ERASMUS UNIVERSITY ROTTERDAM
ERASMUS SCHOOL OF ECONOMICS
MSc Economics & Business
Master Specialisation Financial Economics

The New York Effect

An Analysis of New York Analysts' Performance

Abstract

This study examines the performance of New York-based equity analyst relative to US analysts located elsewhere. By running regressions of de-meaned absolute forecast errors on a host of analyst characteristics - using a sample of thousands of quarterly and fiscal year earnings forecasts issued in 2007 on 27 Dow Jones firm - I do not find evidence of an accuracy advantage for New York analysts. However, among the tercile with the most active analysts, I find that New York analysts are 14.3% more accurate than other analysts. This suggests a New York informational advantage through the gathering of private information, translating into better performance. Further I confirm that especially the analysts' experience and the size of their brokerage firm have a positive effect on their forecast accuracy. An approach whereby dividing the forecasts into several forecast periods proves to effectively control for the disparity of forecast horizons between different (analysts') forecasts.

Keywords: equity analysts, information asymmetry, geographic proximity, private information, de-meaned variables

Author:	M.E. van Laar
Student number:	312532
Thesis supervisor:	Dr. M. Gabarro Bonet
Finish date:	May 2016

ACKNOWLEDGEMENTS

First and foremost I would like to thank my parents for their support during all of my studies. I am very grateful for all they have done to help me finish this project. I also thank my supervisor Marc Gabarro Bonet for getting me started on my thesis, being patient, and his helpful suggestions and comments throughout. Finally, I am thankful to all of my friends and family for various support given.

1. Introduction

I investigate the role geographic location of equity analysts plays in relation to their performance. Specifically, I explore the impact of analysts' proximity to New York City on their ability to predict a company's earnings per share. Differences in forecast accuracy as a result of this proximity to New York may uncover information asymmetries between New York analysts and others.

The vast number of papers dedicated to analyst forecasts and stock recommendations in itself proves that they matter. The analysts' reports serve as an important tool for investors to assess a company's current and future performance. To illustrate this Beyer et al. (2010) show that analyst forecasts provide 22% of accounting-based information to investors.[1] In this light it is important what investors take into account when assessing analyst forecasts and making investment decisions. I aim to contribute to this by focusing on this particular group of analysts.

Previous papers have established a link between geographic distance and information flow. In the field of my focus - analyst earnings forecasts - a modest amount of papers investigate the role of geography. Most, however, focus on the difference of domestic versus foreign analysts. Depending on the markets that are the focus of their attention (e.g. emerging or developed markets), results are mixed. Malloy (2005) is the first to test a pure distance effect on a sample of U.S. equity analysts covering U.S. firms. He finds that analysts that are based geographically proximate to the company they follow are more accurate than others when it comes to forecasting earnings per share. I build on this research and try to find analyst accuracy advantages as a results of proximity to New York.

The findings in Malloy (2005) that geographically proximate analysts outperform their more distant counterparts are the starting point for my research. His findings are indicative of local analysts having an information advantage over other analysts. This train of thought has led to the main research question of this study; are New York based analysts more accurate in predicting earnings per share than other analysts? With New York arguably the financial capital of the United States (if not the world), there are reasons to believe this may be the case.

The main argument for superior performance of New York analysts is that in New York analysts could benefit from a superior flow of information. In this information flow story, information arrives faster in (or stems from) New York City. Analysts can also benefit from being close to providers of finance and financial services, which are abundantly present in New York. Furthermore, it is important to note that the majority of analyst forecasts is issued by New York-based analysts. This means that analysts in New York have an increased number of fellow analysts

[1] They do so by analysing cumulative abnormal returns on days on which accounting-based information becomes public. The five accounting-based information sources being: earnings announcements, earnings pre-announcements, management forecasts, analyst forecasts and SEC filings. Only management forecasts provide more accounting-based information than analyst forecasts (55%).

around that are following the same stock, with whom they are able to share information and discuss insights with one another. This could in effect improve the information environment for New York-based analysts relative to other analysts.

As a result, analysts located in New York may benefit from a superior flow of information, receiving first-hand and more timely information compared to other analysts. Therefore New York analysts would be able to issue more accurate earnings forecasts than analysts that are based outside the New York region. There are two ways to approach this hypothesised information advantage of New York analysts. Namely from the perspective of private information and cost of gathering information. The former builds on the assumption that New York analysts may benefit from absorbing the local culture and maintaining contact with local financial/business leaders, and thus are able to obtain information that other analysts cannot. The latter builds on the idea that (the same) information is simply easier to obtain for New York analysts as a result of an improved information environment, hence resulting in reduced cost of gathering information.

In my statistical research I use a methodology similar to the one used in Malloy (2005). This consists of running regressions of de-meaned absolute forecast errors on a host of analyst characteristics, including location based variables. Using a sample of five thousand quarterly and fiscal year forecasts issued in 2007 on 27 Dow Jones firm, I do not find evidence of an accuracy advantage for New York analysts. However, I do provide evidence that among specialised and the most active analysts, New Yorkers outperform others that are based more than 100 kilometres from New York City. This suggests a New York informational advantage through the gathering of private information. Additionally, among low-status and less active analysts I find that New Yorkers perform worse than their more distant counterparts. This is indicative of increased cost of information gathering. However, these results of worse New York performance are not robust to all model adjustments. In the end, only the better performance of New York analysts among high-coverage analysts proved to be robust to all model and sample alterations. This result showed that among the most active analysts New Yorkers are 14.3% more accurate than others, translating in a $0.014 per share accuracy advantage.

Of the other analyst characteristics, my results indicate that especially the analysts' experience and the size of their brokerage firm have a positive effect on their forecast accuracy. Finally, I also employ an alternative approach that effectively controls for the disparity of forecast horizons between different (analysts') forecasts.

The remainder of this paper is organised as follows. In section 2 I provide a review of the relevant scientific literature. In section 3 I present my research methodology and data in detail. Section 4-5 report my empirical results including some robustness checks. In section 6 I discuss shortcomings and opportunities for future research, before concluding. After that only the references and the appendix remain.

2. Literature review

In this section I discuss previous research in the field of equity analyst performance. The literature provides a multitude of papers that focus on the role of financial analysts in the allocation of economic resources. Far too many to discuss in this paper.[2] Therefore, I limit myself to those studies that researched factors of influence on analysts' forecast accuracy. I do so on the basis of the factors for which I control in my statistical research later on.

2.1 Geography in Analyst Performance

My research sets out to find a link between information flow and proximity to New York. Research into geography and information flow is not widespread and almost exclusively focuses on geographic proximity to company headquarters. Coval and Moskowitz (2001) present two arguments in support of information asymmetry dependent on geographic proximity to company headquarters. They do so in the context of mutual fund holdings and performance. They argue that the observed informational advantage as a result of geographic proximity results from both a decreased cost of information gathering and access to private information. Where lower cost of information expresses itself in the form of improved monitoring capabilities, visiting the firm's operations and better knowledge of the local market conditions. While the private information follows from the investors running in the same social circles as company executives. These arguments neatly apply to equity analysts, as Malloy (2005) argues in his research into the connection between an analyst's forecast accuracy and their proximity to company headquarters: "I would argue that the ability of local analysts to make house calls rather than conference calls, during which time they can meet CEOs face to face and survey the firm's operations directly, provides them with an opportunity to obtain valuable private information. Following this logic, geographic proximity is a sensible proxy for the quality of analyst information" (p. 721).

Malloy (2005) finds that, for a large US data set stretching from 1983 to the end of 2002, geographically proximate analysts are more accurate than other analysts when it comes to forecasting earnings per share. The 2.77% difference in accuracy translates into a $0.025 per share advantage that local analysts have over other analysts. While the magnitude of the accuracy advantage is larger in various subsamples, up to $0.141 per share. Besides this apparent better performance, Malloy (2005) also investigates the impact of forecast revisions on stock prices. In this regard he finds that both strong negative as well as strong positive local analyst revisions are met with an incremental average excess return in the three days surrounding the revision. -0.128%

[2] Ramnath, Rock and Shane (2008) give a very comprehensive overview of the financial analysts forecast literature.

and 0.095%, for negative and positive revisions respectively. Both these results indicate that local analysts have an informational advantage over other analysts as their forecast bring new information to the market.

Malloy (2005) also investigates how the remoteness of covered companies and the size of analysts' portfolios relate to this local effect. He places firms in three categories; Metro, Nonmetro and Remote, dependent on the firm being either located in one of the 20 most populated cities in the US, near one or further away. He finds that the magnitude of the accuracy advantage of local analysts is significantly larger for smaller stocks in remote areas. Additionally, the effect is not present among large and Metro stocks. Furthermore he finds that the local effect is strongest among analysts who cover a relatively large amount of (local) stocks. Suggesting that the local effect has more to do with private information rather than reduced cost of information gathering.

A small number of other papers investigate the relation between locality and analyst performance, all focusing on cross-border effects though. Research where performance is either measured by forecast accuracy or stock picking abilities (via stock recommendations). For example, Bae et al. (2008) investigate for 32 countries in the period 2001-2003 whether local analysts perform better than foreign analysts, where local simply means located in the same country as company headquarters. In all specifications of their model they find that the average forecast accuracy of local analysts exceeds the average forecast accuracy of foreign analysts. Ranging from 4.5% to even 9% more accurate than analysts from other continents. However the evidence in the international context is mixed. For instance, Bacmann and Bolliger (2001) find that in Latin American emerging markets foreign analysts are more accurate than domestic analysts.

2.1.1 Proximity to Financial Centres

Prior research into the effect financial centres might have on the performance of equity analysts is scarce. There is, however, one recent paper that documents on earnings forecasts by analysts working at brokerage firms headquartered in one of the three main economic centres in China (Beijing, Shanghai, and Shenzhen). In their research Bartholdy and Feng (2013) do not present unambiguous results about the performance of such analysts in comparison to other Chinese securities analysts. They find both better and worse performance in different sub periods. However they suspect that the superior performance may be caused by another informational advantage. In their sample period the Chinese government implemented big reforms and new regulation in the financial markets were introduced. As a consequence the Chinese government also had a large ownership in listed firms. Bartholdy and Feng (2013) argue that this may have been of influence to their results: "Being close to the regulators and 'owners' may lead to informational advantages" (Bartholdy and Feng, 2013 p. 81).

2.2 Forecast Horizon

Maybe the most important determinant of the accuracy of earnings forecasts is the forecast horizon. This forecast horizon denotes the amount of time between the issuance of an earnings forecast by the analyst and the earnings announcement by the firm. As the day of the announcement of the earnings approaches, analysts' forecasts tend to become more accurate. The logic behind this phenomenon is quite simple, to the extent that analysts have more information at their disposal as time passes. For example, when comparing fiscal year earnings forecasts issued at the start of the fiscal year with forecasts issued close to the announcement of the actuals. In the intermediate period earnings of the first three quarters are made public by the firm, bringing with it a stream of new information or confirmation of earlier expectations about the company. That is, taking in mind that the flow of information during the year is not limited to earnings announcements only.

In this context, Clement (1999) actually reports that relative absolute forecast errors increase by 0.35% for every day a forecast is issued further away from the announcement of the actual earnings. At the very least this indicates that careful control for forecast horizon is needed when comparing different forecasts. The most direct - and most common in prior research – way to incorporate such a control is to add the age of the forecast as an independent variable. A second option would be to pair up forecasts that have virtually the same forecast horizon. In the context of this paper, I would then pair up a with another forecast that immediately precedes it in time. Much like Bae et al. (2008) implement in their research.

