Relative Performance and Security Analysts' Career Outcomes

A thesis

submitted by

Andrew Vincent
Andrew.Vincent@tufts.edu

In partial fulfillment of the requirements for the degree of

Master of Science

in

Economics

Tufts University

May 2012

Advisor: Edward Kutsoati

Abstract

This paper examines the relationship between security analysts' career outcomes and their forecast accuracy. We hypothesize that analysts rely on information gained from private contacts within the management of the firms that they follow in order to make earnings forecasts. We test this by looking at the impact a change in the quality of the brokerage house that employs an analyst has on her forecast error. We find that a move up in quality leads to a lower forecast error, although the evidence in our paper does not provide clear proof that the lower forecast error is a causal effect of the change in quality of employer.

Acknowledgments:

I would like to thank the members of my thesis committee, Marcelo Bianconi, David Garman, and Edward Kutsoati for their support and guidance not only during my thesis, but throughout my career at Tufts.

Table of Contents:

I. Introduction

Financial analysts play an important role in markets. They aggregate information about the health of companies and, using their expertise, produce forecasts of the expected earnings per share for firms. These forecasts form the basis for buy, sell, or hold recommendations that many institutions and individual investors use for investment decisions. Because of the central role of forecasting, it is important to understand how analysts arrive at their conclusions as well as the underlying incentives they have to report their true beliefs.

There is a rich literature on many different aspects of financial analysts' decisions ranging from the statistical distributions of their predictions to the regulatory environment surrounding the forecasts and more. In this paper, we are mostly concerned with determining the main source of financial analysts' information. When making predictions, do they rely more on some theoretical modeling techniques, or on private information received from personal contacts within a firm?

In order to answer this question, we use career changes to set up a sort of natural experiment. Our basic theoretical question is when an analyst changes jobs, particularly when she moves from a lower quality brokerage house to a higher quality one, what happens to her accuracy? One would expect that if we observe a substantial increase in accuracy it would more likely be the result of increased access to the management within a firm rather than from some specific modeling technique. We wouldn't expect this increase from a modeling technique because the only thing changing is the move to a new brokerage house, and it is less likely that a high quality brokerage house would have access to some drastically improved techniques rather than more contacts compared to other brokerage houses.

In addition to this question, we try to determine some of the factors that lead to career changes (i.e. moving to high quality brokerage houses, or moving down to lower quality ones). This part of the analysis is particularly important because it helps us sort out the issue of causality in our main question. Basically, the analysts who are more accurate should be the ones who move up to higher quality brokerage houses. If this is the case, it will be difficult to determine if the increase in accuracy after the change in quality is a result of new connections as our hypothesis states or of an inherent difference in forecasting ability. In a later section we will discuss this issue at greater length.

We find that there is an effect of forecast error on the likelihood of an analyst moving up. It is a statistically significant relationship of a fairly small magnitude so it is difficult to say which way to causal relationship between error and moving up in quality goes. Interestingly, we do not find a statistically significant coefficient on error when we regress on the likelihood of moving down in quality.

The results of our main regressions show that if an analyst moves down in quality, there is a reduction in her forecast error. This suggests that there may be some benefit of access to new contacts when moving to a higher quality brokerage house. We also find a positive relationship between moving down in quality and forecast error. This relationship is much weaker, which suggests that it is unclear what happens to an analyst's access to private information when leaving a high quality brokerage house. Overall it seems that there are some benefits of access to new contacts in higher quality brokerage houses and those contacts may or may not follow analyst's if they move out of those brokerage houses.

The rest of the paper will follow in four sections. We continue a review of the literature related to financial analysts in section II. Next, we develop our hypothesis for all of our regressions in Section III. Section IV presents the data, including our set up for the various variables used in our analysis. In section V we discuss our results. We provide concluding remarks in section VI.

II. Literature Review

There is quite a large body of research related to financial analysts that covers many different areas. Some of the earliest work in the 1990s focused on questions related to the distributional characteristics of analysts' predictions (Sundaresh 2008). Those authors immediately saw the opportunity to research questions that relate to the incentives analysts face and the behavioral motivations behind their predictions.

One of the most relevant questions for our research that arose in the literature was whether analysts have a systematic optimism bias in their forecasts. If they do have an optimism bias, it suggests that they are using non-quantitative information in their analyses, which is an important concept in our research. Easterwood and Nutt (1999) examined what type of biases analysts have. Before their work there were many questions about how analysts incorporated new information into their forecasts. The authors found that analysts underreacted to negative news and overreacted to good news, which indicates a systematic optimism bias.

After this finding, some economists objected to the implication that this bias was necessarily irrational. Lim (2001) and Gu and Xue (2005) found rational explanations for the documented bias in forecasts. Gu and Xue suggest that extreme good news in earnings tends to be followed by greater uncertainty and that leads to more forecast optimism. However, they do not find any such similar explanation for analysts' underreaction to bad news. Lim takes a different approach to the question. He looks at a different type of model for forecast bias that suggests analysts face a trade off between optimism, which could lead to greater access to management, and forecast accuracy. He finds that empirical results match the model, and this suggests that the optimism bias may not be irrational as previous works had suggested, but instead a function of that trade off.

