= -[1-\alpha] U'(\cdot) < 0
\end{aligned}
$$

we require $\alpha > 0$ in the first case and $\alpha < 1$ in the second.■

Proof of Proposition 1.3.

First we can write (1.5) as

$$L(e_{t+1}) = e^{-\rho\overline{w}}[e^{\rho\alpha|e_{t+1}|+\rho(1-2\alpha)|e_{t+1}|1_{\{e_{t+1}<0\}}} - 1]$$

1. Given the form above, the higher the value of $\overline{w}$ the lower will be the value of the loss incurred due to forecast errors $\therefore$ 1 follows.

2. Given the expression above, errors with different signs but same absolute magnitude are penalized equally only when α=0.5. Therefore if α $\neq$0.5 then the loss is asymmetric $\therefore$ we have 2.

The two arguments conclude the proof.■

Proof of Proposition 1.4.

From Christoffersen and Diebold (1997), we know that the optimal forecast $f^*_{t+1} = \mu_{t+1|t} + \lambda_{t+1|t}$, where $\lambda_{t+1|t}$ depends on the loss, $L(\cdot)$, $G_{t+1|t}$ and $\sigma_{t+1|t}$ but not on $\mu_{t+1|t}$. We know the optimal forecast should minimize the expected loss. The first order condition of the risk problem is given by

$$\int_{-\infty}^{\infty}[1_{\{\varepsilon_{t+1}<\frac{\lambda_{t+1|t}}{\sigma_{t+1|t}}\}}-\alpha]e^{\alpha\rho|\sigma_{t+1|t}\varepsilon_{t+1}-\lambda_{t+1|t}|+\rho(1-2\alpha)|\sigma_{t+1|t}\varepsilon_{t+1}-\lambda_{t+1|t}|1_{\{\varepsilon_{t+1}<\frac{\lambda_{t+1|t}}{\sigma_{t+1|t}}\}}}dG_{t+1|t} = 0$$

This first order condition is achieved by using the Dirac delta function. If $\lambda_{t+1|t} = 0$, then we have

$$[1-\alpha]\int_{-\infty}^{0} e^{(\rho-\alpha\rho)\sigma_{t+1|t}|\varepsilon_{t+1}|}dG_{t+1} = \alpha\int_{0}^{\infty} e^{\rho\alpha\sigma_{t+1|t}|\varepsilon_{t+1}|}dG_{t+1}$$

Given that we assumed the density of ε_{t+1} to be symmetric, this result is valid only if $\alpha = 0.5$. Therefore, by the counterpositive, if $\alpha \neq 0.5$ then the forecast should be biased, since $\lambda_{t+1|t} \neq 0$.■

Proof of Proposition 1.5.

If θ^* is the solution of the minimization problem

$$min_{\theta\in\Theta}E[e^{-\rho(-\alpha|Y_{t+1}-\theta'Z_t|-(1-2\alpha)|Y_{t+1}-\theta'Z_t|1_{\{Y_{t+1}-\theta'Z_t<0\}})}-1]=min_{\theta\in\Theta}\Sigma(\theta) \quad \text{(A.1)}$$

with $\Sigma(\theta)$ continuously differentiable on Θ, and $\theta^* \in \dot{\Theta}$, then θ^* satisfies the first order condition $\nabla_\theta\Sigma(\theta^*) = 0$. Let $\Sigma_{t+1}(\theta) \equiv [e^{-\rho(-\alpha|Y_{t+1}-\theta'Z_t|-(1-2\alpha)|Y_{t+1}-\theta'Z_t|1_{\{Y_{t+1}-\theta'Z_t<0\}})}-1]$. The function $\Sigma_{t+1}(\theta)$ is continuously differentiable on $\Theta \setminus A_{t+1}$ where $A_{t+1} = \{\theta \in \Theta : Y_{t+1} = \theta'Z_t\}$. Let $\nabla_\theta\Sigma_{t+1}(\theta)$ be the gradient of $\Sigma_{t+1}(\theta)$ on $\Theta \setminus A_{t+1}$. We have, by the law of iterated expectations,