2.3 Analyst Experience

Clement (1999) showed for the first time that an analyst's experience is positively associated with forecast accuracy. He argues that the analyst's experience - measured by the number of years he or she has been active - acts as an indicator of the analyst's ability and skill. The argument for experience being a good proxy for ability and skill is twofold. In his reasoning he makes use of the way the analyst labour market functions: "The analyst labor market is assumed to function as a tournament in which stronger performers continue, while the weaker performers are forced out of the profession. According to Milgrom and Roberts (1992), the only performance information needed or used in a tournament is the relative, ordinal information about who did better."(Clement 1999, p288). He then argues that this is similar to the way brokerage firms (and their clients) evaluate their analysts. From here it is not a stretch to assuming such evaluations affect an analyst's tenure. Logically one would then expect more experienced analysts to be more accurate, as they have been through this process repeatedly. The second reason for assuming experience would have a positive effect on forecast accuracy is more straightforward. Being that analysts simply gain skills and knowledge as the years go by. For instance, more experienced analysts might be better at identifying economic trends or analysing financial statements, but presumably also have a more extensive professional network.

Up till know I have only talked about general experience and the increase in skills that come with it. However, besides the general skills, analysts also acquire skills and knowledge that are specific to the firms they follow. This includes better interpretation of their earnings reports and more generally how the firm functions financially. Additionally, as the years pass by, an analyst may be able to gain better access to information through establishing contacts with company insiders. Although one has to be careful not to confuse this with the selective disclosure of information by companies, especially since the adoption of the Fair Disclosure Regulation in the autumn of 2000. However a good network with company management, or even suppliers and clients, can at the very least help the analysts to better interpret public information.

Clement (1999) shows that in a 1985-1994 sample the expected absolute forecast error is 0.40% lower for every additional year of general forecasting experience. The coefficient of the firm-specific experience is even stronger with absolute forecast errors expected to drop with 0.94% with every additional year the analysts has been covering a company.

2.4 Resources

The amount of resources analysts have at their disposal also warrants a control when investigating their performance. Again, Clement (1999) introduces the relation between forecast accuracy and the resources the analyst has at their firm. To test this he uses the size of the brokerage firms that employ the analysts. The reasoning being that larger brokerage firms are able to provide superior resources for their analysts. This can come in the form of more advanced data sets, but also having more (capable) associates and other employees that support their analysts. It is also suggested that analyst who work at large firms may have better access to private information from managers at the firms they follow. He finds that analysts employed by the 10% largest brokerage firms are 7.7% more accurate than analysts employed by the other 90%. This amounts to a $0.045 difference in forecast error between the two types of analysts.

Since then many researchers in this field have controlled for analyst resources. For instance, Malloy (2005) also finds a strong negative relation between absolute forecast error and a continuous variable measuring brokerage size. In his 1994-2001 sample the brokerage variable is strongly negative in all specifications of his model.

2.5 Analyst and Firm Reputation

The performance of analysts can also related to their reputation as Stickel (1992) shows. In this respect, high reputation analysts are analysts who secure a spot on the All-American Research Team, as published by *Institutional Investor* magazine. Stickel (1992) follows on a debate in the literature about the performance of All-American analysts. A position on the Research Team could be seen as a proxy for relative reputation and compensation. However there have been debates about its relation to performance, with critics stating that the yearly rankings are nothing more than "popularity contests". In his 1992 paper, Stickel shows that forecasts by analysts with a high reputation are indeed more accurate than others by 2.8 cents per share. With a sample average earnings per share of $2.81, yielding a 1% difference in performance.

Leone and Wu (2007), in a 1991-2000 sample, show that greater forecast accuracy increases the likelihood of analysts being ranked as an All-American. While at the same time poor forecast accuracy by an All-American analyst increases the likelihood that he or she will lose his or her ranking. More recently, Fang and Yasuda (2014) show similar results, although focusing on stock recommendations rather than earnings forecasts. Using 1994-2009 data, they show that trading on All-American analysts' buy and sell recommendations yields an annualised risk-adjusted return difference of 7% for institutional investors who have advance access to stock recommendations. But also individual investors who only have access to public information can yield an 4% return when quickly trading on recommendations by top-rank All-American analysts. Further Jackson (2005), focusing on the Australian equity market over the period 1992-2002, finds that the analyst's reputation is significantly positively related to relative accuracy the same year. Reputation is measured in a similar way by looking at analyst rankings. Finally, Malloy (2005) also finds All-Star analysts to be more accurate than other analysts (by 3.83%).

Malloy (2005) additionally controls for reputation of the brokerage firm an analyst works for. The brokerage firm rankings follow from the same magazine, listing the brokerage houses with the most All-Americans as "The Leaders". In his model Malloy (2005) finds a positive relation between forecast accuracy and brokerage firm status, stronger than the analyst's own status even. He finds forecast by analysts employed at high-status brokerage firms to be 5.82% more accurate than forecasts by other analysts. Hong and Kubik (2003) also use Institutional Investor's house rankings to measure firm status. They report that past forecast accuracy is positively related to career outcome. Where a move from a low-status firm to a high-status firm is seen as the positive outcome (as wages at high-status houses are substantially higher). Performance in the top 10% of the distribution increases the chances of moving up the ladder 62%, while it decreases the chance of a move down the ladder for analysts already at a high-status brokerage house by 32%.

2.6 Optimism

An analyst's optimism can also lead to observed differences in performance. In the context of this paper, it may be that New York analysts are more accurate than others simply because they are less optimistic. At least that would be the expected relation, as the literature tells us. DeBondt and Thaler (1990) report on a systematic optimism bias when it comes to analysts' earnings forecast. They do so by regressing actual changes of earnings per share on the forecasted change in earnings by analysts. In such a regression a positive intercept signals (unrealistic) optimism and, generally speaking, the closer the beta is to one, the more informative the forecast. If analysts were rational, one would expect the intercept to be equal to zero. A beta below (above) one is associated with forecasts being too extremely high (low). The deviation of beta from one could possibly be explained by either under- or overreaction to strong/weak earlier earnings performance or new information as DeBondt and Thaler (1990) put it. Although Abarbanell and Bernard (1992) argue that these too extreme forecasts, found at the start of the fiscal year, are not an overreaction to the prior year's earnings release.

Easterwood and Nutt (1999) actually find that analysts overreact to positive earnings surprises and underreact to negative earnings surprises, concluding that analysts appear to be systematically optimistically interpreting new information. For one-year ahead forecasts DeBondt and Thaler (1990) find that analyst forecasts are significantly optimistic (intercept -0.094) and too extreme (slope of 0.648).

One could argue, though, that the optimism bias in analyst forecasts is inherent to the way the industry functions: "Optimism bias also has a plausible agency interpretation. Many analysts work for brokerage houses that make money by encouraging trading. Since every customer is potentially interested in a buy recommendation, while only current stockholders (and a few willing to go short) are interested in sell recommendations, optimistic forecasts may be preferable. Indeed, it is well known that buy recommendations issued by brokerage houses greatly exceed sell recommendations" (DeBondt and Thaler 1990, p55). On the other end of the spectrum, negative analyst reports on a company's stock may put a strain on the relation with company executives and access to information.

Following Malloy (2005), in my statistical tests I use a control for overall optimism as well as an explicit control for the relative optimism specific to the stock in question. Additionally, I try to distinguish bias from informativeness, with respect to New York analysts and others.

2.7 Affiliation

Analysts are assumed to have an underwriter affiliation with the stock they cover if the brokerage house they work for is lead underwriter for a recent SEO or IPO of the covered firm. And agency problems may arise as a consequence. There are several reasons to assume that affiliated analysts would issue more favourable forecasts compared to other analysts. First off, concerning the relation between the brokerage house and the issuing firm, negative forecasts by the research department could put a strain on the relation between the two firms. Similarly, one could argue that when a brokerage firm is willing to be lead underwriter, it logically follows that they would have a positive view on the issuing company's prospects. Or, at least, showing the opposite in their earnings forecasts of the firm would seem counterintuitive to increasing their investment banking revenues. This argument is strengthened by focusing on the issuing firm's choice: "If issuers select underwriters on the basis of the favourableness of the terms underwriters offer and these terms are related to their analysts' views, then the chosen underwriters' analysts are more likely to have favourable views of issuing companies' prospects" (Lin and Mcnichols 1998, p. 102).

In their extensive research into underwriter relationships, Lin and McNichols (1998) find that stock recommendations by affiliated analysts are significantly more favourable than recommendations by unaffiliated analysts. And some evidence that five-year earnings growth forecasts are marginally higher (21.29% for affiliated versus 20.73% unaffiliated, at a one-tailed probability value of 0.10). However, they find there to be no difference between affiliated and unaffiliated analysts when it comes to current and subsequent year earnings forecasts. In particular, the mean EPS forecast as a percentage of price is not significantly different between the two groups of analysts. It is striking that the affiliation bias only seems to be present in the longer term analyst expectations and not in the short-term earnings forecasts.[3]

Although there may not be a bias in earnings forecasts, it could well be that affiliated analysts issue more accurate forecasts than others. The reason behind this could be the inherent information advantage affiliated analysts have through the underwriting relation with the stock. An information advantage that would arise from, for example, better access to the company's management. If not from superior information as a result of the extensive analysis, done prior to the decision to act as lead underwriter.

---

[3] Lin and Mcnichols (1998) offer two interpretations of the differing results between stock recommendations and earnings forecasts, in relation to analyst affiliation. First, issuing firms are more concerned about stock recommendations than earnings forecasts. That way recommendations between the two types of analysts may differ, but forecasts not necessarily. And on the other hand, they argue that that manipulation of an stock recommendation is more difficult for investors to detect than manipulation of an earnings forecast. It could therefore be less costly to an analyst to issue overoptimistic recommendations than overoptimistic earnings forecasts.

Malloy (2005) opts to include a control for underwriter affiliation in his forecast accuracy model. He finds that affiliated analysts are 3.99% more accurate in their current year earnings forecasts. The effect is no longer significant when using only 2-year ahead forecasts.

2.8 Coverage Decisions and Analyst Effort

Finally, which and the number stocks analysts cover can have an effect on their forecast accuracy. From Malloy (2005), among others, I identify three specific possible effects on performance. Namely focus, specialisation and effort. Starting with focus, it is hypothesised that when analysts cover fewer firms, they may benefit from the increased focus of their attention relative to analysts who cover many stocks. Possibly leading to superior performance. A similar story can be told about analyst specialisation. Analysts who are able to direct all their attention to one industry may perform better than more broadly focused analysts. Finally, a control for the analyst's effort is also warranted when analysing analysts' forecast accuracy. Where the analyst's forecast frequency serves as the proxy for the analyst's forecast effort, as introduced by Jacob et al. (1999). The reasoning behind this is rather straightforward; more forecasts issued per firm indicates higher effort by the analyst, they argue. And subsequently, they prove that analysts who showcase this higher effort, are relatively more accurate in their forecasts than other analysts. Contrary to Jacob et al. (1999), who measure firm-specific forecast effort, Klettke et al. (2014) control for general forecast effort. In this recent paper, Klettke et al. find that this measure has higher explanatory power regarding differences in forecast accuracy than does the firm-specific measure (7.48% vs. 2.92%).[4]

[4] "We argue that analysts who generally devote high [low] effort are likely to devote high [low] effort to a specific firm, although this might not be captured by a firm-specific measure. Our general effort measure helps to reduce the measurement error in proxies based on only one firm." (Klettke et al. 2014, p. 16).

3. Data and Methodology

To test whether New York based analysts are more accurate than others, I use the analysts' relative forecast errors. By comparing analysts' predictions of EPS with the reported numbers, I can evaluate the accuracy of the analysts in my sample. Using relative forecast errors as my dependent variable, I run several regressions in order to identify a potential significant difference in performance between New York analysts and others. Of course while controlling for the factors mentioned in the preceding section. In this section I describe all the sources of data and the methodology of my research. Firstly, I describe all data sources and explain how I use the found data to create my variables. After that, I outline my regression methodology using variables that are de-meaned by forecast-period means.

To avoid confusion, when I mention earnings forecasts by analysts hereafter, I am talking about analysts' estimations of what the earnings per share will turn out to be.[5] Further, I will use the term firm periods to refer to specific company earnings. Hence, each firm has five firm periods per year; four quarters with associated quarterly earnings, and the fifth is the entire fiscal year with accompanied earnings per share.