The most important feature of this debate over the optimism bias in analysts is that both sides accept that such a bias exists. This indicates that analysts exert some control over their analyses that need not come from some advanced technique. If the bias is, in fact, in alignment with Lim's suggested model, then it seems that access to management is potentially a highly important part of the information that an analyst use to make forecasts. This finding supports our hypothesis that moving to a high quality brokerage house, which comes with greater access to management, affects analyst performance.

Another important area in the research has to do with the concept of being bold or herding in one's forecasts. Most of the literature found that herding (i.e. deviating toward the mean in an analyst's forecasts) suggests inexperience, and less weight is put on those forecasts (Hong, et al., 2000). It also appears to be the case that inexperienced analysts, who are generally judged more harshly on forecast accuracy, understand this relationship and are much more likely to herd and make more revisions to their original forecasts. This concept of boldness extends to timeliness as well. Cooper et al. (2001) find that what they call lead analysts (those who post the earliest forecasts) tend to be more accurate and have a larger impact on the market. Much like the systematic optimism bias, these findings suggest that analysts are incorporating non-quantitative information in their forecasts.

The last component that is important for our research are the determinants of forecast accuracy and career outcomes for analysts. Some of the earliest research found that individual aptitude, experience, employer size, and the number of firms and industries followed all factor in to accuracy (Clement 1999, Jacob et al. 1999). Other research suggests that past accuracy works just as well at predicting future accuracy as a more complex model of individual characteristics (Brown 2001). The fact that brokerage house characteristics factor in to accuracy lends support our hypothesis of the effect of moving up on accuracy, and that Brown found that past accuracy is a good measure of future accuracy indicates that our results should not suffer from not having a wealth of information about the characteristics of individual analysts.

Perhaps the most influential paper for our research relates career concerns and biases in analysts' forecasts (Hong and Kubik 2003). They found that increased accuracy does increase the chance of better career changes, like moving up to a high quality brokerage house. They also find that being more optimistic increases the likelihood of positive changes, and that optimism can even become more important than accuracy if an analyst is following a stock that her brokerage house is underwriting. This all suggests that analysts can be rewarded for providing biased forecasts. This means that analysts may not be basing forecasts solely on rigorous, analytical information, and that employers may not care much if analysts are using non-relevant information in forecasts.

Taken all together, the evidence from the literature suggests that there is more to financial analysts's forecasts than one might think initially. There is certainly an important role for access to management in formulating predictions. There are also some strategic advantages to timing forecasts, adding revisions, and not revealing all available information. Our question about the extent to which non-quantitative information affects an analyst's performance arises naturally from all of these considerations.

III. Hypothesis Development

In our analysis we use two different types of models. First, we look at a simple probit model to determine some of the factors that contribute to an analyst either moving up to a higher quality brokerage house, or moving down to a lower quality one. The basic model was used was given by the equation,

$$Move_i = \lambda_o + \lambda_1 E_i + \lambda_2 Ten_i + \lambda_3 Ten_i^2 + \lambda_4 O_i + \lambda_5 R_i + \lambda_6 Time_i + \lambda_7 W_i + \lambda_8 W_i^2$$

where *Move* is a dummy where 0 represents the case of remaining in the same quality brokerage house as last period, and 1 represents the case of changing quality for an individual analyst. E is our measure of the forecast error, *Ten* is the number of years an analyst has been working in the industry, O is a measure of optimism, R is the number of revisions made to the original forecast, *Time* is the number of days between the initial forecast and the date earnings were reported, and W is the number of firms an analyst follows. We ran two separate probit models for the two cases: moving up in quality and moving down in quality.

First, let's discuss the variables in the case of moving from a low quality brokerage house to a high quality one. We define error such that a higher value means a greater deviation from actual realized value. This means that we expect there to be a negative coefficient on error. If an analyst has a higher error, that implies that she is less accurate. If she is less accurate then there is less value in her predictions and that should reduce the chance that a higher brokerage house would want to hire her.

Tenure is simply a measure of how many years an analyst has been working in the industry. We would expect that tenure would have a positive coefficient since spending more time in the industry proxies for experience, which higher quality employers would find valuable. However, there certainly are some diminishing returns to tenure. A person who has spent more time in the industry will probably demand a higher salary, and so one would expect to see a point where the marginal benefit of paying the higher wage to an older analyst may not be balanced by the marginal cost of that analyst. For this reason we have included tenure squared, and we expect to see a negative coefficient there.

Optimism is a measure of whether an analyst had a more positive forecast or a more negative one. It is calculated for an individual relative to other analysts following the same stock so a high optimism score simply means that she was more optimistic than the rest of the analysts. It does not necessarily mean that she predicted positive returns. There is a documented optimism bias in financial analysts where they are rewarded more for being optimistic. As such, we would expect that there would be a positive coefficient on optimism, however this need not be the case.

Confidence can be an important signal for an analyst to send to employers in the job market. We measure it in two ways. We count the average number of revisions an analyst made in a year. We expect that fewer revisions means higher confidence and so we expect to see a negative sign on that measure. However, the number of revisions could have a positive impact on the chance of moving up to a high quality brokerage house because more revisions may be interpreted as an analyst incorporating more information. If that were the case, then an employer may view more revisions as a positive and so we could observe a positive coefficient on the number of revisions.