$$\Sigma(\theta) = E[\Sigma_{t+1}(\theta)] = E[E_{t+1}[\Sigma_{t+1}(\theta)]],$$

so that

$$\begin{aligned}
\nabla_\theta\Sigma(\theta) &= E[\nabla_\theta E_{t+1}[\Sigma_{t+1}(\theta)]] \\
&= E[\nabla_\theta E_{t+1}[\Sigma_{t+1}(\theta)\cdot 1_{\{\theta\in A^c_{t+1}\}}]] + E[\nabla_\theta E_{t+1}[\Sigma_{t+1}(\theta)\cdot 1_{\{\theta\in A_{t+1}\}}]] \\
&= E[\nabla_\theta\Sigma_{t+1}(\theta)E_{t+1}[1_{\{\theta\in A^c_{t+1}\}}]] + E[\nabla_\theta\Sigma_{t+1}(\theta)E_{t+1}[1_{\{\theta\in A_{t+1}\}}]]
\end{aligned}$$

where $E_{t+1}[1_{\{\theta\in A^c_{t+1}\}}] = \mathcal{P}(A^c_\theta)$ with $A^c_\theta \equiv \Omega \setminus A_\theta$ and $A_\theta \equiv \{\omega \in \Omega : Y_{t+1}(\omega) = \theta'Z_t(\omega)\}$. Hence, $E_{t+1}[1_{\{\theta\in A^c_{t+1}\}}] = 1$ and $E_{t+1}[1_{\{\theta\in A^c_{t+1}\}}] = 0$. $\Sigma(\theta)$ is therefore continuously differentiable on Θ and we have

$$\begin{aligned}
\nabla_\theta\Sigma(\theta) &= \rho E[\nabla_\theta(\alpha|Y_{t+1}-\theta'Z_t| + (1-2\alpha)|Y_{t+1}-\theta'Z_t|1_{\{Y_{t+1}-\theta'Z_t<0\}}) \\
&\quad e^{-\rho(-\alpha|Y_{t+1}-\theta'Z_t|-(1-2\alpha)|Y_{t+1}-\theta'Z_t|1_{\{Y_{t+1}-\theta'Z_t<0\}})}] \\
&= \rho E[((1-2\alpha)|Y_{t+1}-\theta'Z_t|\nabla_\theta 1_{\{Y_{t+1}-\theta'Z_t<0\}} \\
&\quad +(\alpha+(1-2\alpha)1_{\{Y_{t+1}-\theta'Z_t<0\}})\nabla_\theta|Y_{t+1}-\theta'Z_t|) \\
&\quad e^{-\rho(-\alpha|Y_{t+1}-\theta'Z_t|-(1-2\alpha)|Y_{t+1}-\theta'Z_t|1_{\{Y_{t+1}-\theta'Z_t<0\}})}]
\end{aligned}$$

Note that

$$\nabla_\theta 1_{\{Y_{t+1}-\theta'Z_t<0\}} = Z_t \cdot \delta(\theta'Z_t - Y_{t+1})$$

where δ represents the Dirac delta function, i.e. for all $x \in \mathbb{R}^*$, $\delta(x) = 0$ and $\int_{\mathbb{R}} \delta(x)\, dx = 1$. Knowing that for any real value function $\varphi : \mathbb{R} \longrightarrow \mathbb{R}$ we have $\int_{\mathbb{R}} \delta(x)\varphi(x)\, dx = \varphi(0)$ we can find the optimal condition,

$$E[Z_t(1_{\{e_{t+1}(\theta^*)<0\}} - \alpha)e^{-\rho w_{t+1}(\theta^*)}] = 0$$

.

And this proves the Proposition 1.5.■

Proof of Proposition 1.6.

We derive the conditions for $\theta^* \in \Theta$ to be a solution of

$$min_{\theta\in\Theta} E[e^{-\rho w_{t+1}(\theta)} - 1] = min_{\theta\in\Theta}\Sigma(\theta)$$

where $w_{t+1}(\theta) = (-\alpha|Y_{t+1} - \theta' Z_t| - (1-2\alpha)|Y_{t+1} - \theta' Z_t| 1_{\{Y_{t+1}-\theta' Z_t<0\}})$ and $e_{t+1}(\theta) = Y_{t+1} - \theta' Z_t$.