3.1 Data

The backbone of my data set is a collection of analysts' earnings forecasts of Dow Jones firms. My sample is restricted to forecasts made in the year 2007, as I solely have the information needed to pinpoint the analysts' locations for that particular year. I merge this sample with several other datasets to arrive at my final dataset with all the needed variables and controls. Through Wharton Research Data Services (WRDS) I can access Institutional Brokers' Estimate System (IBES). The IBES database holds earnings forecasts and recommendations by sell-side analysts for U.S. companies since 1976. I make use of the Detail History and Recommendation History files. The first has the earnings forecasts by individual analysts for all 27 Dow Jones firms, issued in the year 2007. The sample consists of both fiscal year earnings forecasts, as well as analyst forecasts of quarterly earnings. The fiscal year forecasts are only current year's forecasts, so that means forecasts that are made after the fiscal year has started. Quarterly earnings forecasts are also restricted to a maximum of one-year ahead, ranging from forecasts on current quarter earnings up to three quarters ahead.

Along with all the analysts' earnings projections the Detail History file contains the issue date of the forecasts, the actual earnings as announced by the companies and the date of those

[5] Though analysts write extensive reports on the firms they cover, the stock recommendations and, in this case, the dollar amount of the earnings forecasts are a good testing ground for quantitative research.

earnings announcements. Finally, there are several tickers for the forecasted stocks and indicators for the issuing analyst and his or her brokerage house. These can used to merge this data with other data sets and to identify individual analysts.

The first order of business is to identify the analysts, by their last name, and their brokerage firms. To do this I use another file downloaded from IBES containing all analyst stock recommendations in 2007, which also contains names of the analysts and brokerage firms.[6] Using the Analyst Masked Code that is in both the Detail History file and the Recommendation history file, I add the names and brokerage firms of the analysts to my sample.

3.1.1 Analyst's and Firm's Location

The analysts' locations are derived from area codes found in Nelson's Directory of Investment Research. Among other things, this directory has yearly information on which analysts cover a firm. Along with the names of all analysts that are known to follow a firm, it also lists the phone number of the office they work at. I use the 2007 issue, which uses data as of November 2007. This means I assume that the same area codes hold for the previous months of 2007, and they have not moved just prior to the data collection. Beside the odd move to another office within a large brokerage house, this assumption is not that much of a stretch.

Using the area codes from the phone numbers I can then determine each analyst's location.[7] I then use a file from the U.S. Census Bureau, containing latitudes and longitudes of all city centres in the U.S., to pinpoint the approximate location of each analyst. After several transformations in Excel I am able to calculate the physical distance between analysts and the New York financial district in Lower Manhattan. The distance from the analyst's location to the headquarters of the company in question is done similarly. Using the area code of the firm's headquarters in 2007[8], I determine the city where companies were based in 2007. All distances are measured as the crow flies.

[6] I also include stock recommendations issued in 2006 and 2008, as a means to identify more analysts. For these analysts I check, by hand, whether the brokerage house corresponds with Nelson's Directory. As it could be that they have moved on to another employer. Checking for analyst moves within 2007 is an unfeasible task. I can only rely on the Analyst Masked Code, which analysts take with them when they switch brokerage firms.

[7] Websites like http://www.areacodelocations.info/ are really helpful to associate the area codes with the correct cities.

[8] I use Nelson's Directory for this and verify these records on company websites.

3.2 Variable Construction & Other Data Sources

Many of my variables depend on information that comes with the forecasts from the IBES database. In this subsection I describe how I transform some of the data to arrive at the variables I can use for regression. If additional data sources are used, they are stated accordingly.

3.2.1 Location Dummies

Using the distance to New York and company headquarters, I construct several dummy variables. Firstly, I create a dummy variable that separates all New York analysts from others. This variable (NYC) is equal to one if an analyst is located within 100 km of the financial centre in Lower Manhattan, New York. By using such a dummy in my regression methodology I am able to identify a potential forecast differential between the two groups of analysts.

At first I planned on solely using this one proximity measure that could capture the difference between New Yorkers and non-New Yorkers. However, with the results of Malloy (2005) among others in mind - namely that local analysts outperform others - I additionally establish a distinction between two types New York analysts; 'normal' New York analysts and New York-local analysts. This distinction is needed because 10 out of the 27 firms in my final sample are headquartered in New York. Subsequently it follows that when allowing for only one type of New York analyst, results may be biased by a such a locality effect.

Analysts are considered 'normal' New York analysts (NY) if they are located within 100 km of the financial centre in Lower Manhattan, and the covered firm's headquarters is not within a 100 km radius. Whenever the firm covered is headquartered within 100 km of the analyst's location, the analyst is marked as a New York-local (NY_LOC). For both variables it holds that they are equal to one when the analyst is a New York(-local), and zero otherwise.[9] Essentially NY and NY_LOC are a rearrangement of the NYC dummy. I will use the pure New York effect dummy (NYC) and the two separate dummies alongside each other as interchangeable measures. Mainly because my main research question is whether New York analysts (all of them) outperform others. Nonetheless, I do not want to ignore the implications that findings about the effect of locality have on my research.

The last location-based dummy variable controls for local analysts based outside of New York. This dummy is equal to one when a non-New York analyst is located within 100 km of company headquarters of the particular stocks he or she covers, and zero otherwise. Finally, I also employ an alternative measure of proximity to New York, namely (the natural logarithm of 1 plus) the physical distance between analyst i and New York City (LOGD). The obvious downside of this

[9] Note that an analyst located in New York may be marked a New York analyst for one stock and a New York-local for another. The distinction is not made on an analyst basis, but is dependent on the location of the headquarters for the particular stock in question.

measure is that it is not possible to make the distinction between the two types of analysts, so it only serves as an alternative for the NYC variable.

3.2.2 Analyst and Brokerage House Reputation

Information about individual analyst's and brokerage house reputation is derived from *Institutional Investor* magazine. I create dummy variables for both, which are equal to one when status is high and zero otherwise. In the absence of the 2007 data, all individual All-Star nominations are from the October 2006 issue. This means that analysts included in the All-American Research Team in late 2006, remain marked as high-status analysts throughout my sample.[10]

For each of 71 industries the magazine names a First, Second and Third Team of All-American analysts, as well as in some cases runners-up. Following Malloy (2005), analysts in either of those four categories will be referred to as All-Americans (c.q. high-status analysts). The dummy variable for analyst status (STAR) is equal to one if the analyst is an All-American, and zero otherwise.

For the brokerage firm's status I create a dummy variable (HIGH) that is equal to one if the analyst's brokerage firm is in the top 5 of brokerage firms and zero otherwise. As the rankings are published in the same October issue of *Institutional Investor* magazine, I use the 2006 ranking for the first ten months of 2007. For November and December of 2007, I use the newly published ranking.[11]

3.2.3 Age

The age variable (AGE), that controls for differences of forecast horizon, is equal to (the natural logarithm of 1 plus) the number of days that is between the date the forecast is issued and the date the actual earnings are released by the firm. Using information that is available in the IBES Detail History file.

3.2.4 Size

The brokerage size variable (SIZE), which proxies for available resources, is equal to (the natural logarithm of 1 plus) the number of analysts that issue earnings forecasts in 2007 for the particular brokerage firm in question. To get to this number I have downloaded an IBES Detail History file containing all earnings forecasts made for U.S. firms in 2007. After identifying all brokerage firms in

[10] Ideally, the last two months of 2007 would use the rankings from the October 2007 issue. However this data was not available to me.
[11] There is one mutation. In the new ranking UBS drops out of the top five and is replaced by Merrill Lynch.

my sample, I was able to calculate how many unique analysts had issued forecasts for each brokerage house.

3.2.5 Experience

I adopt two controls for the analyst's experience; overall and firm-specific experience. I am interested in two things: how long the analyst has been working as an analyst, and how long he or she has been covering a particular stock. To this purpose I use the IBES Detail History database again. This time using a file containing the complete database of all forecasts made for U.S. firms, issuance dating as far back as 1982.

The general experience (EXP) is measured as (the natural logarithm of 1 plus) the number of years in which the analyst has issued at least one earnings forecast, for any stock. The firm-specific experience (FEXP) is similarly measured as (the natural logarithm of 1 plus) the number of years in which analyst *i* has issued at least one earnings forecast for firm *j*.

3.2.6 Optimism

To create the optimism variable I inspect all the forecasts analysts in my sample made in the year 2007, forecasts for firms that are not in my sample included. The statistic of interest is the overall tendency towards optimism the analyst exhibits. Following Malloy (2005) this optimism variable (PCTOP) is measured as the percentage of company earnings (quarterly and fiscal) for which an analyst's latest forecast was above actual earnings. The reason that only the analyst's last forecast for each firm's quarterly or fiscal year earnings is considered, is to prevent bias as a result of differing numbers of forecast revisions among analysts and firm periods.

3.2.7 Affiliation

To control for underwriting affiliations I need information about IPOs and SEOs of all firms in my sample. More specifically, which brokerage firms acted as lead underwriter for those offerings. Via the Thomson One database I obtained information of all such offerings between 2005 and 2007 by the 27 companies in my sample, including subsidiaries.

Whenever a brokerage firm has been lead underwriter for an IPO or SEO in that three-year period, that brokerage firm is defined as having an underwriting relationship with the issuing company. The dummy variable for underwriting affiliation (AFFIL) is equal to one if the analyst's brokerage firm has an underwriting relationship with the stock on which he or she issues an earnings forecast, and zero otherwise.

3.2.8 Analyst's Coverage

After initial baseline regressions, I also investigate whether the companies an analyst chooses to cover are of influence on the observed results. Specifically, I am interested in the total number of companies an analyst covers and whether some analysts are more specialised than others. For the former, I use the same IBES file with all forecasts issued in 2007 by the analysts in my sample, and for each analyst I compute the number of unique firms they have issued at least one forecast for. The latter demands some additional information on how to define analyst specialisation.

Using the Center for Research in Security Prices (CRSP) U.S. stock database, I add the corresponding Standard Industrial Classification (SIC) codes to every company that is covered my sample. Following Malloy (2005) I use the first two digits of the SIC code, indicating to which major industry group a company belongs. Analyst are deemed Expert analysts when they cover firms in just a single major industry group. Accordingly, the dummy variable for analyst specialisation (SPEC) is equal to one if analysts only cover firms in a single major industry group, and zero otherwise. The variable for the number of firms covered (COV), which proxies for workload and (relative) focus, is measured as (the natural logarithm of 1 plus) the number of unique firms the analyst covers in 2007.

Finally, to be able to control for the general effort of the analyst, I calculate how many forecasts an analyst issues per firm covered. From this, the average number of forecasts an analyst issues per firm is derived. Such that the analyst's general effort (FREQ) is measured as (the natural logarithm of 1 plus) the mean number of forecasts issued per covered company.

3.3 Methodology

In order to compare (groups of) analysts I measure performance by looking at the analysts' capabilities to predict earnings per share. To correctly test this I compute the relative forecast error of each forecast, where smaller forecast errors obviously indicate superior performance. Throughout this paper terms like analyst's performance, forecast accuracy and relative forecast errors are used alongside of each other. They all serve the same purpose, namely to point out differences between analysts' capabilities to predict company earnings. I solely investigate whether different groups of analysts have significantly lower relative forecast errors. All the while controlling for other characteristics that may be of influence on these errors. Intuitively, the perfect forecast has a forecast error equal to zero, when forecasted earnings are equal to actual earnings. Terms like superior forecast accuracy or analyst performance, hence, are derived from a difference in relative forecast errors between analysts.

In this subsection, I describe my research design in detail.[12] First I explain my use of forecast periods. After that I outline the regression methodology where I use forecast period means instead of firm period means to de-mean the variables.

