Ambiguity aversion, company size and the pricing of earnings forecasts

Constantinos Antoniou[1]

Emilios C. Galariotis

Daniel Read

November 2011

Abstract: Several authors have reported an unconditional size effect in returns around earnings announcements. In this study we show how this finding can be understood as resulting from ambiguity aversion. We hypothesize that analyst forecasts for smaller companies are relatively more ambiguous; hence they are priced pessimistically by ambiguity-averse investors. As the quarter comes to a close and ambiguity gradually subsides, the stock prices of smaller companies rise to correct this pessimism, creating the size effect. Our results support these hypotheses.

Keywords: Ambiguity aversion, size premium, analyst earnings forecasts.

JEL Classification: D03, D81, D84, G11, G14

1. Introduction

Neoclassical valuation theory postulates that investors calculate a projection of earnings, determine its present value, and trade accordingly. An assumption of this framework is that the distribution of expected earnings can be fully calculated from available information, and so decisions are made in conditions of *risk*. However, in some cases the available information may be of low quality, and investors will be unable to confidently pin down the exact distribution of expected earnings. In these cases decisions are made in conditions of *ambiguity*. Epstein and Schneider (2008) show theoretically that in such conditions pricing will be initially pessimistic due to ambiguity aversion, [2] with this pessimism being corrected with an upward price movement as ambiguity gradually subsides.

A pattern observed in stock market data is a "size effect" around earnings announcements (Ball and Kothari, 1991; Atiase 1985; Foster et al., 1984; Bergman et al., 2008). Ball and Kothari (1991) find that smaller firms have higher abnormal returns around earnings announcements compared to larger firms, *irrespective* of the sign or magnitude of the earnings surprise. This size effect is puzzling because it is of considerable economic magnitude and does not appear to reflect compensation for risk. Foster et al (1984) extend this result and show that the unconditional upward movement in the prices of small companies begins months *before* the earnings announcements. In this paper we provide an ambiguity based explanation for this finding. Our hypothesis is that the analyst forecasts for smaller companies used by investors to form earnings expectations[3] are relatively more ambiguous because there is less 'hard' information with which to estimate forecast accuracy

---

[2] Ambiguity aversion is one of the most well-established results in behavioural economics. A large literature, starting with Knight (1921) and Keynes (1921), and continuing through Ellsberg (1961) and up to the present day (Ahn et al., 2009), shows that situations that involve *ambiguity* are treated differently from those that involve risk. Hsu et al. (2005) present evidence that ambiguous situations produce a unique neurological fingerprint, suggesting that ambiguity aversion is rooted in the fundamentals of human cognition. See Camerer and Weber (1992) and Keren and Gerritsen (1999) for reviews on the evidence on ambiguity aversion.

[3] The importance of analyst forecasts to investors is evident from the robust evidence that they exert a powerful influence on asset prices (Givoly and Lakonishok, 1979; Stickel, 1992; Gleason and Lee, 2003). The impact of analyst forecasts is so large that it is common in the finance literature to treat analysts' forecasts as a proxy for market expectations (Livnat and Mendenhall, 2006).

(something which we estimate quantitatively).[4] These forecasts are consequently priced pessimistically due to ambiguity aversion. As the quarter comes to a close and more information about upcoming earnings is revealed, ambiguity is gradually resolved and so the stock prices of smaller companies drift upwards to correct this pessimism, generating the documented size effect.

Our first test examines whether smaller firms are likely to be perceived as more ambiguous by investors. It is motivated by Daniel Ellsberg's (1961) proposal that when there is little reliable information about the relative likelihood of events, decision makers will feel they face ambiguity.[5] In our framework the event of interest is a forecast, and the likelihood of interest is its accuracy. When this accuracy can be predicted from available information, there is no ambiguity. On the other hand, when the information is of low quality and cannot produce reliable estimates of accuracy there is ambiguity. In this paper we estimate expected accuracy using the model proposed by Clement and Tse (2003), and then examine how this estimate predicts actual accuracy for small and large companies. We show that the relationship between estimated and actual accuracy is significantly weaker for smaller companies, which amounts to higher ambiguity.

We then examine whether forecasts for small companies elicit pessimistic valuations. Our hypothesis, based on ES, is that when investors receive a forecast for a small company they will overweight the worst-case scenario and respond pessimistically. If, therefore, the forecast brings "bad news" (i.e., an earnings decrease), the worst case scenario is that it is very accurate whereas if it brings "good news" (i.e., an earnings increase), the worst case scenario is that it is very inaccurate. In both cases the resulting stock price would be *too low* relative to the price that would be achieved in the absence of ambiguity. Therefore, as ambiguity subsides when the quarter comes to a close prices will rise to correct the previous pessimism. On the contrary, when the forecast is targeted toward large

---

[4] Along these lines various authors have suggested that the information environment of smaller companies is poorer (Atiase, 1980, 1985; Zhang, 2006; Hong et al., 2000).

[5] Frisch and Baron (1988) proposed that "ambiguity is uncertainty about probability, created by missing information that is relevant and could be known" (P. 1988). Einhorn and Hogarth (1985) suggest that ambiguous situations arise when the available information is vague, and does not allow one to confidently rule out alternative possibilities, while Gärdenfors and Sahlin (1982, 1983) argue that feelings of ambiguity are produced when the relevance of the available information is low. For all these authors the underlying theme is that ambiguity is negatively related to what might be called the "richness" of the information that can be used to compute the likelihood of interest.

companies the initial response will not be pessimistic; therefore such a pattern should not arise. Our results confirm this hypothesis. The appendix presents a stylized theoretical model of subjective beliefs that formalizes this prediction.

Our study makes two contributions. Firstly, our evidence highlights that the distribution of expected earnings of smaller companies is (relatively) more ambiguous. Moreover, the evidence that the pricing of these forecasts is consistent with pessimism suggests that investors have (negative) preferences toward this ambiguity, contradicting Subjective Expected Utility (SEU) and Bayesian updating. Jointly, this evidence supports the notion that the size effect around earnings announcements is related to ambiguity aversion.