To be a strict minimum of $\Sigma(\theta)$ on $\dot{\Theta}$ we need $\nabla_\theta \Sigma(\theta^*) = 0$ and $H_\theta(\Sigma)$ has to be positive definite, H being the hessian. Recall

$$\nabla_\theta \Sigma(\theta) = -\rho E[Z_t(1_{\{e_{t+1}(\theta)<0\}} - \alpha)e^{-\rho w_{t+1}(\theta)}]$$

By an argument similar to the previous proof:

$$\begin{aligned} H_\theta(\Sigma) &= -\rho E[Z_t Z_t' \delta(\theta' Z_t - Y_{t+1})e^{-\rho w_{t+1}(\theta)}] \\ &\quad +\rho^2 E[Z_t(1_{\{e_{t+1}(\theta)<0\}} - \alpha)\nabla_\theta w_{t+1}(\theta)e^{-\rho w_{t+1}(\theta)}] \\ &= \rho^2 E[Z_t Z_t'(1_{\{e_{t+1}(\theta)<0\}} - \alpha)^2 e^{-\rho w_{t+1}(\theta)}] \\ &\quad -\rho E[Z_t Z_t' \delta(\theta' Z_t - Y_{t+1})e^{-\rho w_{t+1}(\theta)}] \end{aligned} \qquad \text{(A.2)}$$

By assumptions (A0)-(A1), the quadratic form above is positive definite and we achieve a strict local minimum.■

Proof of Proposition 1.7.

The proof of this result follows the proof in Elliott, Komunjer, and Timmermann (2005b). The idea is to invoke the Implicit Function Theorem and try to derive the main conclusion. Define,

$$h_\rho(\alpha, \theta) = -\rho E[Z_t(1_{\{e_{t+1}(\theta)<0\}} - \alpha)e^{-\rho w_{t+1}(\theta)}]$$

so the first order condition is equivalent to $h_\rho(\alpha, \theta) = 0$.

In order to apply the implicit function theorem we need to show that (i) $h_\rho(\alpha, \theta) : (0, 1) \times \Theta \longrightarrow \mathbb{R}^k$ is continuously differentiable on $(0, 1) \times \Theta$ and (ii) $\forall \alpha \in (0, 1)$, the $\mathbb{R}^k \times \mathbb{R}^k$ matrix $\partial h_\rho(\alpha, \theta^*)/\partial\theta$ is nonsingular, i.e., $[\partial h_\rho(\alpha, \theta^*)/\partial\theta]^{-1}$ exists. Clearly, $h_\rho(\alpha, \theta)$ is continuous differentiable with respect to α since is a product of continuously differentiable functions. Now, the focus turns to $h_\rho(\alpha, \cdot) : \Theta \longrightarrow \mathbb{R}^k$. First, we note that $\frac{\partial h_\rho}{\partial \theta}(\alpha, \theta) = H_\theta(\Sigma)$, where $H_\theta(\Sigma)$ is (A.2). Since, the function $\partial h_\rho(\alpha, \cdot)/\partial\theta : \Theta \longrightarrow \mathbb{R}^k \times \mathbb{R}^k$ is an integral is clearly continuous on Θ. Therefore (i) is shown.

From the previous proof we know that $\Sigma(\theta)$ is positive definite $\forall(\alpha, \rho, \theta) \in (0, 1) \times \mathbb{R}_+ \times \Theta$, $\therefore$ for any $\rho \in \mathbb{R}_+$. We conclude that $[\partial h_\rho(\alpha, \theta^*)/\partial\theta]^{-1}$ exists $\forall\alpha \in (0, 1)$ and (ii) is verified. Invoking the implicit function theorem, we conclude that for a given $\theta^* \in G$ $\exists!\alpha \in (0, 1)$ such that $\theta^* = \theta_\rho(\alpha)$.$\blacksquare$