We also look at the timeliness of the first estimate as measured by the number of days between the first estimate and the report date of earnings. We expect to see a positive coefficient on timeliness because the sooner an analyst makes an estimate, the more likely she is making her forecast based on private information. In other words, a later forecast may be incorporating information from other analysts and would therefore be less valuable.

The last variable in our probit model is the work load of an analyst. Work load is a measure of how many stocks an analyst follows in a given year. We expect a positive coefficient on work load. This would mean that analyzing more firms is a positive signal to employers of an analyst's ability, because it requires more skill to follow more stocks. However, one would expect that much like tenure there is some point when the marginal benefit of following another firm would be lower than the marginal cost. As such we have included work load squared and we expect to see a negative coefficient here.

Now we will examine the other probit model where we are looking at the factors that determine the likelihood of going from a high quality brokerage house to a low quality one. For our measure of error, we expect to see a positive coefficient because a higher error means a worse performance which means a greater likelihood of a negative career outcome. We expect to see a negative coefficient on tenure because more experience should reduce the chance of having to move down in the hierarchy of brokerage houses. The coefficient on the number of revisions could be either positive or negative because it is unclear whether more revisions would be interpreted as a positive or a negative signal to employers. If it were viewed as a positive signal, then we would expect to see a negative coefficient for it in this probit model. We expect to see a negative coefficient on timeliness. Lastly, work load's coefficient should be negative.

Moving on from our probit models, we next look at the regressions for our main question of how much analysts rely on private information via contacts within a firm compared to some analytical modeling technique. Again, we broke up our analysis into two sub samples: one with only those who were in low quality places and either stayed low or moved up, and those who were high and either stayed high or moved down. This is the basic model

$$E_i = \beta_0 + \beta_1 Move_i + \beta_2 Ten_i + \beta_3 Ten_i^2 + \beta_3 Time_i + \beta_4 W_i + \beta_5 W_i^2$$

The coefficient for our variable of whether or not an analyst moved to a brokerage house of different quality is the coefficient of interest. The other variables are to control for some of the other variation in error. We expect that if an analyst were to move up from a low quality brokerage house to a high quality one that we would observe a higher accuracy in her predictions. A higher accuracy is the same as lower error, so we would expect to see a negative coefficient. In the opposite case of a move down in quality we would expect to see a positive coefficient for error, although it would be expected that the magnitude would be smaller. We expect that there is less of a negative impact on contacts when moving down to a lower quality brokerage house than there is a positive impact when moving up in quality.

IV. Data

For our analysis we use the Institutional Broker Estimate System (I/B/E/S) data set. This is a panel data set that follows the predictions and performance of analysts across different brokerage houses every year. The analysts make predictions about the earnings per share for firms. For each prediction there are really three dates of interest, the date the prediction was made (called estimate date or revision date depending on whether it is the first estimate or a revision of said estimate), the date the forecast is for (called forecast date and it usually falls on the last day of the year), and the date the actual value is reported (called the report date).

In our analysis we compare analysts each year against their peers following the same stock, and then we average that performance across all stocks followed for each year. That is to say that for each analyst we end up with one observation for each year that they made a prediction. We also average optimism, and timeliness across stocks followed for each year.

A. Measure of Error

First, we will discuss our measure of forecast error. Initially we take the absolute value of the deviation of the error from the reported value given by

$$Err_{i,s} = \left|Forecast_{i,s} - Actual_{i,s}\right|$$

This gives us a measure of the magnitude of the error by an analyst for a given stock. Since analysts often revise their initial forecasts several times, we look only at this value for the last revision of the year. We choose to look only at the last revision because that should incorporate the full information set available to the analyst and therefore covers all of the factors that she uses to make her predictions, including all of the modeling techniques used and private information gained from contacts. We rule out forecast errors that are greater than $10 per share. Without doing so there are data entry errors that lead to extreme outliers (forecast errors up to $28,000 per share). There are a handful of observations that are dropped by this exclusion, and as such, there is no loss of generality of our results.

After we have the forecast error for each analyst for a given stock we rank all of the analysts following that stock. The highest ranked analyst (i.e. the one with a rank of 1) has the lowest error for that year and the lowest ranked has the highest error in absolute terms. Ranks are also generated such that if two analysts have the same error they receive the same rank. At this point we have a count for the number of analysts that are following that stock and we save that number as the popularity of a stock.

Because we are ranking individuals and want to compare against other analysts who are not following the same stock we need to control for the number of people following each stock. First, we drop all stocks that have fewer than five analysts following it. This rules out problems of punishing accurate analysts or helping inaccurate analysts that follow low popularity stocks. Next, we normalize the ranking using the following formula

$$E_{i,s} = \frac{(Rank_{i,s} - 1)}{(Popularity_{i,s} - 1)} \times 100$$

This gives us a number from 0 to 100 for all analysts for each stock they follow. This number can still be interpreted as a measure of error since the analyst with the lowest prediction error will have normalized score of 0, and as error increases so does the magnitude of normalized score to the lowest ranked analyst who receives a score of 100. This means that an analyst who falls right in the middle of the pack should have a score close to 50. One might consider using the deviation from the mean to rank analysts, and this approach has been used in the literature. However, we feel that simply using the ranking works just as well and may be better if there is some skewness in the distribution of errors attributed to some herding of analysts or excessively bold forecasts.