3.3.1 Seasonality

To a great extent, the methodology I apply is analogous to the one used by Malloy (2005). As stated earlier, this paper was the starting point of my research and its design fits my research question very well. The most obvious change I make is the addition of New York variables, though I also add several other elements and leave some out as well. A less apparent deviation from his framework, is my use of forecast periods.

My sample consists of over 8000 analyst forecasts and forecast revisions issued in the year 2007. By revising their forecast analysts adjust their forecast as more information becomes available. Along with Malloy (2005), in most prior forecast accuracy papers researchers have opted to only use the analyst's latest forecast before earnings announcement. This is done to avoid potential problems caused by autocorrelation in the analyst's revisions. Such autocorrelation is suggested to be the result of an analyst's delayed response to an economic trend impacting firm earnings. For instance, imagine an analyst taking several revisions to incorporate the full magnitude of such a trend in their forecasts.

However, I am working with a small sample of 27 firms and can only allow for one year of forecasts. And by dismissing all but the last forecast before the earnings release, I would decrease the size of my sample considerably. Furthermore I feel that analysts that revise their forecasts irregularly may have a considerable effect on the results in small sample. To the extent that the control variable that captures the forecast horizon of forecasts (AGE) is not able to effectively correct for the amount of information that is available to analysts at different points in time. In other words, it is possible that the differences in forecast accuracy caused by big forecast horizon disparity cannot be adequately explained by the overall trend of increasing accuracy with the passing of time. To overcome this, I develop an alternative method to control for the amount of available information at different points in time.

To this purpose I identify multiple forecast periods in which there is a level playing field among analysts with regards to the amount of publicly available information. All forecasts are then assigned to these different forecast periods. From each forecast period I only use the analyst's most recent forecast issued. The most obvious advantage is that the decrease of my sample size is relatively low, and I overcome some of the stated issues. Furthermore, using multiple forecast

[12] Unless stated otherwise, all modifications to my data set, and the statistical tests later on, are performed in the statistical software package Stata.

periods opens up the possibility to investigate a time component to the hypothesised New York effect, and inherent information flow.

3.3.2 Forecast Periods

Following Stickel (1990) and Agrawel, Chadha and Chen (2006) I use the pattern of forecast issuance to sort the forecasts into periods. The former divides a six-year sample in 144 semi-monthly periods in his model to predict individual analyst forecasts. While the latter use an early and a late forecast period in their research into the impact of the Fair Disclosure Regulation (2000) on the accuracy and dispersion of analysts' earnings forecasts. Keeping this in mind, I take a closer look at the timing of forecast issuance.

Figure 1: The Pattern of Forecast Issuance by Forecast Horizon

This figure shows the timing of the forecast issuance relative to the announcement date of the actual earnings. It shows the number of forecasts issued in 10-day windows of forecast horizon. Clearly there are four spikes around 90, 180, 270 and 360 days prior to announcement. This uses information on 7456 forecasts by identified analysts.

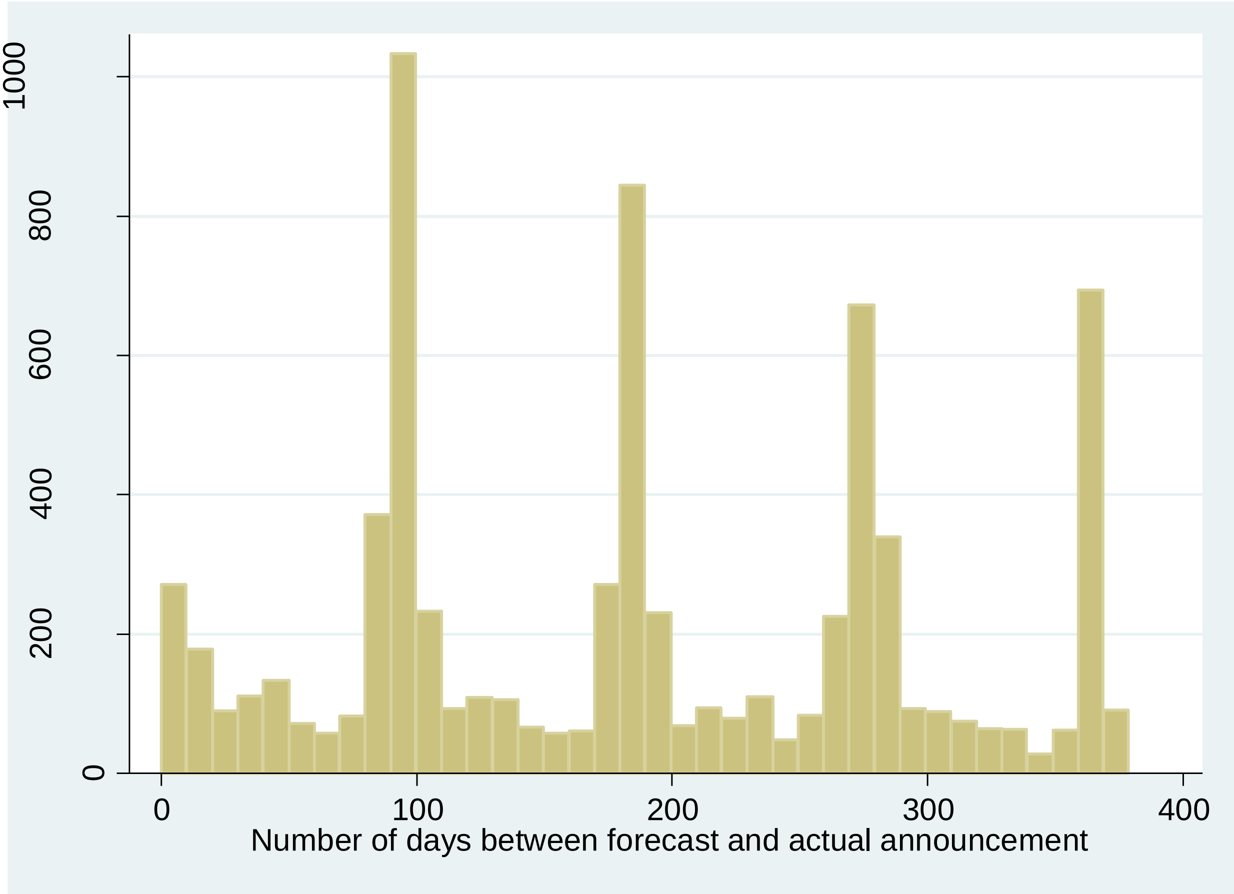

In my sample, I also identify a clear seasonality in the issuance of forecasts by analysts. I recognise a broad pattern in the timing of forecast issuance that I can use to group certain forecasts together. Figure 1 illustrates this pattern, with four pronounced spikes in the number of forecasts issued. As the analysts in my sample usually submit forecasts for several firm periods of the same firm at the same time, the same pattern applies to both quarterly and fiscal year earnings forecasts.

Further inspection of my sample learns that the pattern shows analysts releasing their first forecast approximately one year prior to earnings announcement. This is right at the start of the fiscal year for fiscal year and 4th quarter forecasts. By that time the previous year's earnings have been announced as well. The first forecasts on the 1st, 2nd and 3rd quarters are also generally issued one year ahead.

In addition to issuing a forecast for a new firm period every three months, analysts also revise their earlier forecasts for other firm periods at that time. These three-monthly revisions take place around the time the companies announce the earnings of the fiscal period that has just ended. There is an obvious explanation for the observed pattern, however, as spikes in forecasts issuance coincide with the four instances per year companies announce earnings. In conclusion, the data shows a clear pattern where every three months analysts submit a forecast for a new period, whilst at the same time updating their earlier forecasts. As a result, when it comes to forecast horizon, a large portion of forecasts concentrate around the 90, 180, 270 and 360-day mark (see Figure 1).

Given this pattern, I use five forecast periods: Early Forecast Period, 1st Revision Period, 2nd Revision Period, 3rd Revision Period and Late Forecast Period. To demonstrate this, I use an example of fiscal year forecasts on a firm which fiscal year ends on the 31st of December. Following Agrawal, Chadha and Chen (2006) the Late Forecast Period consists of forecasts issued within two months of the earnings announcement. On the other side of the spectrum, the Early Forecast Period has forecasts with a forecast horizon of 330 days or more. This roughly comes down to forecasts that are issued in the first two months of the fiscal year, and well before the earnings of the first quarter are announced. The three revision periods are formed around the announcement dates of Q1, Q2 and Q3 earnings.

Figure 2: Timeline of Forecast Periods

This figure depicts how forecast are generally divided among the five forecast periods. The number of days indicates the forecast horizon of the forecast. This is an example of how 4th quarter and fiscal year forecasts are distributed when a company's fiscal year ends on the 31st of December. For other quarterly earnings and/or companies only the dates change.

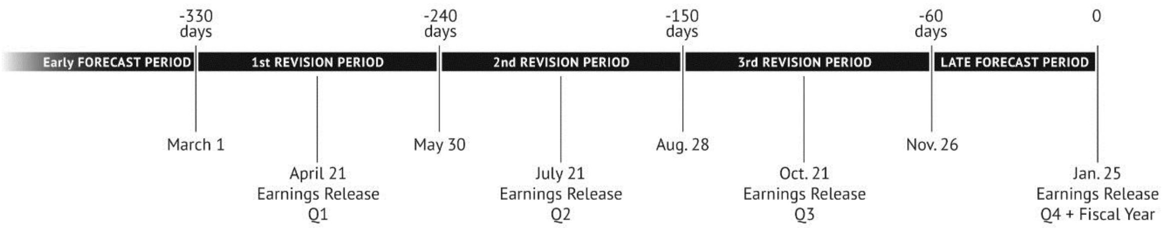

Figure 2 gives a graphical overview. The scale shows the range of forecast horizons for each of the five forecast periods. The announcement dates follow from the average number of days between the end of a firm period and the announcement of earnings by the firm. In my sample these are 21 days for quarterly earnings and 25 days for fiscal year earnings. These forecast periods are intended to create a level playing field in regards to publicly available information when comparing earnings forecasts. With this in mind I only use an analyst's last forecast per forecast period, to limit potential autocorrelation in forecast revisions.

3.3.3 Dependent Variables

One of the most important attributes of the framework used by Malloy (2005) is the control for firm-year effects. In this paper, with the presence of quarterly fiscal year earnings, this would become a control for firm-period effects. However, since I use multiple forecast periods for each firm period, I control for forecast period effects. In the forecast period effect lies information about how easy (or hard) certain earnings are to predict. Corporate events like voluntary management disclosures, mergers and strikes are examples of factors that can be of influence. By controlling for forecast period effects I am effectively able to compare forecasts on the same firm period that are issued in the same forecast period.

The forecast period control is achieved by comparing the forecast error to the average forecast error made by all analysts in the same time window. Following Malloy (2005), I compute the absolute forecast error as a percentage of stock price 12 months prior to the beginning of the firm period (AFE_{ijt}). Specifically, AFE is equal to the absolute value of an analyst's forecasted earnings minus actual company earnings, as a percentage of the stock price. I use these values to construct the dependent variables[13].

The next step is to calculate the mean absolute forecast error ($MAFE_{jt}$). Where MAFE is equal to the mean absolute forecast error for firm *j* for forecast period *t*. This average error thus captures the forecast-period effects one wants to control for. Using AFE and MAFE, I create two relative forecast error variables. Namely the de-meaned absolute forecast error ($DAFE_{ijt}$) and the proportional mean forecast error ($PMAFE_{ijt}$). DAFE is equal to the absolute forecast error (AFE_{ijt}), minus the mean absolute forecast error ($MAFE_{jt}$). While PMAFE is equal to DAFE divided by MAFE. DAFE tells us by how much an analyst's forecast is more (or less) precise than the average forecast issued in that forecast period for a particular firm period. Where negative values of DAFE indicate lower than average forecast errors, and better than average performance. While positive values represent worse than average performance. PMAFE adds another dimension to this. By deflating DAFE by MAFE, it tells us how much an analyst's forecast error is below (or above) the mean forecast error, in terms of that same mean forecast error. As a consequence, values of PMAFE are limited to -1 on the downside, as the perfect forecast (forecast error equal to zero) is always one

[13] All forecast error variables depicted in capitals are calculated as a percentage of stock price.

mean below the mean absolute forecast error. However, there is no such restriction on the other side of the spectrum. Because, strictly speaking, a forecast error can be any multiple of the mean forecast error for some particular earnings per share.