Secondly, our study contributes to the literature on ambiguity aversion, which mainly uses theoretical or experimental tools of analysis (Camerer and Kunreuther, 1989; Sarin and Weber, 1993; Chen and Epstein, 2002). One exception to this is an important recent study by Anderson et al (2009), who decompose market returns into a risk and an ambiguity premium, and conclude that the latter is more important. Anderson et al (2009) use a market-wide measure of ambiguity and thus focus their analysis on market returns. Our study complements theirs by illustrating the importance of ambiguity and ambiguity aversion at the company level.

The remainder of the paper is organized as follows: Section 2 describes our methodology, defines the variables used and describes the sample. Section 3 presents and discusses the results, and Section 4 concludes.

2. Methodology and Data

2.1 *Are small companies more ambiguous than larger ones?*

As explained in the introduction we use the model proposed by Clement and Tse (2003) to estimate forecast accuracy using available information, and then examine how well this estimate predicts actual

accuracy for companies of different size. Firstly, we use their method to standardise the variables used to predict forecast accuracy (Clement and Tse, 2003; 2005):

$$\text{Accuracy}_{ijt} = \frac{\max(\text{AFE}_{it}) - \text{AFE}_{ijt}}{\max(\text{AFE}_{it}) - \min(\text{AFE}_{it})} \qquad (1)$$

Where *AFE* is *Absolute Forecast Error*, defined as |(Forecasted earnings–Actual earnings)/Actual earnings|, and i indexes the specific company, j the analyst and t the period of the forecast. The accuracy of a given forecast is given by the absolute value of its forecast error (AFE_{ijt}), subtracted from the maximum AFE for that company and period issued by *any* analyst, divided by the range of absolute errors for forecasts for that period. Accuracy ranges from 0 to 1, where a higher value indicates higher accuracy. For each forecast we then measure eight variables that Clement and Tse (2003) have shown to predict accuracy. In our context these variables constitute the information set from which investors can estimate expected accuracy:

1) *Forecast horizon*: the days that separate a forecast from its corresponding earnings announcement date.

2) *Days elapsed*: the days intervening between a forecast and the previous forecast issued by any analyst for the same company and period.

3) *Lag accuracy*: the value of *Accuracy* for the analyst j's last forecast for a company i in the previous period.

4) *Firm experience*: the number of quarters that analyst j has been covering company i.

5) *Broker size*: the number of analysts employed by the brokerage house in which the analyst belonged in period t.

6) *Forecast frequency*: the number of forecasts issued by analyst j for company i in period t.

7) *Industries*: the number of industries (4-digit SIC) that analyst j is following in period t.

8) *Companies:* the number of companies that analyst j is following in period t.

All these variables (except *Lag accuracy*) are standardised to vary between 0 and 1 using the following formula[6]:

$$St.Var._{ijt} = \frac{Raw\,Var._{ijt} - \min(Raw\,Var._{it})}{\max(Raw\,Var._{it}) - \min(Raw\,Var._{it})} \tag{2}$$

We use the standardized variables to estimate the predicted accuracy of each forecast using the following cross sectional regression in each year *t* (subsuming analyst and company subscripts):[7]

$$Accuracy_t = \alpha_t + \beta_{1t}Comp.Exp_t + \beta_{2t}Bro\,kerSize_t + \beta_{3t}Companies_t \\ + \beta_{4t}For.freq._t + \beta_{5t}Days_t + \beta_{6t}For.Horizon_t + \beta_{7t}Industries_t + \beta_{8t}Lag\,Acc._t + u_t \tag{3}$$

Using these coefficients and constant, we then derive the *Expected Accuracy* of forecasts issued in year *t+1* as in Brown (2001):

$$Expected\,Accuracy_{t+1} = \alpha_t + \beta_{1t}Comp.Exp_{t+1} + \beta_{2t}Bro\,kerSize_{t+1} + \beta_{3t}Companies_{t+1} \\ + \beta_{4t}For.freq._{t+1} + \beta_{5t}Days_{t+1} + \beta_{6t}For.Horizon_{t+1} + \beta_{7t}Industries_{t+1} + \beta_{8t}Lag\,Acc._{t+1} \tag{4}$$

Expected Accuracy is therefore an estimate of forecast accuracy using information available to the investor. As explained, ambiguity arises to the degree that expected accuracy (Equation 4) deviates from actual accuracy (Equation 1). To examine whether company size predicts ambiguity in this sense, we regress *Actual Accuracy* on *Expected Accuracy* using dummy variables for company size:[8]

$$Actual\,Accuracy_{ijt} = \alpha + \alpha_1 Middle_{it} + \alpha_2 Small_{it} + \theta_1 ExpectedAc\,curacy_{ijt} \\ + \theta_2 ExpectedAc\,curacy_{ijt} * Middle_{it} + \theta_3 ExpectedAc\,curacy_{ijt} * Small_{it} + u_{ijt} \tag{5}$$

---

[6] For example, if an analyst who issues a forecast for IBM in the first quarter of 2000 follows 3 industries, and the IBM-covering analyst who, for the same period, covers the fewest industries covers 2 industries, and the one who covers the most industries covers 10, then for this analyst *Industries* would equal $(3 - 2)/(10 - 2) = 0.125$.

[7] For simplicity we estimate the regression yearly, as opposed to quarterly, because accounting quarters frequently correspond to different calendar months for different companies.

[8] To identify the levels of ambiguity we split our sample into quintiles based on company size (using the breakpoints from Kenneth French's website in month t to assign the forecasts in month t+1).

Where *Middle* is a dummy that equals 1 if the size of the company belongs in the 2nd, 3rd or 4th quintile (i.e., is of intermediate size) and *Small* equals 1 if company is in the 1st quintile. The coefficient of "predictability", θ_1, shows how well expected accuracy (Equation 4) predicts actual accuracy (Equation 1) for the low ambiguity (i.e., *largest*) companies. We predicted that as ambiguity increases, this relationship would become weaker, so that $\theta_3 < \theta_2 < 0$. That is, the smaller the company, the less predictable is accuracy from available information, and therefore the greater the ambiguity.