The last step in calculating the error that we use in our regressions is to take the average error for all of the forecasts done by each analyst for each year using the following formula.

$$Error_i = \frac{1}{NumStocks_i} \sum_{s \in S} E_{i,s}$$

That is to say that if an analyst has made forecasts on five different stocks in one year, we have five normalized scores for her for that year. We then take the mean of her normalized scores and call that her average error. If she were the best performer amongst all of the stocks she follows then she would have an average error of 0. Keep in mind that an average error of 0 does not mean that she actually made correct predictions on all of her stocks. Instead it means that she was the most accurate compared to her peers on all of her stocks. That does not rule out that she could have some positive or negative forecast error.

The reason we take the average of the normalized scores for each year for an analyst is that it gives us a measure of an analyst's overall performance. Surely brokerage houses make decisions about hiring and firing based partially on overall performance. In the literature, some have taken average rankings in a rolling average across three consecutive years to cut down the noisiness of the data. While there may be some information that signals the quality of an analyst to employers by looking farther back in her record than one year, employers definitely weigh the most recent year's performance more heavily than earlier years. As such, we feel that not including the previous two years as well does not impact our results.

B. Quality Measures and Changes in Quality

Next we will discuss our measures of the quality of brokerage houses and the movements of analysts between high and low quality employers. With the data available there is not a direct measure of quality. However, in the literature size is commonly used as a proxy for quality. Size is measured by the number of analysts working at a given brokerage house in a given year. By using size as a proxy for quality we are assuming that if a brokerage house employs more analysts, they must have more resources available to aid analysts' forecasts. One shortcoming of this measure is that it doesn't capture boutique brokerage houses, which are generally smaller, more highly focused brokerage houses of high quality in a specific industry. But without a direct measure of quality it really isn't possible to capture those brokerage houses, and fortunately they are more of an exception to the rule. It is also the case that larger brokerage houses generally cover a wider range of stocks (i.e. they may focus on more industries, or may be able to allocate analysts to unpopular stocks). Without direct information about the industries of the stocks that are followed, using a measure of the diversity of stocks as a proxy for quality would not be possible. It is also not clear that such a measure would be any better a measure of quality than size since they are both indirect measures of the unobservable characteristic of quality.

Once we have counted the number of analysts employed by each brokerage house, we then rank the brokerage houses. Again we are ranking within each year. We take the top ten largest to be considered high quality, so we assign those brokerage houses a value of 1 in a variable we called quality and assign 0 to all others. Because we used the top 10 each year, there is a potential problem of the top 10 changing each year. This problem could mean that in one year the brokerage house that is tenth gets bumped to eleventh the next year even though they may not have changed the number of employees. To fix this we go through the data set and checked whether a brokerage house got bumped out of the top 10. If they had been bumped and their size did not decrease by more than 10 employees, then we reassign them to high quality.

Using this dummy variable for the quality of a brokerage house, we can now begin to look at the movements of analysts between high and low quality employers. By each year we identified whether an analyst was employed by a high quality or low quality brokerage house. We defined a dummy variable to compare those analysts who remained in a low quality brokerage house to those who moved up. For all analysts in a low quality place we defined a variable for the move up to be 0. If an analyst is now in a high quality brokerage house and last year was in a low quality brokerage house, we set our move up variable to 1. In this way we are comparing those who moved up only to those who remained in the low quality brokerage houses. Similarly we defined another variable for the opposite case. In this case we set our move down variable to 0 if an analyst was in a high quality brokerage house and if they dropped in quality we set the move down variable equal to 1.

By defining the movements up and down in quality by separate dummy variables we allowed for a much simpler interpretation of our results. Also, it allows us to compare the results of a move with only the relevant subsample that acts as the control group. Another way to run this analysis would be to compare moves up and down to the subsample of those who changed employers. That is to say that one could look only at analysts who are working at a different brokerage house this year from last year. However, the interpretation of our results would not be as clear in this case. We want to see the effect of moving from low quality to high quality on forecast error, so the obvious comparison is to look at those who move up compared with those who stayed low. We are not interested in how the accuracy of analysts who move up compares to those who move in general because there could be numerous reasons why an analyst would chose to change brokerage houses. In other words, we do not need to exclude those who remain in the same job or at the same brokerage house for our comparison.

C. Other Control Variables

In our analysis we have several control variables in our main regressions as well as in our probit models. Here we will briefly discuss the design and the reasoning behind the variables we have included. First, we look at a measure of the relative optimism of analysts. We start by calculating the forecast error for each analyst for each year. Similar to our construction of error rankings, we used the last revision for each analyst for each stock. The difference here arises from the fact that we simply rank the error, not the absolute value of the error. So the number one ranked individual had the largest negative deviation from the actual value, and the lowest ranked individual had the greatest positive error. Obviously if all of the prediction errors were positive, then the person with the number one rank would be closest to zero on the positive side, and if all prediction errors were negative than the lowest ranked individual would have the error closest to zero on the negative side.

The next steps are identical to the procedure for our normalized error ranking. We take the optimism rankings and use the following formula to normalize ranks between 0 and 100.

$$Opt_{i,s} = \frac{(Rank_{i,s} - 1)}{(Popularity_{i,s} - 1)} \times 100$$

The analyst who was relatively the most pessimistic about a stock receives a score of 0 and the most optimistic receives a score of 100. A useful feature of this measure of optimism is that scoring close to a 100 does not necessarily mean that an analyst's prediction was a large positive error. It simply means that it was a more positive prediction than her peers following the same stock. So, if there were some shock in the market or some event that caused there to be an overall negative outlook for a firm, we can still look at the relative optimism of analysts, even if all predictions were negative.