Appendix B

Construction of a Contiguous Series of Earnings Forecasts - Chapter 2

We follow the dating convention of I/B/E/S and refer to the third Thursday of a given month as the statistical date, labeled s. Our analysis focuses on the forecast of earnings for the current fiscal year, T, denoted A_T^j for firm j. Upon the announcement of earnings figures for the current fiscal year (T), a shift to the new fiscal year ($T+1$) occurs. Let $f_{s,T}^j$ be the consensus earnings estimate for fiscal year T on statistical data s that applies to firm j. For statistical dates within the fiscal year (T), we compute revisions to the earnings estimate as shown below:

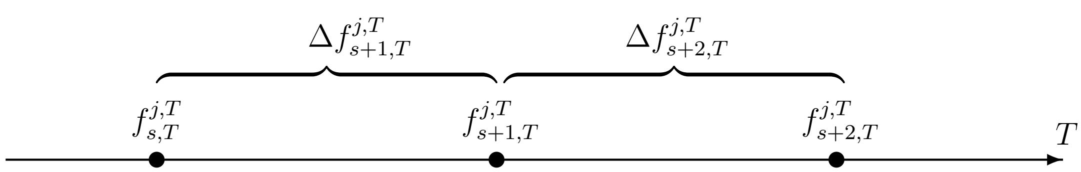

More precisely, when $s+1 \leq T$, the earnings revision for firm j, $\Delta f_{s+1,T+1}^j$, for the fiscal year T is computed as the difference between the earnings estimates on the statistical dates $s+1$ and s, scaled by the initial stock price on day s.

During months where the fiscal year changes from T to $T+1$ we compute the revision to the earnings forecast by comparing the previous month's forecast of earnings for fiscal year $T+1$ to the current month's forecast of this value. This allows us to create a contiguous time-series.

The figure below demonstrates this convention. The figure assumes that $s+1 < T < s+2$ so the end of fiscal year T falls between statistical dates $s+1$ and $s+2$.

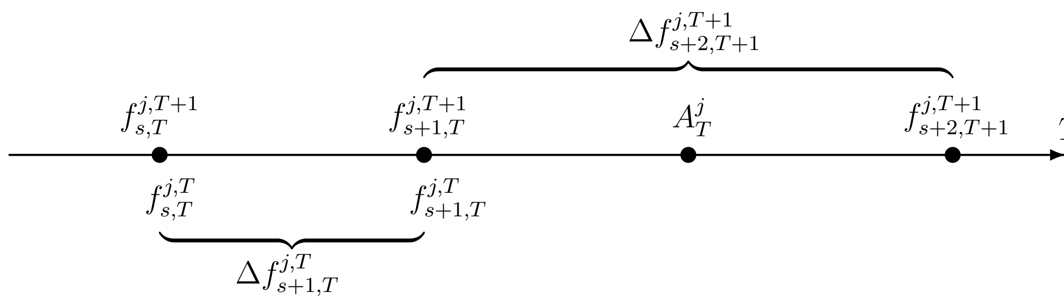

The earnings revision between the statistical dates $s+2$ and $s+1$ corresponds to the fiscal year ending in $T+1$. To create a time-series of forecast revisions and calculate $\Delta f^{j}_{s+2,T+1}$, we use the consensus earnings estimate for fiscal year $T+1$ issued before the end of fiscal year T.

Bibliography

ABARBANELL, J. (1991): "Do Analysts' Earnings Forecasts Incorporate Information in Prior Stock Price Changes," *Journal of Accounting and Economics*, 14, 147–165.

ABARBANELL, J., AND V. BERNARD (1992): "Tests of Analysts Overreaction/Underreaction to Earnings Information as an Explanation for Anomalous Stock Price Behavior," *Journal of Finance*, 57, 1181–1207.

ABARNANELL, J. S., AND B. J. BUSHEE (1997): "Fundamental Analysis, Future Earnings, and Stock Prices," *Journal of Accounting Research*, 35, 1–24.

AIOLFI, M., AND M. RODRIGUEZ (2006): "Financial Analysts'Forecast Revisions and Macroeconomic Information," Working Paper, UCSD.

AIOLFI, M., M. RODRIGUEZ, AND A. TIMMERMANN (2006): "Managing Earnings Expectations: Persistence, Asymmetry and Predictability in Analysts' Earnings Forecasts," Working Paper, UCSD.

AITKEN, M., A. FRINO, AND R. WINN (1996): "Consensus Analysts' Earnings Forecasts and Security Returns," *Asia Pacific Journal of Management*, 13, 101–110.