Our test of ambiguity is based on the notion that the market uses available information to estimate forecast accuracy. If so, it then identifies whatever relationship exists between firm size and the predictability of forecast accuracy. Although there is already evidence that investors use the information cues from Equation 4 to estimate forecast accuracy (see Clement and Tse, 2003; Bonner et al., 2003; Mikhail et al., 2003, 2004), we also test whether these relationships emerge in our sample using the regression proposed by Clement and Tse (2003). Specifically we regress event time returns (-1, 1) interacted with revision magnitude on the accuracy-related variables in equation 3. If the market is using this information we should observe that variables that predict higher accuracy have a higher impact on returns. The regression is of the form:

$$
\begin{aligned}
CAR_{ijt} = \alpha_t &+ Re\,vMagn_{ijt} * (b_1 + \beta_2 Comp.Exp_{ijt} + \beta_{2t} Bro\,kerSize_{ijt} \\
&+ \beta_{3t} Companies_{ijt} + \beta_{4t} For.freq._{ijt} + \beta_{5t} Days_{ijt} + \beta_{6t} For.Horizon_{ijt} \\
&+ \beta_{7t} Industries_{ijt} + \beta_{8t} LagAcc._{ijt}) + u_t
\end{aligned}
\tag{6}
$$

2.2. *Market reaction to ambiguity*

To examine whether a size effect arises in response to analyst forecasts, we group these forecasts according to the size of the targeted company and then perform a two-stage event study.[9] The post-forecast trading period is divided into an impact stage during which we expect pessimism due to ambiguity aversion to be priced, and an adjustment stage during which the pessimistic response will

---

[9] Company size is measured as price × shares outstanding three days prior to the release of a forecast revision.

be corrected. The impact stage includes trading days -2 to 2 where 0 is the day the forecast is issued, and the adjustment stage includes days 3 to 20.[10] For each stage we compute risk adjusted returns using the modified market model.

In the adjustment period ambiguity is partly resolved because the end of the quarter nears and new information is released about upcoming earnings. This rationale is supported by the empirical evidence in Bernard and Thomas (1989) and Campbell et al (2008), who show that prices move in the direction of the earnings surprise months before the announcement, consistent with an expansion in the information set. We argue that this expansion prompts the market to revise its expectation to a less pessimistic level so the prices after ambiguous forecasts will on average drift upward, *irrespective* of the sign of these forecasts. The model in the appendix demonstrates that such an upward movement is not consistent with a Bayesian model with risk, but can be explained by a framework which incorporates ambiguity and ambiguity aversion.

An analogous adjustment period is used by ES in their example regarding the response of the market to the 9/11 events. ES explain that initially investors faced ambiguity in terms of the economic consequences of these events and, being averse to ambiguity, they priced stocks pessimistically. With the passage of time this ambiguity was resolved and the market upwardly adjusted it's valuation to a less pessimistic level. Our event study design is aimed to capture an analogous (but much more mundane) pattern of pessimism-and-correction around analyst forecasts.

2.3 Multivariate analysis

Our univariate analysis is complemented by controls for other variables known to affect post-revision returns (Clement and Tse, 2003). CAR, in the equation below, is the market adjusted cumulative abnormal return for company *j* at time *t* after the forecast revision issued by analyst *i*. Our variable of interest is the inverse logarithm of company size. Our control variables are: average quoted bid ask

---

[10] When examining the impact of an event it is customary to define a window around the event as the event period to capture information leakages and gradual information diffusion (e.g. Sorescu and Subrahmanyam, 2006). We use a somewhat wider window (i.e., -2, 2, instead of -1, 1) to allow the price to absorb the information.

spread for the twenty-day period prior to the issuance of the forecast (to control for potential liquidity effects), forecast horizon (the days intervening between a forecast revision and the earnings announcement), broker size (the number of analysts employed by the brokerage house employing the analyst who issued a forecast) and the absolute value of the revision magnitude ((new revision − old revision)/old revision). In addition, we control for analyst forecast bias (signed forecast error) to ensure that the pricing patterns we document are due to ambiguity aversion and not to analysts systematically under predicting earnings (Matsumoto, 2002). The multivariate regression is of the form:

$$\text{CAR}_{ijt} = \alpha_0 + \beta_1 \log \text{size}_{ijt}^{-1} + \beta_2 \text{BidAsk}_{ijt} + \beta_3 \text{For.Horizon}_{ijt} + \tag{7}$$
$$\beta_4 \text{Bro ker Size}_{ijt} + \beta_5 \text{For.Error}_{ijt} + \beta_6 \text{Re vMagn}_{ijt} + u_{ijt}$$

We conduct two regressions: for returns during the adjustment stage separately for upward and downward forecasts. To reduce the effect of correlated residuals due to the panel nature of our dataset in all regressions we cluster on the firm level and use year fixed effects.

2.4 *Sample construction and descriptive statistics*

Data on all quarterly US analyst forecast revisions and actual earnings are collected from the IBES detail files. The sample period spans from January 1995 to March 2008.[11] Data on returns and shares outstanding are from CRSP. We apply the following filters to the sample: First, a company must have data from both CRSP and IBES. Second, if the forecast estimation date is after the earnings announcement date the observation is deleted. Third, all revisions greater than 100% and all revisions which entail an error in excess of 100% of the actual earnings are removed from the sample as these are likely to be erroneous according to Capstaff et al. (1995). Fourth, all revisions equal to 0 are deleted. After the above standardizations, and consistent with other studies (Clement and Tse, 2003; Clement, 1999; O'Brien, 1990), only the last forecast issued by each analyst-firm pair for each quarter is retained. In addition, we impose a minimum forecast horizon of 30 days (to eliminate the effect of

---

[11] Similar to Clement and Tse (2003) we choose 1995 as the cut-off point to ensure that the issuance date of the revision reported in the IBES files is accurate.

the earnings announcement), and a maximum of 90 days (to avoid stale forecasts). After these filters are applied, we end up with 264,963 earnings forecast revisions. This number may change slightly depending on the variables required for each test, and this information is provided in each table.

Table 1, Panel A, presents descriptive statistics for the variables used. The means and medians for the accuracy related variables are comparable to those of Clement and Tse (2003). The average revision is -4.34%, and the average forecast error is -4.51%. These figures are consistent with prior evidence (e.g., Richardson et al., 2004; Matsumoto, 2002). In terms of company size we observe that the mean size decile rank (using end of previous month breakpoints) for the firms in the sample is 6.5, which indicates the sample is slightly tilted towards large companies. This is a common finding when IBES data are used because large companies have larger analyst coverage. However, the sample does include small firms, as 25% are below a decile rank of 4. Panel B, shows the annual frequency of analyst forecasts according to revision direction and company size. Our sample contains 150,554 downward forecasts and 114,409 upward ones.