The advantage of the PMAFE specification is how it relates the forecast error to the mean forecast error, also making the coefficient more intuitive then the DAFE specification. Furthermore, as Clement (1998) shows that firms with large earnings per share have higher variation in their DAFE's, deflating DAFE by MAFE can also reduce heteroscedasticity.

3.3.4 Restrictions

Following the seasonality pattern, I only use each analyst's latest forecast per forecast period. This means that all but the analyst's last forecast issued in each forecast period will be dropped from the sample. Subsequently, I employ some further restrictions to arrive at the final sample. Following Clement (1999)[14], I control for analysts that appear to be mimicking the forecasts by other analysts and analysts that are not actively following a stock. Effectively this means that I drop instances when an analyst only issues forecasts (for a certain firm period) with forecast horizons less than 30 days, or when he or she only issues forecast in the early forecast period.

Furthermore, I drop three out of thirty Dow Jones firms from my sample to end up with 27 firms in total. General Motors, Citigroup and American International Group are deleted from the sample because forecast errors for these firms are extremely high and as a results not representative. The average forecast error of these three firms combined amounts to 34.26% of stock price, compared to 0.27% of stock price for the other 27 firms.

3.3.5 Outliers

I do not winsorise DAFE or PMAFE because observations that one might categorise as outliers actually hold valuable information. Careful inspection of these observations made it clear that most of the 'extreme' forecast errors are made by a small number of analysts. As such, I rule out the possibility of data entry errors in IBES and feel that information about these 'bad' analysts is relevant and should be kept intact. Nonetheless, as a robustness check, I also run tests using a sample where these observations are omitted from the sample.

[14] "The 11-month requirement is imposed based on the assumption that active analysts would supply forecasts for the firms they follow during this period. An analyst who only releases forecasts more than 12 months prior to period end is not likely to be following companies very closely. Similarly, an analyst who only releases forecasts less than 30 days prior to period end is more likely to be mimicking the forecasts of other analysts rather than following the companies himself." (Clement 1999, p.292).

3.4 Regression Model

To test my hypotheses I run numerous pooled, cross-sectional regressions. The regression as depicted in Eq. (1) serves as a baseline specification. To correctly estimate the model all independent variables are also forecast-period mean-adjusted. This way all independent variable values are relative to the mean value for its particular forecast period (i.e. relative to the same observations as the dependent variables). As a consequence of solely using mean adjusted variables, no intercept is needed in the regression equations.[15] As such the baseline regression equation takes the form of:

$$\begin{aligned} \text{PMAFE}_{ijt} \text{ (or DAFE}_{ijt}) = {} & \beta_1 \text{ DLOC}_{ijt} + \beta_2 \text{ DAGE}_{ijt} + \beta_3 \text{ DSIZE}_{ijt} \\ & + \beta_4 \text{ DSTAR}_{ijt} + \beta_5 \text{ DHIGH}_{ijt} + \beta_6 \text{ DEXP}_{ijt} \\ & + \beta_7 \text{ DFEXP}_{ijt} + \beta_8 \text{ DPCTOP}_{ijt} \\ & + \beta_9 \text{ DAFFIL}_{ijt} + \beta_{10} \text{ DNYC}_{ijt} + \varepsilon_{ijt}, \end{aligned} \tag{1}$$

where the D preceding each variables stands for de-meaned.[16] And where, in other setups, the NYC variable may be replaced by either the log distance variable or the two separate New York dummy variables. I use different specifications of this model, as well as adding other variables to it. This baseline regression equation is also used when running robustness test.

[15] A constant is necessary in Stata though, to make use of robust errors. All intercept coefficients in these regressions are logically equal to zero.
[16] See appendix for an overview of variable definitions.

4. Results

In this section I present the results of my statistical research. First I will describe my final sample in more detail. After that I proceed with the results of running the baseline model. Subsequently, I run further tests using alternate models and sample. Specifically, I look at bias and informativeness in forecasts, add controls for analysts' coverage decisions and do a breakdown of the sample by certain characteristics. Finally, I also do some robustness checks.

4.1 Descriptive Statistics

Originally I started off with over thirteen thousand forecasts provided by the IBES database. After merging those with many other data sources and imposing restrictions, the final sample consists of 4985 forecasts of fiscal year and quarterly earnings. These are forecasts on the earnings of 27 out of 30 firms that made up the Dow Jones Industrial in 2007, made by 299 identified equity analysts. Table 1 displays the descriptive statistics of my final sample. This concerns values before all variables are de-meaned and (some) scaled for regression purposes.

Panel A has the means and distribution characteristics of several forecast characteristics. It reports on the numbers after having dropped all but the last forecast per analyst and forecast period, as previously described. It appears that when looking at the whole sample, company earnings are on average underestimated by \$0.031 per share. The minimum and maximum forecast error values are part of a larger group of forecasts mainly issued in the Early and 1st Revision Period. One can imagine that much can change over the course of a year, especially when taking in mind the turbulent financial and economic conditions in 2007-2008. The mean absolute forecast error of the sample, which I later need to be able to attach Dollar amounts to coefficients, is equal to \$0.102. The average forecast horizon is approximately equal to six months, while the maximum is 371 days.[17] Additionally, the median and 25th and 75th percentile numbers are striking. Although these are a result of the four peaks in forecast issuance around the 360-, 270-, 180- and 90-day mark, as reported earlier. Furthermore it seems that the analysts in my sample are not overly optimistic. On average the analysts in my sample will overestimate company earnings roughly four out of ten times, while underestimating earnings in the other instances.[18] Panel B gives a breakdown of the variable means means by the analysts' locations, while also adding three more

[17] Note that while I only use one-year ahead forecasts, there are forecasts that are issued more than 365 days before earnings release. This follows from the fact that the age of the forecast is measured in relation to the date of the earnings release by the firm. While one-year forecast in IBES are forecast made within one year of the end date of the fiscal period. The delay between period end and earnings release causes this discrepancy.
[18] This concerns all stocks covered by the analyst in 2007, not limited to those in my sample. As this variable functions as an indicator of the overall tendency towards optimism.

Table 1: Descriptive Statistics

This table shows the summary statistics of my main sample. Panel A reports the statistics for all fiscal year and quarterly earnings forecasts made in 2007. The forecast age or forecast horizon is equal to the number of days between forecast issuance and earnings release. The brokerage size is equal to the number of active analysts working at a firm in 2007. The distance between the analyst's location and New York (NY) and the covered firm's headquarters (HQ) is calculated as the crow flies in kilometres. Pct. Optimistic is equal to the percentage of company earnings for which the analyst's latest forecast was above actual earnings. Panel B gives the means by analyst location, while adding the percentages of forecasts by affiliated analysts, All-Stars and analysts at high-status brokerage firms. Finally the means for the two types of New York analysts are shown.

Panel A: Summary Statistics of Final Sample (*n* = 4985)

VARIABLES	Mean	sd	Min	25th	Median	75th	Max
Forecast Error ($)	-0.031	0.169	-1.620	-0.080	-0.030	0.010	1.330
Abs. Forecast Error ($)	0.102	0.139	0	0.020	0.060	0.120	1.620
Forecast Age	182.2	103.1	0	91	181	272	371
Distance to NY	646.3	1341	0	0	0	307.9	4186
Distance to HQ	1195	1302	0	34.09	917.7	1827	4336
Brokerage Size	71.90	41.73	1	31	92	105	138
Firm-specific Exp.	6.888	4.557	1	4	6	8	26
General Exp.	10.66	5.788	1	7	9	13	26
Pct. Optimistic	39.17	14.06	0	29.17	37.70	47.52	78.69

Panel B: Means by Analyst Location

	NYC = 1	NYC = 0	DIFF = 0	NY = 1	NY_LOC = 1	DIFF = 0
	(*n = 3542*)	(*n = 1443*)	*Sign.*	(*n = 2149*)	(*n = 1393*)	*Sign.*
Forecast Error ($)	-0.029	-0.036	0.19	-0.036	-0.019	0.01
Abs. Forecast Error ($)	0.104	0.096	0.04	0.099	0.112	0.01
Forecast Age	180.1	187.4	0.02	1.79.2	181.5	0.53
Distance (km.) to NY	2.324	2227	0.00	2.346	2.289	0.89
Distance (km.) to HQ	1113	1398	0.00	1821	20.30	0.00
Brokerage Size	82.33	46.29	0.00	81.72	83.27	0.25
Firm-specific Exp.	7.074	6.432	0.00	7.230	6.832	0.02
General Exp.	10.89	10.09	0.00	10.94	10.82	0.58
Pct. Optimistic.	38.94	39.75	0.07	39.07	38.74	0.50
Pct. Affiliated	4.376	0	0.00	1.349	9.045	0.00
Pct. All-Star Status	46.81	17.67	0.00	48.95	43.50	0.00
Pct. High-Status	32.98	12.61	0.00	33.04	32.88	0.92
No. of Analysts	199	100		152	90	
No. of Covered Firms	27	27		19	10	
No. of Brokerage Firms*	43	47		37	31	

* In total there are analysts from 69 brokerage firms in the final sample. Evidently, larger brokerage firms with multiple offices in different cities are included as one unique brokerage firm.

(dummy) variables. On the left-hand side of the table forecasts by analysts located in New York (NYC = 1) are compared to all others (NYC = 0). Out of the 299 analysts in the final sample 199 are located in New York, and the other 100 are located more than 100 kilometres from downtown Manhattan. From the statistics it follows that the mean absolute forecast error is actually higher for New York analysts compared to others, although the difference is very small. Also the New York analysts in my sample have both more firm-specific as well as general experience. Both differences are significant and amount to 0.6 and 0.8 additional years of experience, respectively. Further, as to be expected New York analysts work at much larger brokerage houses than others. The New York analysts in my sample on average work at brokerage houses who employ roughly 82 analysts, while for all other analysts this averages at roughly 46. Also the percentages of forecasts made by All-Star analysts and analysts working at high-status brokerage firms strongly differ between the two groups of analysts. Almost half of all New York forecasts are issued by All-American analysts, compared to less than one in five for all other forecasts. The difference in high-status firms is slightly smaller, bit still considerably large. 32.98 percent of New York forecasts are made by analysts working at high-status brokerage firms, while that percentage is only 12.61 for all other forecasts. It seems that in general New York analyst are slightly less optimistic than other analysts, although this difference is rather small and less significant compared to the other characteristics. Finally all affiliated forecasts are issued by analysts from New York. This outcome is not surprising in the sense that any Dow Jones firm in my sample is likely to partner up with a big brokerage firm. Their sheer size rules out smaller brokerage firms to act as lead underwriter in any IPO or SEO one of these firms wants to pursue. Moreover, it is plausible to assume most of this business runs through New York, where all large brokerage firms have their main offices. Subsequently it makes sense that analysts from that office are to analyse these firms.

The right-hand side of Panel B further breaks down all forecasts by New York analysts into two groups; forecasts on firms headquartered within 100 kilometres of the analysts (NY_LOC = 1) and forecasts on firms located further away (NY = 1).[19] Most notable is the percentage of affiliated forecasts. Forecasts by affiliated analysts are for a large part concentrated in the category of New York analysts covering local firms. This is mainly down to two New York based firms with multiple underwriter relationships with different brokerage firms.[20] Apart from that only the differences in firm-specific experience and the percentage All-American analysts are significant. On average New York analysts have more firm-specific experience when covering non-local stocks, though this difference is small (7.2 vs. 6.8 years of covering a particular firm). And the percentage of All-Star

[19] As New York analysts can be part of both subcategories when covering different firms, the total amount of unique analysts in these subcategories exceeds the total of 199 New York-based. Likewise, the total number of firms covered in both categories exceeds the number of firms in the sample (29 vs. 27). This follows from one New York analyst who happens to be more than 100 kilometres away from two (New York-headquartered) firms he covers. All other New York analysts covering these stocks are within that range.
[20] Merck & Co. And General Electric make up for 116 out of 155 affiliated forecasts.

analysts is also higher among New York analyst covering non-local stocks, compared to New York analysts covering local stocks (49 vs. 43.5 percent). While significant, the differences in mean (absolute) forecast error do not hold much value as they are comparing forecasts on two different sets of firms.