ALI, A., A. KLEIN, AND J. ROSENFELD (1992): "Analysts' Use of Information about Permanent and Transitory Earnings Components in Forecasting Annual EPS," *The Accounting Review*, 67, 183–198.

ANDERSEN, T. G., T. BOLLERSLEV, AND F. X. DIEBOLD (2006): "Parametric and Nonparametric Volatility Measurement," in *Handbook of Financial Econometrics*, ed. by L. P. Hansen, Amsterdam. North-Holland.

ANDREWS, D. W. K. (1991): "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, 49, 817–858.

BAI, J., AND S. NG (2006): "Confidence Intervals for Diffusion Index Forecasts and Inference for Factor-Augmented Regressions," *Econometrica*, 74, 1133–1150.

BASU, S., AND S. MARKOV (2003): "Loss Function Assumptions in Rational Expectations Tests on Financial Analysts' Earnings Forecasts," Working Paper, Emory University.

BERNHARDT, D., M. CAMPELLO, AND E. KUTSOATI (2004): "Analyst Compensation and Forecast Bias," Working Paper, University of Illinois.

BERNHARDT, D., AND E. KUTSOATI (1999): "Can Relative Performance Compensation Explain Analysts' Forecast of Earnings?," Working Paper, University of Illinois.

BRANDT, M. W., AND Q. KANG (2004): "On the Relation of Between the Conditional Mean and Volatility of Stock Returns: A Latent VAR Approach," *Journal of Financial Economics*, 72, 217–257.

BROUS, P. (1992): "Common Stock Offering and Earnings Expectations: A Test of the Release of Unfavorable Information," *Journal of Finance*, 4, 1517–1536.

BROWN, L., AND M. ROZEFF (1978): "The Superiority of Analyst Forecasts as Measures of Expectations: Evidence from Earnings," *Journal of Finance*, 33, 1–16.

BROWN, L. D. (1993): "Earnings Forecast Research: Its Implications for Capital Markets Research," *International Journal of Foercasting*, 9, 295–320.

BROWN, P., AND R. BALL (1967): "Some Preliminary Findings on the Association Between the Earnings of a Firm, Its Industry and the Economy," *Journal of Accounting Research*, 5, 55–77.

BROWN, P., G. FOSTER, AND E. NOREEN (1985): "Security Analyst Multi-Year Earnings Forecasts and the Capital Market," *Studies in Accounting Research*, 21.

BUTLER, K., AND L. LANG (1991): "The Forecast Accuracy of Individual Analysts: Evidence of Systematic Optimism and Pessimism," *Journal of Accounting Research*, 29, 150–156.

CARTER, R. B., F. H. DARK, AND A. K. SINGH (1998): "Underwriter Reputation, Initial Returns and the Long-Run Performance of IPO Stocks," *Journal of Finance*, 53, 285–311.

CHANEY, P., C. HOGAN, AND D. JETER (1999): "The Effects of Reporting Restructuring Charges on Analysts' Forecast Revisions and Errors," *Journal of Accouting and Economics*, 27, 261–284.

CHEN, Q., J. FRANCIS, AND W. JIANG (2004): "The Role of Forecast Patterns in Conveying Analysts' Predictive Ability," Working Paper, Duke University.

CHEN, Q., J. FRANCIS, AND W. JIANG (2005): "Investors Learning About Analyst Predictive Ability," *Journal of Accounting and Economics*, 39, 3–24.

CHRISTOFFERSEN, P., AND F. DIEBOLD (1997): "Optimal Prediction under Asymmetric Loss," *Econometric Theory*, 13, 808–817.

CLARKE, J., S. FERRIS, N. JAYARAMAN, AND J. LEE (2006): "Are Analyst Recommendations Biased? Evidence from Corporate Bankruptcies," *Journal of Financial and Quantitative Analysis*, 41, 169–196.

CLEMENT, M., AND S. TSE (2003): "Do Investors Respond to Analysts' Forecast Revisions as If Forecast Accuracy Is All That Matters?," *The Accounting Review*, 78, 227–249.