[Please insert Table 1 here]

3. Results

3.1. *Estimating forecast ambiguity*

In this section we examine whether small companies are indeed more ambiguous than larger ones. Before moving on to the main test we replicate previous results that forecast accuracy can be partly predicted from previous information. Table 2 shows the output from applying the model shown in regression 3 to the full sample. Consistent with earlier research (Clement and Tse, 2003), we find that accuracy increases with broker size, forecast frequency and lag accuracy, and decreases with the number of industries and companies followed by the analyst, days elapsed and forecast horizon. These relationships are highly significant, which confirms that forecast accuracy is partly predictable from available information.

[Please insert Table 2 here]

We next examine whether the market responds to factors that predict forecast accuracy. We estimate Equation 6, whereby we regress the market adjusted cumulative return from days -1 to 1 (0 being the forecast release date) on the accuracy factors interacted with revision magnitude. The results show that the market responds more strongly to forecasts issued from larger brokerage houses, to analysts with greater lagged accuracy and to analysts who issue more forecasts during the quarter. It responds less strongly to analysts who follow many industries, and to forecasts with more days elapsed.[12] These results suggest that investors use these accuracy-related factors to estimate expected earnings.

[Please insert Table 3 here]

We now estimate whether company size is associated with the predictability of forecast accuracy. Table 4 shows estimates of the model in Equation 5. The coefficient of predictability, θ_1, for large companies (low ambiguity forecasts) is 0.716. Since a coefficient of 1 indicates a perfect fit, this suggests that the available information for larger companies can be used to predict forecast accuracy fairly precisely. The coefficient is significantly lower for mid-sized companies as it decreases by -0.104 (p-value <0.0001), and much lower for small ones as it decreases by -0.237 (p-value <0.0001). In other words, the predictability of forecast accuracy from expected accuracy drops by 33% when moving from the largest to the smallest companies, hence the forecast accuracy of smaller companies is less predictable than that of larger ones. This means that the factors that the market is using to predict forecast accuracy are of *lower* quality when the forecast is targeted toward smaller companies. According to the definitions of ambiguity provided by Ellsberg and others this implies that ambiguity rises as one moves from larger to smaller companies.

[Please insert Table 4 here]

---

[12] However, as reported by previous studies such as Clement and Tse (2003), we find that in some cases the market responds more strongly to variables that are less predicative of accuracy. For example, the market reacts more strongly to forecasts issued earlier in the quarter. See Clement and Tse (2003) for possible explanations.

3.2. *Impact and adjustment stage returns by ambiguity proxy*

Table 5 presents market adjusted cumulative abnormal returns for equally weighted portfolios formed according to company size for the impact (-2, 2) and the adjustment (3, 20) stages for upward and downward forecasts respectively. Responses to ambiguity should be visible on impact period returns as well, i.e., the impact of downward forecasts for small companies should exceed that of upward forecasts. This prediction, however, can only be made if we assume that the *only* dimension that varies when we compare upward and downward forecasts is the sign of the revision. In reality however, this assumption is not likely to hold, as upward forecasts may on average be viewed more positively than downward forecasts are viewed negatively. [13] So, even if returns after downward forecasts for smaller companies do not exceed in absolute value returns after upward forecasts (as we actually observe in the data), it may still be the case that the response to the former (latter) is stronger (weaker) than what it should be *in the absence* of ambiguity. For this reason we test our hypothesis based on patterns of adjustment in prices that are indicative of previous mispricings, as is common in the literature (i.e., Baker and Wurgler, 2006).

Adjustment period returns are analyzed in Panel B. For upward forecasts, return continuations occur exclusively amongst smaller companies, while for larger ones adjustment stage returns are close to zero. This generates an economically and statistically significant size effect, equal to 0.72%. A similar picture emerges when we consider downward revisions. For large companies, adjustment stage returns for all magnitude categories are negative, indicating that prices continue to drop in the adjustment stage. However, as we move toward smaller companies the returns become positive, creating a significant size effect, equal to 0.81%. This again shows that prices in the adjustment stage for small companies increase substantially. In short, and as predicted, the returns of small companies

---

[13] This is consistent with a well known finding in the analyst forecast literature is that analysts initially issue optimistic forecasts, and gradually walk them down as the quarter comes to a close (Matsumoto, 2002). Therefore, at least some of the downward forecasts in our sample do not bring new information and only reflect a correction to optimism, which suggests that the average information content of upward and downward analyst forecasts *is not* equal.

after analyst forecasts exhibit a strong upward movement, whereas the returns of large companies do not.[14]

To examine whether the adjustment period has completely eliminated the underpricing, or whether some of it is resolved around earnings announcements we also analyze market adjusted cumulative returns around earnings announcement (trading days -5, 0; as in Cornell et al., 2002). The results are shown in Panel C. For both upward and downward forecasts we observe that for smaller companies there is an upward movement in prices irrespective of the earnings surprise (as documented by Ball and Kothari, 1991), which creates a size effect of 50 basis points. This suggests that some of the pessimism induced by the preceding analyst forecasts is corrected around the announcement date.

Zhang (2006), also studies investor response to analyst forecasts when information uncertainty is higher. One of his proxies for uncertainty is company size because, similarly to our intuition, he suggests that the accuracy of forecasts for smaller companies is more difficult to predict. He reports a size effect *only* in response to upward revisions, but no size effect after downward revisions. Our analysis, however, is slightly different because we are testing whether a size effect exists *only* in the adjustment period, whereas Zhang measures total post-revision returns. If we combine impact and adjustment periods the total average returns of small (large) companies after a downward revision is -2.11% (-1.94%), which implies an annualized size effect comparable to Zhang's findings.

Several authors have documented an aggregate cross sectional size effect (Banz 1981, Reinganum 1981, Fama and French 1992). Under our hypothesis this size effect may also be partly related to ambiguity. [15] Anderson et al (2009) provide some supportive evidence for this claim as they

---

[14] Constructing the portfolios on value weighted basis does not change any of these conclusions. These results are available upon request.
[15] The notion that the aggregate size premium reflects ambiguity has been first proposed in the literature by Olsen and Troughton (2000).

report that the correlation between their ambiguity measure and the Small minus Big (SMB) portfolio is 0.24.