The final step for our optimism measure was to take the mean of the optimism scores for each analyst for each year.

$$Optimism_i = \frac{1}{NumStocks_i} \sum_{s \in S} Opt_{i,s}$$

Exactly like our measure of forecast error, we now have one number for overall optimism for a given year for each analyst. Again, the benefit of this average measure is that we can see how optimistic an analyst was compared to all other analysts regardless of stocks followed, and regardless of what shocks may have affected the outlook for the whole market.

We include optimism in our probit model because we want to try to capture any things that might motivate an employer to hire a new analyst. There has been some documented optimism bias in analysts and analysts are sometimes rewarded for being more optimistic. As such, it seems worthwhile to include since being optimistic might improve the probability of a high quality brokerage house hiring an analyst from a lower quality one.

We also wanted to look at measures of confidence of analysts. One of the best measures for confidence would be to look at how many revisions an analyst makes to a forecast. We counted this by simply taking the number of times an analyst made a prediction about a stock and subtracting one for their initial forecast. This directly gives us the number of revisions. We then averaged across all stocks followed by an analyst and got her average number of revisions for the year.

We also include the number of days between an analyst's first forecast and the date that the company announced its earnings. We call this measure an analyst's timeliness, and it can also be viewed as a measure of confidence. The earlier an analyst makes her first prediction, the more confident she must be about her forecast. We also average timeliness of a prediction for each analyst across all stocks covered in a year. This again gives us an idea of how confident analysts were each year.

The last control variable that we use is a measure of how much stress an analyst was under. We look at the total number of stocks that an analyst followed each year. It could be the case that following too many stocks would negatively impact accuracy, and could also impact whether or not a higher quality brokerage house would want to hire an analyst.

V. Analysis

The first step in our analysis is to determine whether there is any difference in the error between those analysts who have moved and those who have remained in the same quality brokerage houses. The simplest way to see this information is to look at the summary statistics for error broken down by the different cases for moving. We look at the mean of the error for the four different cases: moving from a low quality to a high quality, remaining at low quality, moving from high quality to low quality, and remaining at high quality.

We see that the mean for those who moved up is 50.43, while the mean for those who remained low is about 51.80. This indicates that there is some difference in the mean error between these two groups. While this doesn't tell us about the impact of moving up on an analyst's accuracy, it does suggest that there is a meaningful distinction between these two subsamples. Similarly, the mean error of those analysts who move down in quality is about 50.46, compared to the mean error of those who stayed at a high quality brokerage house with a mean of 49.91. This again suggests that there is a valid distinction between these two groups with regards to their performance in terms of forecast error.

From these simple considerations we must deal with a serious concern. When an analyst gets hired by a high quality brokerage house and moves up in quality, the decision to be hired must be at least partially based on past accuracy. To the extent that past accuracy is an indicator of present accuracy, we have to worry about the causality in our models. Does forecast accuracy (or error as we measure it) determine whether an analyst moves up or down in quality, or does a move up or down in quality factor in to forecast accuracy?

The ideal solution to this problem would be to randomly assign analysts to high quality brokerage houses. In that case we would clearly know that the career move would have a direct effect on accuracy. Unfortunately it is not possible to set up such an experiment. One could also conceive of using regression discontinuity to sort out this problem. Basically, if one could determine whether an analyst barely got hired or didn't get hired by a high quality brokerage house then we could hypothesize that the hiring decisions around that border would be randomly assigned and we could get the local average treatment effect for moving up. Unfortunately that data is not available, so we have to try to directly measure whether accuracy impacts the likelihood of an analyst moving up in quality. If it does, then it is not clear that any impact that moving up in quality has on accuracy is not just a function of selection bias.

To sort out this issue, we begin by estimating a probit model to look at the factors that determine whether an analyst moves up or down in terms of the quality of her brokerage house. If we find that accuracy plays a large role in this change, then we would have a harder time identifying the impact of changes in quality on accuracy. We examine both cases of moving up in quality and moving down in quality to see if the same factors affect the likelihood of those moves.

For the case of moving up in quality, we regressed tenure, tenure squared, average error, average optimism, the average number of revisions, average timeliness, work load, and work load squared on our dummy variable for an analyst moving up. We find that the coefficients on the average optimism and the average number of revisions we not significant at a 10% confidence level. We also find that the coefficients on work load and on work load squared we not statistically significant. This is a somewhat surprising result for the number of revisions and an analyst's work load because we expected that those two variables would be a signal to employers of an analyst's abilities. It is also interesting that optimism does not affect the likelihood of moving up in quality since the literature suggests that optimism leads to better career outcomes. The most likely explanation for why we don't see an effect here comes from the fact that we average optimism for the whole year. It is probably the case that being optimistic in general is not important to brokerage houses, but that being optimistic for specific stocks only would be highly important. Averaging by year prevents us from seeing the impact of optimism from individual stocks.