CONRAD, J., B. CORNELL, W. LANDSMAN, AND B. ROUNTREE (2006): "How do Analyst Recommendations Respond to Major News?," *Journal of Financial and Quantitative Analysis*, 41, 25–49.

COOPER, R. A., T. E. DAY, AND C. M. LEWIS (2001): "Following the Leader: A Study of Individual Analysts' Earnings Forecasts," *Journal of Financial Economics*, 61, 383–416.

COWEN, A., B. GROYSBERG, AND P. HEALY (2003): "Which Types of Analyst Firms Make More Optimistic Forecasts?," Working Paper, Harvard University.

DE BONDT, W., AND R. THALER (1990): "Do Security Analysts Overreact," *American Economic Review*, 80, 52–57.

DECHOW, P., A. HUTTON, AND R. SLOAN (1999): "The Relation Between Analysts' Forecasts of Long Term Growth and Stock Price Performance Following Equity Offerings," Working Paper, University of Michigan.

DIAMOND, D., AND R. E. VERRECCHIA (1981): "Information Aggregation in a Noisy Rational Expectations Economy," *Journal of Financial Economics*, 9, 221–235.

DREMAN, D., AND M. BERRY (1995): "Analyst Forecasting Errors and Their Implications for Security Analysis," *Financial Analysts Journal*, 51, 30–42.

DUGAR, A., AND S. NATHAN (1995): "The Effects of Investment Banking Relationships on Financial Analysts' Earnings Forecasts and Investment Recommendations," *Contemporary Accounting Research*, 12, 131–160.

EASTERWOOD, J., AND S. NUTT (1999): "Inefficiency in Analysts' Earnings Forecasts: Systematic Misreaction or Systematic Optimism?," *Journal of Finance*, 54, 1777–1797.

EHRBECK, T., AND R. WALDMAN (1996): "Why Are Professional Forecasts Biased?," *Quarterly Journal of Economics*, 111, 21–40.

ELLIOTT, G., I. KOMUNJER, AND A. TIMMERMANN (2005a): "Biases in Macroeconomic Forecasts: Irrationality or Asymmetric Loss?," Working Paper, UCSD.

——— (2005b): "Estimation and Testing of Forecast Rationality under Flexible Loss," *Review of Economic Studies*, 72, 1107–1125.

ELLIOTT, G., AND A. TIMMERMANN (2004): "Optimal Forecast Combinations Under General Loss Functions and Forecast Error Distributions," *Journal of Econometrics*, 122, 47–79.

ELLIOTT, J. A., D. R. PHILBRICK, AND C. WEIDMAN (1995): "Evidence from Archival Data on the Relation Between Security Analysts' Forecasts Errors and Prior Forecast Revisions," *Comtemporary Accounting Research*, 11, 919–938.

ELTON, E., AND M. GRUBER (1972): "Earnings Estimates and the Accuracy of Expectational Data," *Management Science*, 18, 409–424.

FAMA, E., AND K. FRENCH (1988): "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–25.

——— (2000): "Forecasting Profitability and Earnings," *Journal of Business*, 73, 161–175.

FAMA, E., AND G. SCHWERT (1977): "Asset Returns and Inflation," *Journal of Financial Economics*, 5, 115–146.

FERSON, W., AND C. HARVEY (1991): "The Variation of Economic Risk Premiums," *Journal of Political Economy*, 99, 385–415.

FRANCIS, J., AND D. PHILBRICK (1993): "Analysts' Decision as Products of a Multi-Task Environment," *Journal of Accounting Research*, 31, 216–230.

FRIED, D., AND D. GILVOY (1982): "Financial Analysts' Forecasts of Earnings: A Better Surrogate for Market Expectations," *Journal of Accounting and Economics*, 4, 85–107.

GLEASON, C. A., AND C. M. C. LEE (2002): "Analyst Forecast Revisions and Market Price Discovery," Working Paper, Cornell University.

GLOSTEN, L., R. JAGANNATHAN, AND D. RUNKLE (1993): "On the relation between the Expected Value and the Volatility of the Nominal Excess returns on Stocks," *Journal of Finance*, 48, 1770–1801.

GLOSTEN, L., AND P. MILGROM (1985): "Bid, Ask and Transaction Prices in a specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71.