Overall, the behaviour of returns in the adjustment stage is as predicted by the ambiguity aversion hypothesis: because of ambiguous forecasts, investors overweight pessimistic scenarios, and consequently set prices too low, generating an economically and statistically significant size effect in the period around earnings announcements.

[Please insert Table 5 here]

3.3. *Multivariate analysis*

We now examine the robustness of results in a multivariate setting. Table 6 Panels A and B, show the analysis for downward and upward forecasts respectively. The results are consistent with the univariate analysis. Adjustment period returns for upward and downward forecasts increase when the revision is targeted towards smaller companies. We also find that the upward movement in prices increases with the forecast horizon and decrease with signed forecast error.[16]

[Please insert Table 6 here]

4. Concluding remarks

In neoclassical theories investment decisions are made under conditions of risk, where investors have access to sufficient information to calculate the distribution of expected returns. However, in reality, the available information in some cases will be scanty or unreliable, preventing precise identification of this distribution. Several authors have documented that investors are averse to such ambiguity and price assets pessimistically, which generates an ambiguity premium.

In this study we examine whether the previously documented unconditional size effect around earnings announcements reflects an ambiguity premium. Our hypothesis is that analyst forecasts for

---

[16] Long horizon forecasts can also be seen as more ambiguous as it is more difficult to estimate their accuracy due to lack of information. Therefore, the positive relationship between horizon and adjustment period returns may reflect also pessimism.

smaller companies are relatively more ambiguous; therefore investors respond pessimistically towards them due to ambiguity aversion. Then, as the quarter comes to a close and ambiguity is gradually resolved, prices rise to correct previous pessimism and generate the documented size effect.

In ours tests we first confirm that the accuracy of forecasts for smaller companies is more difficult to estimate from available information. According to the definitions of ambiguity provided by Daniel Ellsberg and others, this amounts to higher ambiguity. Second, we demonstrate that after an initial "impact" period, stock returns of small companies show evidence of pessimism, by rising significantly during an "adjustment" period *regardless* of the sign of the revision. No such increase is observed after revisions for larger companies. These results support our hypothesis that ambiguity aversion relates to the documented size effect around earnings announcements.

Appendix

The exposition below follows closely the model by Kelsey et al. (2010). It also draws on the model by Epstein and Schneider (2008). The model aims to illustrate the precise mechanism underlying our hypothesis, and to highlight that our results cannot be explained in a risk-only Bayesian framework. In the model we have three periods, one representative risk-neutral agent and one risky asset, which pays all earnings as dividends.

In period 0 the agent has some prior about the dividend, and sets the price. In period 1 the agent receives a signal (either risky or ambiguous), and revises the price. In the third period we assume that risk or ambiguity are partly resolved, so the price changes once more. The intuition of the resolution is that as the terminal date approaches (i.e., when the dividend is revealed) the market is exposed to more information, which helps resolve some of the existing risk or ambiguity. We examine price changes in the adjustment period and highlight that our empirical findings, i.e., the upward movement in prices after good and bad news forecasts, only obtains when we allow the signal to be ambiguous and the investor to be ambiguity averse.

To draw a parallel between the model and our empirical work, the first period is where investors receive the analyst forecast, so prices are updated. In the second period, as the quarter comes to a close, more information is revealed which help resolve some of the risk or ambiguity in the forecast, causing another round of price changes (i.e., our adjustment period in the event study). This assumption is backed by empirical evidence such as Bernard and Thomas (1989) and Campbell et al. (2008).

In period 0 the agent has a non-ambiguous prior for dividends as $d \sim N(d_p, Vd_p)$. Abstracting from background wealth, and assuming risk neutrality, we can express the agent's utility function as $U(d) = E(d)$. In each period the demand function of the agent is $D = E(d) - p_0$. The market clearing condition is that in equilibrium excess demand is 0, so $p_0 = E(d)$. Under these assumptions the price in period 0 is $p_0 = d_p$.

Period 1: The agent receives a signal about dividends, $S = d_s + e$, where $S \sim N(d_s, Vd_s)$. The quality of the signal is depicted by its variance, Vd_s. High variance implies low quality (d_s is a weak predictor of d) and vice versa. In situations of pure risk the agent has complete knowledge of this variance. However, in ambiguous cases the variance can vary between a low and a high value, i.e., $Vd_s \in [Vd_{s,MIN}, \dots, Vd_{s,MAX}]$. In this case the agent cannot disentangle whether the signal is of high or low quality, which is the standard exposition of ambiguity (i.e., Epstein and Schneider 2008). We can further assume that the *true* precision of the signal (i.e., the one that the agent uses when the situation only involves risk) lies somewhere between the two extremes, i.e., $Vd_{s,MIN} < Vd_{s,TRUE} < Vd_{s,MAX}$. As explained in Epstein and Schneider (2008) in conditions of ambiguity the ambiguity averse agent updates his beliefs in response to the signal *pessimistically* (i.e., based on the worst case scenario). This means that if the signal conveys good news ($d_s > d_p$) the agent believes that the signal is relatively imprecise, so he updates believing that $Vd_s = Vd_{s,MAX}$. On the contrary, if the signal conveys bad news ($d_s < d_p$) he updates believing that the signal is very precise $Vd_s = Vd_{s,MIN}$.

Pricing in Period 1:

Assume for now that the signal is risky. The price in period 1 using Bayes rule is:

$p_1 = E(d|s) = \left(\frac{Vd_p^{-1}d_p + Vd_s^{-1}d_s}{Vd_p^{-1} + Vd_s^{-1}} \right)$, which is a precision weighted average of the priors, d_p and the new

information, d_s. The price change is equal to: $p_1 - p_0 = \left(\frac{Vd_p^{-1}d_p + Vd_s^{-1}d_s}{Vd_p^{-1} + Vd_s^{-1}} \right) - d_p = \frac{Vd_s^{-1}(d_s - d_p)}{Vd_p^{-1} + Vd_s^{-1}}$.