We find that average timeliness was also significant at a 1% confidence level, however there appears to be a negative correlation between timeliness (i.e. the number of days between an analyst's first forecast and the report date) and the probability of moving up in quality. This is a bit unexpected because one would think that if timeliness were a proxy for overall confidence that it would have a positive impact on the likelihood of moving up. The interpretation of this negative coefficient is not clear. It could simply be that timeliness is not a good signal of confidence, or even that confidence is not a good signal of the information content of forecasts to employers. The reason why the interpretation is not clear is that the coefficient on average timeliness is very small (-0.00117). This suggests that increasing your timeliness by one day only reduces your chance of being hired by a high quality brokerage house is reduced by only .117%.

Next, we look at how tenure impacts the likelihood of a move up in quality. We find that tenure is positively correlated with moving up (0.266) and that it is significant at a 1% confidence level. This suggests that there is some benefit in terms of moving up in quality to remaining in the industry for another year. If an analyst works for an extra year she increases her chance of moving up in quality by 26.6%. We also find a statistically significant negative coefficient on tenure squared that suggests that there are diminishing returns to tenure. It would be interesting to see what parts of experience have the largest value for employers.

Finally, we find a negative coefficient of -0.00280 on average error which was significant at a 5% confidence level. We did expect that there would be a negative coefficient on error. It implies that as error increases, the likelihood of moving to a high quality brokerage house decreases. With a coefficient of -0.0028, it suggests that moving down by 1 ranking in 100 will reduce the chance of moving up to a high quality brokerage house by .28%.

The implications of these results on our main question of determining whether a change in quality effects error is not immediately clear. We can see that tenure has a strong impact on the likelihood of moving up, but our other controls either don't have a statistically significant coefficient (like optimism) or they have a small impact (like timeliness). We do find that there is an effect of increased forecast error decreasing the likelihood of moving up. This does cause some problems for our future analysis. Since forecast error does help determine the likelihood of moving up in quality, we cannot say as much about how a move up in quality will affect forecast error. While this is a fairly serious problem, we do have some further analyses that can help us to better understand this relationship and to determine the impact of moving up in quality on forecast error.

With our other probit model we run the same variables as before on on our dummy for a move down in quality. We find that the coefficients on average error, average optimism, and work load were all not significant at a 10% confidence level. It is somewhat surprising that error does not appear to be an important factor for whether or not an analyst will move down in quality of her brokerage house. It may be the case that there is some difference occurring between the two types of moves. That is to say that a move up must be a hiring decision, whereas a move down could be a firing or a quitting decision. One would expect that hiring is based at least partially on performance, and this is confirmed by a statistically significant coefficient on error in our first model. However, there is a difference between the decision to fire someone and the decision to quit. A move down could also be motivated by a firm that is low quality offering a job that is better for an individual. Because moving down in quality could be motivated by all of these different factors, it is not surprising that the coefficient on average error is not statistically significant.

Overall, we have gained some insight from both models. We see that error does impact the likelihood of moving up in quality. This does suggest that me might have an issue in determining the effect that moving up in quality has on an analyst's error. However, the magnitude of the effect in the probit model suggests that current error does not contribute much to career moves. This is further backed up by the statistically insignificant coefficient on error in the second model. From these results, we feel confident that we will have some meaningful results for our next regressions about the impact of changes in quality on forecast error.

Before continuing on to this new analysis, it is worthwhile to recall the motivation for the following regressions. We want to determine the factors that analysts use to make their forecasts. One can consider analyst information coming from two main classes of information: quantitative and qualitative sources. Quantitative sources would include things like modeling techniques, publicly available information, or other sources that have some mathematical foundations. We expect that such information is either already publicly available or would be difficult to keep secret within specific brokerage houses. Similarly, simply due to analysts changing employers and socializing with one another, modeling techniques would probably be exchanged with some regularity. However, the more qualitative sources like contacts within the management of firms could easily be controlled in terms of who has access to them, and could still carry a lot of useful information. We believe that things such as greater access to these qualitative sources would be more readily available to analysts in higher quality brokerage houses, and they make up a large portion of the information used by analysts when making earnings forecasts. By looking at changes in accuracy when an analyst moves up or down in terms of the quality of her brokerage house we can sort out whether or not qualitative private information is used by analysts.

To follow the logic of our argument, we run a regression of our dummy variable for moving up or down as well as some controls on average forecast error by year. We control for variation in error with variables for how long an analyst has been in the industry (tenure), how confident she was that year (average timeliness), and the number of stocks she followed (work load). In our first regression we look at only those analysts who were either in a low quality brokerage house or those who had just moved up in quality. We defined a dummy variable such that those who worked in a low quality place had a 0 and those who had just moved up had a 1. By defining our variable like that, we are able to compare the effect of moving up compared to those who remained at the same tier of employer.

We find that all of our coefficients in the regression are statistically significant at a 5% confidence level, and in fact all of our controls are statistically significant down to a 1% confidence level. We find that tenure and work load are negatively correlated with error. This means that following more stocks and being more experienced both lower forecast error. Somewhat surprisingly we find that average timeliness is positively correlated with error, which suggests that producing a forecast earlier does not help an analyst's accuracy. This is surprising since timeliness can be thought of as a proxy for confidence, and earlier forecasts are generally considered to have more information content than later forecasts since later forecasts can incorporate the estimates of earlier analysts' work. One possible explanation for this result is that we are using the last revision to measure error, and so any benefit of more informative forecasts from greater timeliness could be lost by the time an analyst produces her last revision. That is to say that the effect of timeliness might be confounded by measuring timeliness by the first forecast and then measuring error by the last forecast.