Table 1.1 reports the correlation matrix of the variables as they will be used for regression testing. This means that some variables are log transformed and all are de-meaned by their corresponding forecast-period mean. The significance levels are in parentheses. First and foremost it stands out that for all three New York dummies the correlation coefficients with the two (dependent) variables that indicate the relative forecast error (DAFE and PMAFE) are negative, although all six are far from significant. However brokerage size (DSIZE) and the three dummy variables (DHIGH, DSTAR and DLOC) for analyst and brokerage status and locality (for non-New York analysts) are all negatively correlated with both forecast error measures. All of which are statistically significant. Concentrating on the New York dummies again, there are the expected positive relations with brokerage size, and analyst and brokerage status. Besides that, all three are negatively correlated with the optimism variable.

Unsurprisingly, analyst status, brokerage size and brokerage status are also mutually positively correlated. Looking at the correlations of the optimism variable, more experienced analysts appear to be less optimistic than other analysts. As both measures of experience are negatively correlated with the optimism variable that captures the analyst's overall tendency towards optimism (DPCTOP). Finally note that, as could be expected, the two experience measures are highly correlated with one another.

4.2 Results of Regressing Relative Forecast Errors

I run a pooled, cross-sectional regressions using either DAFE or PMAFE as the dependent variable, as well as swapping out different proximity measures. The results from estimation of (several specifications of) Eq. (1) are reported in Table 2. One of the first things that catches the eye is the difference in magnitude of the regressions coefficients between the DAFE and PMAFE setups. This is because, although all coefficient values are reported as percentages, the DAFE coefficients report on a change in forecast error relative to stock price, while the PMAFE coefficients portray an effect relative to the mean absolute forecast error. As the mean forecast error represents a much smaller amount than stock price, coefficient values in these regressions are much bigger.

I employ the three different proximity measures in six separate regressions, using both dependent variables for each measure. Values of the adjusted R-squared are very small for all specifications of

Table 1.1: Correlation Table

This table depicts the correlation matrix of the de-meaned regression variables. Shown are the Pearson correlation statistics with the significance levels in parentheses.

	DAFE	PMAFE	DLOC	DAGE	DSIZE	DSTAR	DHIGH	DEXP	DFEXP	DPCTOP	DAFFIL	DNYC	DNY	DNY_ LOC
DAFE	1.00													
PMAFE	0.64	1.00												
	(0.00)													
DLOC	-0.03	-0.04	1.00											
	(0.02)	(0.01)												
DAGE	0.01	0.02	0.02	1.00										
	(0.72)	(0.08)	(0.19)											
DSIZE	-0.12	-0.11	-0.01	0.01	1.00									
	(0.00)	(0.00)	(0.30)	(0.34)										
DSTAR	-0.05	-0.04	0.00	-0.01	0.46	1.00								
	(0.00)	(0.01)	(0.82)	(0.56)	(0.00)									
DHIGH	-0.05	-0.05	0.06	0.02	0.44	0.45	1.00							
	(0.00)	(0.00)	(0.00)	(0.11)	(0.00)	(0.00)								
DEXP	-0.01	-0.03	0.05	-0.01	0.01	0.24	0.10	1.00						
	(0.47)	(0.04)	(0.00)	(0.69)	(0.32)	(0.00)	(0.00)							
DFEXP	0.01	-0.01	0.03	-0.01	-0.08	0.19	0.02	0.71	1.00					
	(0.33)	(0.44)	(0.05)	(0.34)	(0.00)	(0.00)	(0.24)	(0.00)						
DPCTOP	0.00	-0.01	0.08	0.01	-0.13	-0.08	-0.05	-0.13	-0.09	1.00				
	(0.87)	(0.63)	(0.00)	(0.40)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)					
DAFFIL	-0.01	0.00	-0.00	0.02	0.11	0.05	0.06	0.07	0.02	-0.08	1.00			
	(0.69)	(0.82)	(1.00)	(0.17)	(0.00)	(0.00)	(0.00)	(0.00)	(0.29)	(0.00)				
DNYC	-0.01	-0.01	-0.17	0.01	0.31	0.26	0.22	0.02	0.03	-0.16	0.07	1.00		
	(0.72)	(0.66)	(0.00)	(0.38)	(0.00)	(0.00)	(0.00)	(0.20)	(0.05)	(0.00)	(0.00)			
DNY	-0.00	-0.00	-0.20	0.02	0.24	0.22	0.19	0.04	0.04	-0.11	0.03	0.83	1.00	
	(0.93)	(0.82)	(0.00)	(0.21)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.07)	(0.00)		
DNY_LOC	-0.01	-0.01	-0.00	-0.00	0.20	0.13	0.11	-0.03	-0.01	-0.12	0.09	0.55	-0.01	1.00
	(0.61)	(0.66)	(1.00)	(0.77)	(0.00)	(0.00)	(0.00)	(0.06)	(0.43)	(0.00)	(0.00)	(0.00)	(0.47)	

** Each first line portrays the correlation coefficient, significance levels are shown in parentheses.*

Table 2: Regression of Relative Forecast Errors on Analyst Characteristics

This table depicts the results of running six variations of the baseline model (Eq. 1) using yearly and quarterly earnings forecasts of DJIA firms issued in 2007. The variations lie in the use of two different dependent variables and three proximity measures. The reported results are the estimated coefficients from running pooled cross-sectional regressions of de-meaned absolute forecast errors (DAFE) and proportional mean absolute forecast errors (PMAFE) on a number of analyst characteristics. All variables are mean adjusted by their respective forecast period mean (indicated by D for de-meaned). As such, DAFE equals analyst *i*'s absolute forecast error minus the mean absolute forecast error for firm *j* in forecast period *t*. Where absolute forecast error equals the absolute value of an analyst's latest forecast minus actual company earnings, as a percentage of the stock price. PMAFE equals DAFE divided by the mean absolute forecast error for firm *j* in forecast period *t*. LOC is a dummy variable equal to one if analyst *i* is located within 100 kilometres of the headquarters of the covered firm *j* and more than 100 kilometres away from New York City. AGE equals the natural logarithm of 1 plus the age of analyst *i*'s forecast, where the age is equal to the number of days between the issue date of the forecast and the date the actual earnings are released by the firm. SIZE equals the natural logarithm of 1 plus the number of analysts issuing earnings forecasts in 2007 at analyst *i*'s brokerage firm. STAR is a dummy variable equal to one if analyst *i* is an All-American analysts according to *Institutional Investor*. HIGH is a dummy variable equal to one if analyst *i* works at a high-status brokerage firm according to the rankings in *Institutional Investor*. EXP equals the natural logarithm of 1 plus the number of years in which analyst *i* has issued at least one earnings forecast, for any random stock. FEXP equals the natural logarithm of 1 plus the number of years in which analyst *i* has issued at least one earnings forecast for the covered firm *j*. PCTOP equals the percentage of firm periods for which analyst *i*'s last forecast is above actual earnings, looking at all companies followed by analyst *i* in 2007. AFFIL is a dummy variable equal to one if analyst *i*'s brokerage firm has an underwriting relationship with the covered firm *j*. NYC is a dummy variable equal to one if analyst *i* is located within 100 kilometres of the financial centre in New York City. LOGD equals the natural logarithm of 1 plus the physical distance (in kilometres) between analyst *i* and New York City. NY is a dummy variable equal to one if analyst *i* is located within 100 kilometres of New York City and more than 100 kilometres away from headquarters of the covered firm *j*. NY_LOC is a dummy variable equal to one if analyst *i* is located within 100 kilometres of New York City and within 100 kilometres of the headquarters of the covered firm *j*. The coefficients in the DAFE regressions display an effect on the dependent variable in Dollar amounts. Whereas the coefficients in the PMAFE regressions signal the expected change of the dependent in percentages. The robust t-statistics are in parentheses.

Indep. Var.	(1) DAFE	(2) PMAFE	(3) DAFE	(4) PMAFE	(5) DAFE	(6) PMAFE
DLOC	-0.0460****	-18.1295****	-0.0519****	-19.7092****	-0.0463****	-18.2595****
	(-3.72)	(-3.07)	(-4.22)	(-3.32)	(-3.60)	(-3.05)
DAGE	0.0059	8.3160	0.0061	8.3772	0.0059	8.3249
	(0.78)	(1.08)	(0.81)	(1.09)	(0.78)	(1.08)
DSIZE	-0.0246****	-9.5137****	-0.0238****	-9.3045****	-0.0246****	-9.5190****
	(-3.29)	(-4.04)	(-3.23)	(-4.00)	(-3.30)	(-4.05)
DSTAR	0.0025	3.8963**	0.0037	4.2337**	0.0025	3.8999**
	(0.47)	(1.66)	(0.66)	(1.77)	(0.47)	(1.66)
DHIGH	0.0002	-1.6127	0.0011	-1.3515	0.0002	-1.6036
	(0.04)	(-0.71)	(0.21)	(-0.58)	(0.04)	(-0.70)
DEXP	-0.0092*	-4.8650**	-0.0100**	-5.0628***	-0.0092*	-4.8488**
	(-1.61)	(-1.90)	(-1.74)	(-1.98)	(-1.61)	(-1.89)
DFEXP	0.0067	-0.6520	0.0073	-0.5138	0.0067	-0.6544
	(1.10)	(-0.26)	(1.20)	(-0.20)	(1.10)	(-0.26)
DPCTOP	-0.0002	-0.1587	-0.0002	-0.1743	-0.0002	-0.1581
	(-0.35)	(-1.21)	(-0.47)	(-1.33)	(-0.35)	(-1.21)
DAFFIL	0.0088	7.2114	0.0096	7.4509	0.0087	7.1680
	(1.02)	(1.09)	(1.11)	(1.13)	(1.00)	(1.09)
DNYC	0.0113	3.3866				
	(1.23)	(1.18)				
DLOGD			-0.0002	-0.0846		
			(-0.16)	(-0.25)		
DNY					0.0109	3.1938
					(0.97)	(0.97)
DNY_LOC					0.0121*	3.8075
					(1.58)	(0.91)
Observations	4,985	4,985	4,985	4,985	4,985	4,985
Adjusted R-squared	0.015	0.016	0.014	0.015	0.015	0.015

Robust t-statistics in parentheses

**** p<0.01, *** p<0.05, ** p<0.10, * p<0.15

the model. Beside this being common in such research, this is also a result from controlling for forecast-period effects, as Clement remarks about the use of de-meaned variables: "The model's R^2 represents the variation in the dependent variable explained after removing firm-year effects (all variables in the model are adjusted by subtracting firm-year means). This is similar to regressing absolute forecast errors on residuals from a regression of absolute forecast errors on firm-year dummies. Since a large source of variation (i.e., firm-year effect) has been removed from the dependent variable, the R^2 is significantly lower than it would be otherwise." (Clement 1999, p.299).

The results do not indicate that New York analysts outperform others, if anything they imply the opposite. The coefficients of all the New York dummy variables are positive, while the coefficients of the log distance variable are negative. The former implying larger forecast errors for New York analysts, and the latter increasing forecast errors as the distance between analyst and New York decreases. However, only the dummy variable for New York analysts covering local stocks in the DAFE model is significant at the 15% level. Apart from being statistically insignificant, multiplying the NYC coefficient in the PMAFE regression (3.39%) by the mean absolute forecast error from Table 1 ($0.102) translates the accuracy effect in a tiny 0.35 cents per share accuracy disadvantage.