From this equation it is clear that bad news ($d_s < d_p$) cause a drop in the price and good news cause an increase in price. In addition, all else equal, the absolute price change is larger when the variance of the signal Vd_s reduces (i.e., when quality of he signal increases).

Pricing in Period 2:

In this period we assume that some of the uncertainty in the signal is reduced, so its variance decreases, $Vd_{s2} < Vd_s$. This causes a further update to the price equal to:

$p_2 - p_1 = \frac{Vd_s^{-1}[(d_s - d_p)(Vd_{s2}^{-1} - Vd_s^{-1})]}{(Vd_p^{-1} + Vd_s^{-1})(Vd_p^{-1} + Vd_{s2}^{-1})}$. In risky conditions the variance of the signal Vd_{s2} is always

smaller than Vd_s. (This is because we have assumed some resolution). Therefore, the second bracketed term in the numerator is always positive, so prices in the second period *always* move in the direction implied by the signal. That is, if the signal is bad news ($d_s < d_p$), the price change is negative and if it was good news news ($d_s > d_p$) the price change is positive. Note that this *is not* what we observe in the data as we observe an upward movement in prices after both good and bad news. Therefore, the risk-only framework cannot explain our findings.

Let's now assume that the signal was ambiguous, and that it brought bad news. Under the pessimism framework explained above the initial change in price is $p_1 - p_0 = \frac{Vd_{s,MIN}^{-1}(d_s - d_p)}{Vd_p^{-1} + Vd_{s,MIN}^{-1}}$. Note

that in this case Vd_s is replaced with $Vd_{s,MIN}$, which implies that the agent is acting as if the signal is more precise than it really is. This effectively increases the weight put on the signal, and sends the price lower, compared to the risk-only case. Again the second period brings some resolution, but this time to the ambiguity created by the signal. Such a resolution implies that both $Vd_{s,MAX}$ and $Vd_{s,MIN}$ are drawn closer to the true variance Vd_s reducing the amount of ambiguity. Therefore in period 2 pricing is driven by the new lower bound $Vd_{s,MIN2} > Vd_{s,MIN}$, which implies that the agent realizes that the

signal is not an informative as he though in the beginning. The change in prices is then: $p_2 - p_1 =$

$\frac{Vd_s^{-1}[(d_s-d_p)(Vd_{sMIN2}^{-1}-Vd_{sMIN}^{-1})]}{(Vd_p^{-1}+Vd_{sMIN}^{-1})(Vd_p^{-1}+Vd_{sMIN2}^{-1})}$. In this case, since both d_s - d_p and $Vd_{sMIN2}^{-1} - Vd_{sMIN}^{-1}$ are negative, so

the change in price is positive.

Let's now assume that the ambiguous signal brought good news. Under the pessimism the

initial change in price is $p_1 - p_0 = \frac{Vd_{s,MAX}^{-1}(d_s-d_p)}{Vd_p^{-1}+Vd_{s,MAX}^{-1}}$. Note that in this case Vd_s is replaced with $Vd_{s,MAX}$

,which implies that the agent is acting as if the signal is less precise than it really is. The second

period brings some resolution to the existing ambiguity, which implies that in period 2 pricing is

driven by the new upper bound $Vd_{s,MAX2} < Vd_{s,MAX}$. The change in prices is then: $p_2 - p_1 =$

$\frac{Vd_s^{-1}[(d_s-d_p)(Vd_{sMAX2}^{-1}-Vd_{sMAX}^{-1})]}{(Vd_p^{-1}+Vd_{sMAX}^{-1})(Vd_p^{-1}+Vd_{sMAX2}^{-1})}$. In this case, since both d_s - d_p and $Vd_{sMIN2}^{-1} - Vd_{sMIN}^{-1}$ are positive, the

change in price is also positive.

References

Ahn, D., Choi, S., Gale, D. and Karin, S., 'Estimating ambiguity aversion in a portfolio choice experiment', *UC Berkley Working Paper,* 2009.

Anderson, E. W., Ghysels, E., Juergens, J., L.,'The impact of risk and uncertainty on expected returns', *Journal of Financial Economics*, 2009, 94, 233-263.

Atiase, R.K., 'Predisclosure information, firm capitalization, and security price behavior around earnings announcements', *Journal of Accounting Research,* 1985, 21 -36?

Ball, R., Kothari, S., 'Security Returns Around Earnings Announcements', *The Accounting Review*, 1991, 66, 718-738?

Berkman, H., Dimitrov, V., Jain P.C., Koch, D.P., Tice, S., 'Sell on the news: differences of opinion, short-sales constraints, and returns around earnings announcements', *Journal of Financial Economics*, 2009, 92, 376-399.

Bonner, E. S., Walther, R. B. and Young, S. M., 'Sophistication-Related differences in investors models of the relative accuracy of analysts forecast revisions', *The Accounting Review*, Vol. 78, 2003, pp. 679-706.

Brown, L., 'How Important is Past Analyst Forecast Accuracy?', *Financial Analysts Journal*, Vol. 57, 2001, pp. 44-49.

Campbell, J., Y., Ramadorai, T., Schwartz, A., 'Caught on Tape: Intitutional trading, stock returns and earnings announcements', *Journal of Financial Economics*, forthcoming.

Camerer, C.F. and Kunreuther, H., 'Experimental Markets for Insurance', *Journal of Risk and Uncertainty*, Vol. 2, 1989, pp. 265-300.

Camerer, C.F. and Weber, M., 'Recent Developments in Modelling Preferences: Uncertainty and Ambiguity', *Journal of Risk and Uncertainty*, Vol. 5, 1992, pp. 325-370.

Capstaff, J., Paudyal, K. and Rees W., 'The Accuracy and Rationality of Earnings Forecasts by UK Analysts, *Journal of Business, Finance & Accounting*, Vol. 22, 1995, pp. 67-85.

Chen, Z. and Epstein, L. 'Ambiguity, Risk and Asset Returns in Continuous Time', *Econometrica*, Vol. 70, 2002, pp. 1403-1443.

Clement, M., 'Analyst forecast accuracy: Do Ability, Resources and Portfolio Complexity Matter?', *Journal of Accounting and Economics*, Vol. 27, 1999, pp. 285-303.