We also find that the coefficient on moving up in quality is significant at a 5% confidence level and is negatively correlated with error. This means that if an analyst was working in a low quality brokerage house last year and moved to a high quality on this year, they have a lower forecast error compared to those who remained with a low quality employer. The magnitude of this reduction in error (-2.490) is quite a bit larger than the reduction from remaining in the industry for an additional year (-.1.447) or from following an additional stock (-0.622). This does suggest that there is some large impact on error due to moving up in quality. That means that there is some benefit for analysts that comes immediately from being hired by a high quality brokerage house. As we have said before, this strongly suggests that analysts incorporate some qualitative information, such as contacts within a firm's management. This contributes positively to their ability to make forecasts and that also means that increasing the number or quality of contacts within an industry is valuable to analysts.

Now we consider what happens to accuracy when an analyst moves from a high quality brokerage house to a low quality one. If the increase in accuracy we saw during the move up is in fact from some more quantitative resources that are not transferrable by analysts when they move and not something like increased access to management, which would be easily transferrable between jobs (think how easy it is to move a rolodex), then we would expect to see an increase in error after moving down.

When we look at the results of our regression, we do not see any change in an analyst's error. We find that the coefficient on our dummy variable for a move down in quality (where those who stay high have a 0 and those who move down have a 1) is significant at a 10% confidence level and that the magnitude (-2.314) is about as large as the coefficient on moving up in quality. Again we find that all of our controls are statistically significant, as well as being the same sign and nearly the same magnitude as in our earlier regression.

There are a few possible interpretations of the move down in quality is not as statistically significant. One is similar to the argument we made earlier where we discussed the fact that a move down in our data could be caused by many different factors, and that each of those factors could have a different impact on error. For example, if an analyst were simply offered a better job by a low quality brokerage house they may decide to move down, and there may not be any drop in their forecasting accuracy. However, that case supports our idea that the benefits of working for a high quality brokerage house are easily transferrable when an analyst moves, which suggests that they are more likely to be from some qualitative sources like better contacts.

Another case would be that analysts move down in quality because they are let go from a high quality brokerage house for some reason that may or may not be related to performance. If we rule out cases where an analyst were fired for a non-performance based reason (those cases wouldn't affect an analyst's error upon a move down anyway), we can think that a brokerage house would fire an analyst for having too high an error. By the way we calculated error, that is the same as saying that an analyst is performing too poorly compared to her peers. So if an analyst were let go from a high quality firm and was then picked up by a low quality one, there is no reason to think that such a move would lead to an analyst performing any worse compared to her peers.

One last reason for moving down would be that an analyst quit for any number of reasons. If an analyst quit from a high quality brokerage house and left the entire industry, then she would not appear in the regression as having moved down in quality, so we do not have to worry about that case. Any other reason for quitting does not have any implications for why an analyst would experience a decrease in accuracy after the move. All of this together suggests that analysts don't experience a drop in accuracy because when they change employers they retain some characteristics (analytical skills, contacts, general experience, etc.) that move with them.

So we currently have some indication that moving up in quality may reduce forecast error and that moving down may not have as big of an effect of increasing error. This suggests that there may be some benefits that are gained at a high quality brokerage house that remain with an analyst when an analyst changes jobs. We still have to worry about determining the causal impact of our analysis. We run two more regressions for each subsample to help sort out these problems. First, we regress change in quality on the forward forecast error of an analyst. If we see a significant coefficient, it would indicate that there is some benefit to moving up that persists over time. However, we do not find a statistically significant coefficient on a move in either direction. Finally, we run regressions of change in quality on the change in forecast error. Again we do not find statistically significant coefficients on a move in either direction. These regressions do not provide us any more evidence that suggests that moving up in quality has a causal impact on forecast error.

From our main regressions we observe a decrease in error when analysts move up in quality, which suggests that analysts are seeing some benefit from such a move. Unfortunately our analysis does not provide with evidence that this lower error is caused by the move. We do see that there is a weaker effect when moving down in quality, which hints that we may be observing a causal effect, because analysts aren't losing their contacts when they move down. However, further analysis on the persistence of any causal impacts provided no positive evidence, and it also seems that moving does not impact the first difference of forecast error. All together this evidence only suggests that there is some difference in characteristics between analysts of different quality employers, but we cannot say anything definitive about the causal nature of such differences.

VI. Conclusion

Over the course of our analysis, we wanted to determine what kinds of information are valuable to analysts when making earnings forecasts. We thought that analysts might rely more on private contacts within the management of firms rather than on analytical tools unavailable to the public. This would suggest that analysts' skills come mainly from developing and maintaining lists of contacts rather than some special knowledge of advanced mathematical techniques.

To determine whether his hypothesis were true, we looked at what happened to the accuracy of analysts as they moved between high and low quality brokerage houses. We assumed that if our hypothesis were true, that high quality brokerage houses would have greater access to contacts for their analysts, which would in turn increase the accuracy of analysts once they moved into high quality jobs. We found that there was a statistically significant decrease in error (or an increase in accuracy) after analysts moved up. However, we found that a move down in quality did have a negative impact on forecast accuracy, although it is less statistically significant and of a smaller magnitude. This suggested that whatever benefits analysts received in high quality brokerage houses were retained if they moved out of those high quality positions. These two facts both suggest that a large part of the benefits received for working for a high quality brokerage house comes in the form of personal contacts that are valuable to analysts and will follow them throughout the remainder of their careers. So our results do weakly support our hypothesis that analysts rely on private contacts to make their earnings forecasts.