GONEDES, N. J. (1973): "Properties of Accounting Numbers: Models and Tests," *Journal of Accounting Research*, 11, 212–237.

GOYAL, A., AND I. WELCH (2003): "Predicting the Equity Premium with Dividend Ratios," *Management Science*, 49, 639–654.

GRANGER, C. W. J. (1969): "Prediction with a Generalized Cost of Error Function," *Operations Research Quarterly*, 20, 199–207.

——— (1999): "Outline of Forecast Theory Using Generalized Cost Functions," *Spanish Economic Review*, 1, 161–173.

GRANGER, C. W. J., AND M. MACHINA (2006): "Forecasting and Decision Theory," in *Handbook of Economic Forecasting*, ed. by G. Elliott, C. W. J. Granger, and A. Timmermann, Amsterdam. North-Holland.

GRANGER, C. W. J., AND P. NEWBOLD (1986): *Forecasting Economic Time Series*. Academic Press.

GU, Z., AND J. WU (2003): "Earnings Skewness and Analyst Forecast Bias," *Journal of Accounting and Economics*, 35, 5–29.

HANSEN, L. P. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.

HONG, H., AND J. KUBIK (2003): "Analyzing the Analysts: Career Concerns and Biased Earnings Forecasts," *Journal of Finance*, 58, 313–351.

I/B/E/S (2000): *The I/B/E/S Glossary*. I/B/E/S International, Inc.

JAIN, P. C. (1992): "Equity Issues and Changes in Expecations of Earnings by Financial Analysts," *Review of Financial Studies*, 5, 669–683.

KANG, S., J. O'BRIEN, AND K. SIVARAMAKRISHNAN (1994): "Analysts' Interim Earnings Forecasts: Evidence on the Forecasting Process," *Journal of Accounting Research*, 32, 103–112.

KEANE, M. P., AND D. E. RUNKLE (1998): "Are Financial Analysts' Forecasts of Corporate Profits Rational?," *Journal of Political Economy*, 106, 768–805.

KLEIN, A. (1990): "A Direct Test of the Congnitive Bias Theory of Share Price Reversals," *Journal of Accounting and Economics*, 13, 155–166.

KOENKER, R., AND G. BASSET (1978): "Regression Quantiles," *Econometrica*, 46, 30–50.

KOTHARI, S. (2001): "Capital Markets Research in Accounting," *Journal of Accounting and Economics*, 31, 105–231.

LAMONT, O. (1995): "Macroeconomic Forecasts and Microeconomic Forecasts," NBER Working Paper.

LASTER, D., P. BENNETT, AND I. S. GEOUM (1999): "Rational Bias in Macroeconomic Forecasts," *Quarterly Journal of Economics*, 114, 293–318.

LIM, T. (2001): "Rationality and Analysts' Forecast Bias," *Journal of Finance*, 56, 369–385.

LIMA, E. L. (1992): *Curso de Análise Volume 1*. IMPA.

LIN, H., AND M. MCNICHOLS (1998): "Underwriting Relationships, Analysts' Earnings Forecasts and Investment Recomendations," *Journal of Accounting and Economics*, 25, 101–127.

LIU, J., AND J. THOMAS (2000): "Stock Returns and Accountning Earnings," *Journal of Accounting Research*, 38, 71–101.

LOH, R. K., AND G. M. MIAN (2005): "Do Accurate Earnings Forecasts Facilitate Superior Investment Recommendations?," *Journal of Financial Economics*, Forthcoming.

LUDVIGSON, S., AND S. NG (2006): "The Empirical Risk-Return Relation: A Factor Analysis Approach," *Journal of Financial Economics*, Forthcoming.

LYS, T., AND S. SOHN (1990): "The Association Between Revisions of Finacial Analysts' Earnings Forecasts and Security-Price Changes," *Journal of Accounting and Economics*, 13, 341–362.

MAGEE, R. P. (1974): "Industry-Wide Commonalities in Earnings," *Journal of Accounting Research*, 12, 270–287.

MANSKI, C. (1991): "Regression," *Journal of Economic Literature*, 29, 34–50.