Clement, M. and Tse S., 'Do Investors Respond to Analysts' Forecast Revisions as if Accuracy is All that Matters?', *The Accounting Review*, Vol. 78, 2003, pp. 227-249.

Clement, M. and Tse S., 'Financial analysts characteristics and herding behaviour in forecasting,' *Journal of Finance*, Vol. LXI, 2005, 307-341.

Cohen, M., I. Gilboa, Y. J. Jaffray and D. Schmeidler (2000), 'An Experimental Study of Updating Ambiguous Beliefs', *Risk Decision and Policy*, Vol. 5, pp. 123-133.

Ellsberg, D., 'Risk, Ambiguity and the Savage Axioms', *Quarterly Journal of Economics*, Vol. 75, 1961, pp. 643-669.

Einhorn, J. H. and Hogarth, M. R., 'Ambiguity and Uncertainty in Probabilistic Inference', *Psychological Review*, Vol. 92, 1985, pp. 433-462.

Epstein, L. and Schneider, M., 'Ambiguity, Information Quality and Asset Pricing', *Journal of Finance,* Vol. 63, 2008, pp. 197-228.

Fama, F., E., and French, R., K., 'The Cross-Section of Expected Stock Returns', *Journal of Finance*, Vol. 47, 1992, 427-465.

Frisch, D. and Baron, J., 'Ambiguity and Rationality', *Journal of Behavioral Decision Making*, Vol. 1, 1988, pp. 149-157.

Gärdenfors, P. and Sahlin, N. E., 'Unreliable Probabilities, risk taking, and decision making', *Synthese*, Vol. 53, 1982, pp. 361-386.

Gärdenfors, P. and Sahlin, N. E., 'Decision Making with Unreliable Probabilities', *British Journal of Mathematical and Statistical Psychology*, Vol. 36, 1983, pp. 240-251.

Gilboa, I. and Schmeidler, D., 'Maxmin Expected Utility With Non-Unique Prior', *Journal of Mathematical Economics*, Vol. 18, 1989, pp. 141-153.

Givoly, D. and Lakonishok, J., 'The Information Content of Financial Analysts' Forecasts of Earnings', *Journal of Accounting and Economics*, Vol. 1, 1979, pp. 165-185.

Gleason, C. and Lee, C., 'Analyst Forecast Revisions and Market Price Discovery', *The Accounting Review*, Vol. 78, 2003, pp. 193-225.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D. and Camerer, C. F., 'Neural Systems Responding to Degrees of Uncertainty in Human Decision-Making', *Science*, Vol. 310, 2005, pp. 1680-1683.

Keren, G. and Gerritsen, E. M. L., 'On the Robustness and Possible Accounts of Ambiguity Aversion', *Acta Psychologica*, Vol. 103, 1999, pp. 149-172.

Keynes, J. M., A Treatise on Probability, second ed. (1948) (McMillan, London), 1921.

Knight, F.H., Risk, Uncertainty, and Profit (Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Company), 1921.

Livnat, J. and Mendenhall, R., 'Comparing the post-earnings announcement drift for surprises calculated from analyst and time series forecasts', *Journal of Accounting Research*, Vol. 44, 2006, pp. 177-205.

Matsumoto, A. D., 'Management's Incentives to Avoid Negative Earnings Surprises', *The Accounting Review*, Vol. 77, 2002, pp. 483-514.

Mikhail, B. M., Walther, R. B. and Willis, H. R., 'The effect of experience on security analyst underreaction', *Journal of Accounting and Economics*, Vol. 35, 2003, pp. 101-116.

Mikhail, B. M., Walther, R. B. and Willis, H. R., 'Do Security analysts exhibit persistent differences in stock picking ability?', *Journal of Financial Economics*, Vol. 74, 2004, pp. 67-91.

O'Brien, P. C., 'Forecast Accuracy of Individual Analysts in Nine Industries', *Journal of Accounting Research*, Vol. 28, 1990, pp. 286-304.

Olsen, A. R., and Troughton, H. G., 'Are risk Premium anomalies caused by ambiguity?', *Financial Analysts Journal*, Vol. 56, 2000, pp. 24-31.

Richardson, S., Teoh, H. S., Wysocki, D. P., 'The Walk-Down to Beatable Earnings Forecasts: The Role of Equity Issuance and Insider Trading Incentives', *Contemporary Accounting Research*, Vol. 21, 2004, pp. 885-924.

Sarin, R. K. and Weber, M., 'Effects of Ambiguity in Market Experiments', *Management Science*, Vol. 39, 1993, pp. 602-615.

Sorescu, S. and Subrahmanyam, A., 'The cross section of analyst recommendations', *Journal of Financial and Quantitative Analysis*, Vol. 41, 2006, pp. 139-168.

Stickel, S., 'Reputation and Performance among Security Analysts', *Journal of Finance*, Vol. 47, 1992, pp. 1811-1836.

Zhang, F., 'Information Uncertainty and Stock Returns', *Journal of Finance*, Vol. 55, 2006, pp. 105-136.

Table 1: Descriptive statistics for analyst forecasts

This table reports descriptive statistics. Panel A presents the distribution for some of the variables used. Panels B presents forecast frequencies partitioned by revision direction and company size. To assign companies in size quintiles in month $t+1$ we use the breakpoints from Kenneth French's website in month t.

Panel A: Descriptive statistics

Variable	Units of measurement	Mean	Q1	Median	Q3
Company experience	Months analyst follows firm	36.97	10	23	50
Broker size	Number of analysts employed in brokerage in year t	58.76	23	50	90
Companies	Number of companies followed by analyst in year t	17	12	16	21
Forecast frequency	Number of forecasts issued by analyst	3.65	2	3	4
Forecast horizon	Days between forecast and earnings announcement	68.56	49	79	88
Industries	Number of industries followed by analyst in year t	4.7	3	4	6
Revision magnitude	(Forecast-previous forecast)/previous forecast *100	-4.34%	-11.76%	-2.44%	5.55%
Company size decile	Using end of previous month breakpoints	6.5	4	7	9
Forecast error	(forecast - actual)/actual*100	-4.51%	-11.36%	-3.48%	0.00%

Panel B: Frequencies

By revision direction:	Rev > 0	Rev < 0	Total		
	114,409	150,554	264,963		

By size quintile:	Small	2	3	4	Large
Rev < 0	20,378	23,956	26,391	33,846	45,983
Rev > 0	11,388	15,245	19,487	26,969	41,320

Table 2: The Forecast accuracy model.