From the top of the table it is clear that the locality variable is negative and significant in all specifications of the model. From the PMAFE estimations it follows that non-New York analysts covering local stocks are on average more than 18% more accurate. These results translate in a more economically significant per share accuracy advantage of 1.85 cents. Comparing this to the findings about local analysts in Malloy (2005), the magnitude of the local accuracy effect in my sample is much larger in terms of percentages (18.13% vs. 2.77%), while the effect is slightly smaller in dollar amounts ($0.019 vs. $0.025). The reason of the small dollar amount lies in the low mean absolute forecast error of $0.10 in my sample, compared to $0.89 in Malloy (2005). So it appears that the locality effect found in my sample is relatively high, however small in absolute terms. In the context of this high locality effect it is interesting to note that such an effect does not appear to be present among New York analysts whatsoever. The NY and NY_LOC coefficients do not differ from each other in a manner that would suggest any locality effect. Moreover, the magnitude of the negative accuracy effect actually appears to be larger for local New York analysts than non-local New York analysts.

Besides DLOC, also the coefficients of the brokerage size variable (DSIZE) are significant and negative in all regressions. As this variable is a logarithmic function its coefficient portrays the effect of a percentage change in brokerage size. Using the coefficient of regression (2) in Table 2, this suggests that an analyst at a brokerage firm with 105 analysts (the 75^{th} percentile value) is

11.61%[21] more accurate than an analyst at a brokerage firm that employs 31 analysts (the 25th percentile value). Of the two experience proxies, only the general experience coefficients are significant (at the 10% level in the PMAFE specifications). Though these two variables are highly correlated, I do not experience problems with collinearity in these regressions.[22] In a similar fashion to brokerage size, this indicates that an analyst who has thirteen years of forecasting experience (75th percentile) will be 3.01%[23] more accurate than an analyst who has only seven years of experience (25th percentile). The coefficient for the forecast age variable (DAGE) is insignificant in each of the six regressions, while they are all positive as expected. This suggests that the chosen forecast periods effectively control for new information arriving with the passage of time, as in similar research without forecast periods the forecast horizon is the most important determinant of forecast accuracy.

Looking at the two variables that control for reputation, the results are mixed. All DSTAR coefficients are positive, while it was expected that All-Star analysts would be more accurate. Although the results in the DAFE regression are not strong, the results in the PMAFE regressions are significant at the 10% level and indicate that All-American analyst are roughly 4% less accurate than others. The magnitude of this effect is slightly larger than the New York effect discussed above. The results for the brokerage firm dummy (DHIGH) are mixed in sign as well as statistically insignificant in all specifications. Though the results in the PMAFE regressions, which suggest a positive effect of brokerage reputation on forecast accuracy, are stronger than the opposite effect measured in the DAFE setup.

Finally, the coefficients for the general optimism (DPCTOP) and affiliation (DAFFIL) variables are also in the other direction than was expected. Though the results are insignificant, the results would suggest that unaffiliated analysts, who are generally more optimistic, have higher forecast accuracy. However, there are two reasons the expected negative relation between forecast accuracy and optimism may not be found. For one, the expectation is largely based upon analysts being overly optimistic, resulting in an optimism bias. From the summary statistics in Panel A of Table 1 it already became clear that this does not appear to be the case in my sample. Secondly, the general optimism variable includes all firms covered by the analysts, while the sample I use for research only covers a selection of those firms. In the next section I will further investigate the presence of analyst forecast bias in my sample and use an alternative measure of optimism to replace the general optimism variable.

---

[21] Ln(105/31) * -9.5137% = -11.61%

[22] Note from Panel C of Table I that DEXP and DFEXP are highly correlated (0.71). This reduces the model's ability to estimate the unique effects of these two experience variables. When I estimate the model without DFEXP, the accuracy effect of DEXP becomes slightly stronger (β = -5.35%) and significant at the 1% level. With DFEXP as only experience control, its accuracy effect is slightly lower (β = -3.78%) than the general experience variable, but also significant at the 5% level. This further proves that the model is not able to correctly separate the effects of the two experience indicators.

[23] Ln(13/7) * -4.8650% = 3.01%

For the remaining regression tests in this paper I focus on PMAFE as dependent variable, for a couple of reasons. First of all I prefer it over DAFE, as PMAFE actually relates the forecast error to the mean forecast error of the forecasts I want to compare it with.[24] As a result the coefficient values are also much more intuitive because they directly show a difference of forecast accuracy in percentages. Besides that, the regressions results are very similar to those when using DAFE as dependent variable. Similarly, I will focus on the three New York dummies hereafter and disregard the logarithmic distance variable in the remainder.

4.3 A Deeper Look into Forecast Bias and Informativeness

One issue with using forecast errors to try to indicate that certain analysts possess superior information is that there is no distinction between bias and informativeness of the forecasts. For instance, in the case of New York analysts, their lower accuracy may be a result of them being overly optimistic about the stocks they follow. In other words the observed underperformance, albeit insignificant, may be a results of a bias in their forecasts rather than the forecasts being less informative. In this section I run a simple regression that separates bias from informativeness of the forecasts and after that I also apply a specific control for it in the relative forecast error regressions.[25]

To be able to display the forecasts informativeness and identify possible biases I use a methodology used by DeBondt and Thaler (1990). This entails regressing the actual earnings per share on the earnings per share as forecasted by the analysts.[26] For the analysts' forecasts to be rational, it would require the intercept to be equal to zero in a regression of this type. A negative intercept - as found by DeBondt and Thaler (1990) and also Malloy (2005) - signals that on average forecasts are too optimistic, and is known as analyst optimism bias. The slope in such a regression tells us about the informativeness of the forecasts and whether earnings are under- or overestimated. Generally speaking, the closer the slope is to one, the more informative the forecast. A beta below (above) one is associated with forecasts being too high (low). By running such regressions separately for two groups of analysts, one is able to identify differences in bias and informativeness. By regressing both groups together and using dummies, one is able to test the

[24] Furthermore, the drawback of the DAFE specification is that firm periods with high EPS-to-stock price-ratios will produce higher values error values in the DAFE setup than firm periods with lower ratios. While the smaller DAFE values in low EPS firm periods may actually be relatively worse. This is of no concern when deflating by the forecast period mean.

[25] Louis et al (2013) offer some interesting insights into the relationship between the accuracy and 'informativeness' of a forecast. In their paper they suggest that forecasts that are less accurate can actually be more informative (of 'true' earnings), when analysts deviate from management guidance (preannounced earnings) to correct for perceived earnings management. However, in the context of this paper I stick to the informativeness of the forecast about reported (albeit managed) earnings number.

[26] These are just the plain dollar values of both the actual and forecasted earnings per share.

statistical significance of these differences.[27]

I run these regressions for forecasts by New York analysts and others separately and look at differences between the two groups. Apart from comparing all New York forecasts to the forecasts by other analyst, I additionally do the same for (non-)New York firms. Much like how I made the distinction between the two types of New York analyst before. The results are displayed in Table 3. One of the first things that stands out is that all values of beta are above one. The accompanying test-statistic follows from a F-test whether beta is equal to one. So all slopes are significantly higher than one, indicating that on average both groups of analysts underestimate company earnings, in either the entire sample or the subsets. The differences in magnitude of the slopes between New York analysts and others are negligible. The same goes for the values of R-squared. These results indicate that any systematic differences in forecast accuracy between New York analysts and others do not appear to be caused by a difference in informativeness of the forecasts.

The coefficients of the intercepts do produce interesting results. In all three samples, the intercept of the regression of actuals on New York forecasts is not significantly different from zero. Though all three are negative. So it seems that on average the New York analysts in my sample are not hindered by any form of bias. However, the results for non-New York analysts do show analyst forecast bias. When looking at all 27 firms, I find the opposite of the classic optimism bias with the intercept equal to 0.0106 and significant at the 10% level. Interestingly, this bias is even stronger when focusing on non-New York firms (α = 0.0178) and reversed for the subsample with New York firms (α = -0.0287). Both at least significant at the 5% level. The differences in bias between the two types of analysts are displayed at the bottom of the table. Since the New York analysts did not display bias, the results are pretty straightforward. They indicate that New York analysts are generally less pessimistic (t = -1.74) than others, when one looks at it from their perspective. This effect is even stronger when it concerns non-New York firms (t = -2.19). However for New York firms they are actually less optimistic than non-New York analysts (t = 1.50).

These results have implications on the forecast accuracy regressions done so far. It does not necessarily suggest that there should be systematic differences between New York analysts and others. However, it does indicate that bias may be an important influential factor. To investigate whether a systematic optimism bias affected the results of the baseline model from Table 2, I add an explicit control for relative optimism. For this purpose I use the variable ROPT (Relative Optimism). This variable replaces PCTOP, which captures an analyst's overall optimism about all the stocks he or she covers. The construction of this variable is similar to that of the dependent PMAFE, except for not using absolute values of the forecast errors in the calculation. ROPT, then, is equal to an analyst's forecast error (FE_{ijt}) minus the forecast period mean forecast error (MFE_{jt}),

[27] A dummy that identifies one of the two groups (NYC in this case) is added to the equation to capture the difference in the bias. Secondly, the same dummy multiplied by the forecasted earnings is also added to the equation to capture the difference in informativeness between the two groups.

Table 3: Regressing Actual Earnings on Forecasted Earnings

This table shows the intercepts, coefficients and R^2 values of running pooled cross-sectional regressions of actual earnings on forecasted earnings. These regressions are performed separately for New York and non-New York analysts and differences between the two are provided (DIFF). This same process is repeated using either forecasts on New York firms or non-New York firms exclusively. Robust t-statistics are in parentheses.

$(A_{i,j,t} = \alpha_i + \beta_i F_{i,j,t} + \varepsilon_{i,j,t})$

	All Firms			New York Firms			Non-New York Firms		
Analyst Location	Estimate of α (t-statistic for test of $\alpha = 0$)	Estimate of β (F-statistic for test of $\beta = 1$)	R^2	Estimate of α (t-statistic for test of $\alpha = 0$)	Estimate of β (F-statistic for test of $\beta = 1$)	R^2	Estimate of α (t-statistic for test of $\alpha = 0$)	Estimate of β (F-statistic for test of $\beta = 1$)	R^2
NYC = 1	-0.0020	1.0206****	0.9844	-0.0053	1.0159****	0.9852	-0.0007	1.0246****	0.9833
	(-0.45)	(36.91)		(-0.69)	(9.17)		(-0.14)	(31.07)	
NYC = 0	0.0106**	1.0190****	0.9879	-0.0287***	1.0203***	0.9878	0.0178****	1.0228****	0.9878
	(1.85)	(14.00)		(-2.12)	(6.53)		(2.70)	(11.19)	
DIFF (1 - 0)	-0.0126**	0.0016	-0.0035	0.0233*	-0.0044	-0.0026	-0.0185***	0.0017	-0.0045
	(-1.74)	(0.27)	(0.00)	(1.50)	(-0.47)	(0.00)	(-2.19)	(0.22)	(0.00)

Robust t-statistics in parentheses

**** p<0.01, *** p<0.05, ** p<0.10, * p<0.15

divided by the absolute value of MFE_{jt}.[28] As a result ROPT captures the sign (and magnitude) of the forecast error and is scaled the same way as the dependent variable. Positive values of ROPT will then correspond with forecasts of which the value is above the mean forecast of its forecast period, indicating higher relative optimism about the stock. Accordingly the opposite holds for negative values, and thus the variable can serve as a proxy for (optimism) bias.