VII. References

Brown, L. (2001b). How important is past analyst forecast accuracy? *Financial Analysts Journal*, 57, 44–49.

Clement, M. (1999). Analyst forecast accuracy: do ability, resources, and portfolio complexity matter? *Journal of Accounting and Economics*, 27, 285–303.

Clement, M., & Tse, S. (2005). Financial analyst characteristics and herding behavior in forecasting. *Journal of Finance*, 60, 307–341.

Cooper, R., Day, T., & Lewis, C. (2001). Following the leader: A study of individual analysts' earnings forecasts. *Journal of Financial Economics*, 61, 383–416.

Easterwood, J., & Nutt, S. (1999). Inefficiency in analysts' earnings forecasts: systematic misreaction or systematic optimism? *Journal of Finance*, 54, 1777–1797.

Gu, Z., & Xue, J. (2005). Do analysts overreact to extreme good news in earnings? Working paper, Carnegie Mellon University.

Hong, H., & Kubik, J. (2003). Analyzing the analysts: Career concerns and biased earnings forecasts. *Journal of Finance*, 58, 313–351.

Hong, H., Kubik, J., & Solomon, D. (2000a). Security analysts' career concerns and herding of earnings forecasts. *Rand Journal of Economics*, 31, 121–144.

Jacob, J., Lys, T., & Neale, M. (1999). Expertise in forecasting performance of security analysts. Journal of Accounting and Economics, 28, 51–82.

Lim, T. (2001). Rationality and analysts' forecast bias. *Journal of Finance*, 56, 369–385.

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

VIII. Tables

Table 1. Summary Statistics By Change in Quality

Starting Quality	Change in Quality	Variable	Mean	Std. Dev.
Low Quality	Move Up	Avg. Error Score	50.43	16.04
		Avg. Optimism Score	50.53	15.90
		Avg. Num. Revisions	3.33	1.53
		Avg. Timeliness	238.28	111.24
		Work Load	3.77	4.43
	Stay Low	Avg. Error Score	51.80	18.11
		Avg. Optimism Score	50.64	17.61
		Avg. Num. Revisions	3.16	1.59
		Avg. Timeliness	244.47	110.28
		Work Load	2.79	3.27
High Quality	Move Down	Avg. Error Score	51.46	16.23
		Avg. Optimism Score	50.71	15.89
		Avg. Num. Revisions	3.15	1.41
		Avg. Timeliness	251.93	109.44
		Work Load	3.04	3.31
	Stay High	Avg. Error Score	49.91	17.71
		Avg. Optimism Score	49.49	17.02
		Avg. Num. Revisions	3.42	1.74
		Avg. Timeliness	249.04	109.76
		Work Load	3.05	3.60

Table 2. Probit models for the determinants of change in quality of brokerage house

Variables	Move Up	Move Down
Avg. Error Score	-0.00280**	0.00135
Tenure	0.266***	0.280***
Tenure Squared	-0.0168***	-0.0133***
Avg. Optimism Score	7.13e-05	0.00111
Avg. Num. Revisions	0.00409	-0.0741***
Avg. Timeliness	-0.00117***	-0.000995***
Work Load	0.0299	0.0265
Work Load Squared	-0.00112	-0.00281*
Constant	-2.242***	-1.942***
Observations	14,233	5,774

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 3. Regressions of change in quality on forecast error

Starting Quality	Variables	Avg. Error Score
Low Quality	Move Up	-2.490**
	Tenure	-1.447***
	Tenure Squared	0.0613***
	Avg. Timeliness	0.0127***
	Work Load	-0.622***
	Work Load Squared	0.0141***
	Constant	55.47***
	Observations	14,233
	R-squared	0.014
High Quality	Move Down	2.314*
	Tenure	-1.255***
	Tenure Squared	0.0452***
	Avg. Timeliness	0.0140***
	Work Load	-0.999***
	Work Load Squared	0.0489***
	Constant	53.13***
	Observations	5,774
	R-squared	0.015

*** p<0.01, ** p<0.05, * p<0.1

Table 4. Regressions of change in quality on one period forward forecast error

Starting Quality	Variables	Forward Avg. Error Score
Low Quality	Move Up	-0.586**
	Tenure	-0.583***
	Tenure Squared	0.022***
	Avg. Timeliness	-0.005***
	Work Load	-0.113***
	Work Load Squared	0.002***
	Constant	54.160***
	Observations	131,760
	R-squared	0.0074
High Quality	Move Down	1.508***
	Tenure	-0.503***
	Tenure Squared	0.018***
	Avg. Timeliness	-0.003***
	Work Load	0.176***
	Work Load Squared	0.004***
	Constant	52.189***
	Observations	66,340
	R-squared	0.0086

Table 5. Regressions of change in quality on change in forecast error

Variable	Difference in Error
Move Up	-0.189
Move Down	0.264

*** $p<0.01$, ** $p<0.05$, * $p<0.1$