MCLEAN, B., AND P. ELKIND (2003): *The Smartest Guys in the Room: the Amazing Rise and Scandolous Fall of Enron*. Portfolio, New York.

MCNICHOLS, M., AND P. O. O'BRIEN (1997): "Self-Selection and Analyst Coverage," *Journal of Accounting Research*, 35, 167–199.

MENDEHALL, R. (1991): "Evidence of Possible Underweight of Earnings-Related Information," *Journal of Accounting Research*, 29, 170–179.

MICHAELY, R., AND K. L. WOMACK (1999): "Conflict of Interest and the Credibility of Underwriter Analyst Recommendations," *Review of Financial Studies*, 12, 653–686.

MILLER, J., AND L. SEDOR (2006): "The Influence of Observed Stock Price Changes on Analysts' Earnings Forecast Revisions: Experimental Evidence," Working Paper, Notre Dame University.

NEWEY, W. K., AND K. D. WEST (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

O'BRIEN, P. (1988): "Analysts' Forecasts as Earnings Expectation," *Journal of Accounting and Economics*, 10, 53–83.

O'BRIEN, P. C. (1994): "Corporate Earnings Forecasts and the Macroeconomy," Working Paper, Michigan Business School.

PARK, C. W., AND E. K. STICE (2000): "Analyst Forecasting Ability and the Stock Price Reaction to Forecast Revisions," *Review of Accounting Studies*, 5, 259–272.

PATTON, A., AND A. TIMMERMANN (2002): "Properties of Optimal Forecasts," Working Paper, London School of Economics.

PATTON, A., AND A. TIMMERMANN (2006): "Properties of Optimal Forecasts under Asymmetric Loss and Nonlinearity," *Journal of Econometrics*, Forthcoming.

PHILBRICK, D. R., AND W. E. RICKS (1991): "Using Value Line and IBES Analyst Forecasts in Accoutning Research," *Journal of Accounting Research*, 29, 397–417.

RAMANNATH, S., S. ROCK, AND P. SHANE (2006): "A Review of Research Related to Financial Analysts' Forecasts and Stock Recommendations," Working Paper, Georgetown University.

RODRIGUEZ, M. (2004): "Financial Analysts' Asymmetry and Rationality," Working Paper, UCSD.

RODRIGUEZ, M. (2005): "Financial Analysts' Incentives and Forecast Biases," Working Paper, UCSD.

SCHARFSTEIN, D. S., AND J. C. STEIN (1990): "Herd Behavior and Investment," *American Economic Review*, 80, 464–479.

SCHERBINA, A. (2004): "Analyst Disagreement, Forecast Bias and Stock Returns," Working Paper, Harvard University.

SCHIPPER, K. (1991): "Analysts' Forecasts," *Accounting Horizons*, 5, 105–121.

STICKEL, S. E. (1991): "Common Stock Returns Surrounding Earnings Forecast Revisions: More Puzzling Evidence," *The Accounting Review*, 66, 402–416.

STOCK, J. H., AND M. W. WATSON (2002a): "Forecasting Using Principal Components from a Large Number of Predictors," *Journal of the American Statistical Association*, 97, 1167–1179.

——— (2002b): "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business and Economic Statistics*, 20, 147–162.

——— (2004): "Forecasting with Many Predictors," Working Paper, Harvard University.

STOCK, J. H., AND M. W. WATSON (2005): "An Empirical Comparison of Methods for Forecasting Using Many Predictors," Working Paper, Harvard University.

STOCK, J. H., J. H. WRIGHT, AND M. YOGO (2002): "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments," *Journal of Business and Economic Statistics*, 20, 518–529.

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

VARIAN, H. R. (1974): *Studies in Bayesian Econometrics and Statistics in Honor of Leornard J. Savage*North-Holland.

VERONESI, P. (1999): "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model," *Review of Financial Studies*, 12, 975–1007.

WHITE, H. (1980): "A Heteroskedasticity-consistent Covariance Matrix Estimatior and Direct Test for Heteroskedasticity," *Econometrica*, 48, 817–838.

ZELLNER, A. (1986): "Best Predictors, Rationality and the Evaluation of Forecasts?," *Economics Letters*, 21, 45–48.