This table reports estimates from an OLS regression of forecast accuracy (defined in equation 1) on company experience (the years that analyst i has followed company j), broker size (the number of analysts employed in the brokerage of analyst i who issued the revision), companies (the number of companies followed by analyst i in year t), forecast frequency (the number of forecasts issued by analyst i for company j in year t and quarter q), days elapsed (the days that separate the forecast made by analyst i for company j in year t and quarter q with any other forecast by any analyst for the same company and period), forecast horizon (the days that separate the forecast made by analyst i for company j and the corresponding earnings announcement date), industries (the number of 4-digit SIC codes followed by analyst i in year t) and lag accuracy (the absolute error of the last forecast made by analyst i for company j in the previous quarter). These variables (except lag accuracy) are standardised according to equation 2. We cluster standard errors ate the firm level and use year fixed effects.

	Parameter	Estimate	p-value
intercept	a	0.609	<.0001
company experience	β1	-0.003	0.199
broker size	β2	0.002	0.199
Companies	β3	-0.009	0.001
forecast frequency	β4	0.006	0.003
days elapsed	β5	-0.045	<.0001
forecast horizon	β6	-0.155	<.0001
Industries	β7	-0.009	0.001
lag accuracy	β8	0.098	<.0001
Pr>F			<.0001

Table 3 : Regression of Returns on accuracy related variables.

This table reports estimates from an OLS regression of cumulative market adjusted return from trading days -1 to 1 (where date 0 is the date the forecast was issued) on the accuracy related variables defined in table 2. These variables are interacted with revision magnitude ((New value – old value)/old value). We cluster standard errors ate the firm level and use year fixed effects.

		Estimate	p-value
Intercept	a	-0.001	0.087
Rev Magn.	β1	0.046	<.0001
company experience*Rev Magn.	β2	-0.005	0.013
broker size*Rev Magn.	β3	0.007	0.0004
Companies*Rev Magn	β4	0.010	0.0007
forecast frequency*Rev Magn	β5	0.012	<.0001
days elapsed*Rev Magn	β6	-0.029	<.0001
forecast horizon*Rev Magn	β7	0.027	<.0001
Industries*Rev Magn	β8	-0.006	0.027
lag accuracy*Rev Magn	β 9	0.017	<.0001
Pr>F			<.0001

Table 4: Measuring ambiguity.

This table reports estimates from an OLS regression of actual accuracy (defined in equation 1) on expected accuracy (defined in equation 4), with dummy variables that indicate the size quintile of the company for which the forecast is issued. D1 equals to 1 if the size of the company belongs in the 2, 3 or 4[th] quintile, and 0 otherwise. D2 equals to 1 if the size of the company belongs in quintile 1 and 0 otherwise. We use cluster standard errors ate the fir, level and use year fixed effects.

	Parameter	Estimate	P-value
Intercept	a0	0.336	<.0001
Dmidsize (D1)	a1	0.039	0.001
Dsmall (D2)	a2	0.084	<.0001
expected Accuracy	θ1	0.716	<.0001
Expected accuracy*Dmidsize	θ2	-0.104	<.0001
Expected accuracy*Dsmall	θ3	-0.237	<.0001
Pr > F			<.0001

Table 5: Impact and adjustment period returns around analyst earnings forecasts

This table reports mean market adjusted returns partitioned by revision direction and size quintile. Panel A (B) presents returns in the impact (adjustment) period after upward (Rev>0) and downward (Rev<0) forecasts. Panel C presents mean returns around the earnings announcements (-5, 0) associated with upward and downward forecasts. The table provides *p*-values for the significance of the differentials calculated using standard errors adjusted for unequal variances.

Company Size	Panel A : (-2,2)		Panel B: (3,20)	
	Rev <0	Rev > 0	Rev <0	Rev > 0
1(Large)	-1.55	1.19	-0.39	0.05
2	-1.72	1.96	-0.06	0.13
3	-2.35	2.26	-0.03	0.13
4	-2.49	2.69	0.27	0.48
5(Small)	-2.53	2.47	0.42	0.77
Dif. 5-1	**-0.98**	**1.28**	**0.81**	**0.72**
p-value	**<.0001**	**<.0001**	**<.0001**	**<.0001**
Panel C: earnings announcements (-5,0)				
1(Large)	-0.12	0.51		
2	0.26	0.61		
3	0.57	0.68		
4	0.45	0.91		
5(Small)	0.37	1.06		
Dif. 5-1	**0.49**	**0.55**		
p-value	**<0.0001**	**<0.0001**		

Table 6: Multivariate analysis.

This table reports estimates from OLS regressions of market adjusted returns in the adjustment period (3,20) for downward (Panel A) and upward (Panel B) forecasts. The explanatory variables are: the inverse of the logarithm of company size (price 3 days prior to the forecast x shares outstanding), average bid ask spread 20 days prior to the announcement of the forecast, forecast horizon (the days that separate a forecast from the corresponding earnings announcement date), broker size (the number of analysts employed in the brokerage during the year the analyst has issued a forecast), forecast error (the difference between forecasted value and actual value, scaled by actual value) and revision magnitude (the absolute value of the difference between the two most recent forecasts issued by analyst i for company j during period t, scaled by the penultimate forecast). We use cluster standard errors ate the firm level and use year fixed effects.

Variable	Parameter	Panel A: Rev < 0		Panel B: Rev >0	
		Estimate	Pr > |t|	Estimate	Pr > |t|
intercept	a	-0.039	<.0001	-0.021	<.0001
1/log(size)	b1	0.079	0.0006	0.077	0.001
Bid/ask	b2	0.0001	0.902	-0.001	-1.18
Forecast Horizon	b3	0.0002	<.0001	0.0001	<.0001
Broker size	b4	0.000	0.637	0.000	0.471
Forecast error	b5	-0.0185	<.0001	-0.027	<.0001
Rev. magn.	b6	0.001	0.684	0.01	0.002
N			148,520		113,516
Pr > F			<.0001		<.0001