In the regressions with ROPT, there are 21 observations less than in the previous regression. This results from the fact that the relative forecast variable depends on the mean forecast error to be non-zero. However in two cases forecast errors by several analysts cancel each other out exactly, and the mean error is equal to zero. Further, I am forced to add two dummy variables that catch the 1% extreme values of ROPT, both high and low. This follows from forecast periods where the mean forecast error (the denominator in the calculation of ROPT) approaches zero, greatly inflating ROPT values. The results of the regressions are shown in Table 3.1. The first two regressions are reruns of the baseline model, where PCTOP has been replaced by the new control for relative optimism. The results show that, when explicitly controlling for the analysts' relative optimism, most coefficients are virtually unchanged compared to the first run with the analysts' overall tendency to optimism. Only the coefficient of affiliation variable is notably higher, and now significant at the 15% level. The three New York dummies change ever so slightly relative to the regression results presented in Table 2.

Apart from this, both coefficients for ROPT are negatively and strongly significant. Indicating that on average analysts who are relatively optimistic about a stock, are more accurate in their forecasts of said stock. This contradicts the findings by Malloy (2005), who in similar regressions found that relatively optimistic analysts performed worse. Finally, the R-squared value is a lot higher than in the regressions with PCTOP instead of ROPT (0.074 vs 0.016 before). However, this is largely due to the two dummy variables that are used here. Omitting these dummies has little to no effect on the outcome of these regressions. With the exception that it renders the ROPT coefficients insignificant and three times as small in magnitude, while R-squared drops to values similar to the corresponding regressions from Table 2 (0.019).[29]

I also run the same regressions after replacing PMAFE with ROPT as the dependent variable (regressions 3 and 4 from Table 3.1). This produces some interesting results. Although the addition

[28] As before, forecast errors as a percentage of stock price are used here.

[29] The fact that the two dummies explain a significant amount of variation in PMAFE clarifies that the dummies, besides controlling for 'extreme' values of ROPT (when MFE is very close to zero), also single out a number of forecasts with large PMAFE values. So there are also 'extreme' values of ROPT caused by big forecast errors instead of extremely small MFE (or a combination). The only way to separate the two would be to set some arbitrary margin for MFE being too close to zero. However, as the results of the regressions are not affected by the addition of the dummies (apart from the ROPT coefficient and R-squared obviously), it is clear that my results do not depend on some of these high PMAFE observations. Hence adding the two dummies serves its purpose of correctly showing the effect of relative optimism on forecast accuracy. In the robustness section I go into more detail about large values of PMAFE.

Tabel 3.1: Controlling for Relative Optimism

This table reports the estimated coefficients from pooled cross-sectional regressions of proportional mean absolute forecast errors (PMAFE) and relative optimism (ROPT) on a number of analyst characteristics. The dependent variable PMAFE and the independent variables LOC, AGE, SIZE, STAR, HIGH, EXP, FEXP, AFFIL, NYC, NY and NY_LOC are defined as in Table 2. It concerns baseline model used before, where the general optimism variable (PCTOP) has been replaced by a relative measure (ROPT). The new addition ROPT equals analyst *i*'s forecast error (FE_{ijt}) minus the forecast period mean forecast error (MFE_{jt}), divided by the absolute value of MFE_{jt}. In all four regressions two dummies are added two control for extreme values of ROPT, coefficients of these are not reported. Again, all variables are mean adjusted by their respective forecast period means (indicated by D for de-meaned). All coefficients report the expected change of the dependent variables in percentages. The robust t-statistics are in parentheses.

	(1)	(2)	(3)	(4)
Indep. Var.	PMAFE	PMAFE	ROPT	ROPT
DLOC	-17.5206****	-17.5487****	30.2355****	44.3714****
	(-2.96)	(-2.92)	(2.68)	(4.06)
DAGE	8.0331	8.0354	-23.9502***	-25.0085***
	(1.15)	(1.15)	(-2.25)	(-2.34)
DSIZE	-9.0527****	-9.0540****	-8.3771**	-7.7036**
	(-4.15)	(-4.16)	(-1.91)	(-1.75)
DSTAR	2.7482	2.7491	-7.5690	-8.0185
	(1.22)	(1.22)	(-1.13)	(-1.19)
DHIGH	-1.5166	-1.5146	12.2629**	11.1885*
	(-0.68)	(-0.68)	(1.78)	(1.64)
DEXP	-4.1703**	-4.1672**	8.4216	6.7817
	(-1.68)	(-1.68)	(1.16)	(0.94)
DFEXP	-0.0563	-0.0569	0.2227	0.5202
	(-0.02)	(-0.02)	(0.03)	(0.08)
ROPT	-3.9523****	-3.9515****		
	(-3.49)	(-3.47)		
DAFFIL	9.5429*	9.5323*	58.6475****	63.7650****
	(1.48)	(1.48)	(4.44)	(4.73)
DNYC	3.6811		-3.4101	
	(1.34)		(-0.42)	
DNY		3.6386		18.1876****
		(1.16)		(2.59)
DNY_LOC		3.7739		-50.5217****
		(0.91)		(-2.74)
Observations	4,964	4,964	4,964	4,964
Adjusted R-squared	0.074	0.074	0.627	0.628

Robust t-statistics in parentheses
**** p<0.01, *** p<0.05, ** p<0.10, * p<0.15

of ROPT in the regression with PMAFE as independent variable did not change the outcome of the New York effect, two of the New York dummies are strongly significant when ROPT is the dependent. Further the direction of the coefficients of DNY and DNY_LOC supports the findings from regressing actual earnings on forecasted earnings earlier. New York analysts covering local stocks are almost 51% more pessimistic than others, while New Yorkers are roughly 18% more optimistic than others when it concerns stocks outside New York. Certainly the finding that New York locals are that much more pessimistic than other analyst is interesting. Especially so, as local analysts outside of New York are relatively optimistic about the stocks they cover (DLOC = 44.4%, *t* = 4.06). Unsurprisingly, affiliated analysts are significantly more optimistic than unaffiliated analysts. In addition local non-New York analysts are also more optimistic than others.

The results for the age variable, however, are surprising. They are surprising when compared to the results Malloy (2005) finds in a similar regression. He finds a beta of roughly 57 for the log age variable, a result that is indicative of forecasts being too high/optimistic at the beginning of the fiscal year.[30] In a similar regression, however, I find two surprising betas for DAGE of -24 and -25. These indicate that on average forecasts become more optimistic as time passes. There is an important distinction to be made though. Because I make use of forecast periods the age variable supplies information of the forecast's age relative to forecasts in the same forecast period, rather than the trend over the course of a year. Thus these results tell us that on average forecasts are significantly more optimistic at the end of forecast periods. One possibility is that this effect results from quarterly earnings announcements that on average, and thus repeatedly, brought unexpected positive news - or in the broader time frame better yet put as news that was unexpectedly less negative. Another possibility is that the increasing optimism in forecast periods fits in a yearly trend of increasing optimism. I do not find strong evidence for the latter, though.[31] However, this is well beyond the scope of my research and needs closer looking into to draw a meaningful conclusion.

For these two regressions with ROPT as the dependent variable it also holds that the amount of variation that the model is able to explain greatly relies on the two dummies that capture extreme values of ROPT. Without the inclusion of the dummies the value of R-squared actually drops from 0.627 to 0.001. Besides that, coefficients of most of the variables change considerably and some become statistically insignificant, when the dummies are omitted. Confirming that those 2% of extreme (inflated) values of ROPT are responsible for a disproportionate amount of variation in ROPT.

4.4 Controlling for Coverage Decisions and Analyst Effort

In this section I expand the baseline model to control for coverage decisions by the analysts. As analysts decide for themselves which stocks they follow and how intensively they do so, it is possible that this is influencing previous results. Where I did not find any indications of superior New York performance and inherent information asymmetry, controlling for coverage decisions may change this. I introduce three new variables to the baseline model to control for the analysts' coverage decisions, so that the expanded regression equation takes the form of:

---

[30] A result that coincides with the finding of Abarbanell and Bernard (1992) that forecasts at the beginning of the fiscal year are overly extreme due to analysts overreacting to "some information source".

[31] When I run these regressions on a sample without forecast periods (so only one forecast per analyst per firm period) the coefficient of DAGE is also negative, though insignificant and much smaller in magnitude (DAGE = -6.54).

$$
\begin{aligned}
\text{PMAFE}_{ijt} = {} & \beta_1\,\text{DLOC}_{ijt} + \beta_2\,\text{DAGE}_{ijt} + \beta_3\,\text{DSIZE}_{ijt} \\
& + \beta_4\,\text{DSTAR}_{ijt} + \beta_5\,\text{DHIGH}_{ijt} + \beta_6\,\text{DEXP}_{ijt} \\
& + \beta_7\,\text{DFEXP}_{ijt} + \beta_8\,\text{DPCTOP}_{ijt} + \beta_9\,\text{DAFFIL}_{ijt} \\
& + \beta_{10}\,\text{DFREQ}_{ijt} + \beta_{11}\,\text{DCOV}_{ijt} + \beta_{12}\,\text{DSPEC}_{ijt} \\
& + \beta_{13}\,\text{DNYC}_{ijt} + \varepsilon_{ijt}, \qquad (2)
\end{aligned}
$$

where, again, the dependent and all variables preceded by a capital 'D' are forecast period de-meaned.[32] In comparison to the baseline model from Eq. (1), to this expanded model depicted in EQ. (2) the variables DFREQ, DCOV and DSPEC are added. Firstly, DFREQ has information on how many forecasts per firm an analysts issues on average, serving as a proxy for effort. Secondly, DCOV controls for the number of firms an analysts covered in 2007, serving as a control for focus. And lastly, DSPEC is a dummy that controls for analysts that are specialised in covering stocks within one major industry group.

Table 4 displays the results of running the expanded model, depicted next to the results of the baseline model from Table 2. The results show that specialised analysts and those who generally display more effort are significantly more accurate than other analysts. especially the result of DSPEC is strong, indicating that specialised analysts are 11.0% more accurate than others (t = 3.08). However, the results for the New York effect are only marginally affected by the inclusion of the new variables. All three New York dummies remain positive and statistically insignificant.[33]

4.5 Sorting by Characteristics

Up till now I have not found New York analysts to be more accurate than others. However, it remains possible that there are certain areas for which proximity to New York does matter. In this section I provide a breakdown of my sample into several subsample to try to isolate areas where the New York effect may play an important role.

First of all I categorize forecasts according to their signal attributes. Following Gleason and Lee (2003), when a forecast is above the analyst's prior forecast for that firm period and also above the prior consensus[34], I classify that forecast as signalling unambiguously good news (SIGNAL = 1). The opposite holds for forecasts that are both below the analyst's prior forecast and the prior consensus (SIGNAL = -1). All other forecast fall into the third category (SIGNAL = 0). Further, I break the sample down by the number of stocks each analyst covers in 2007. I distinguish between

[32] See appendix for an overview of variable definitions.
[33] For further testing I will use the expanded model.
[34] The mean forecast error (MFE) of the prior forecast period serves as a proxy for the prior consensus forecast.

Table 4: Expanding the Model with Specialisation, Focus and Effort

In this table I report the results of regressing the expanded model (Eq. 2) on the right-hand side. The results of the baseline model on the left-hand side are the same that I reported in Table 2. With the expanded model three variables are added: FREQ equals the natural logarithm of 1 plus the mean number of forecasts analyst *i* issues per covered company in 2007; COV equals the natural logarithm of 1 plus the number of unique firms analyst *i* covers in 2007; and SPEC is a dummy variable equal to one if analyst *i* only covers firms in a single major industry group. Only the coefficients of the three new variables and the location dummies are reported here. PMAFE, NYC, NY, NY_LOC and all unreported variables are defined as in Table 2. All variables are mean adjusted by their respective forecast period means (indicated by D for de-meaned). All coefficients report the expected change of the dependent variable in percentages. The robust t-statistics are in parentheses.

	Baseline Model		Expanded Model	
Indep. Var.	PMAFE	PMAFE	PMAFE	PMAFE
DFREQ			-4.7504***	-4.8638***
			(-1.98)	(-2.03)
DCOV			-4.1644	-4.1450
			(-1.29)	(-1.29)
DSPEC			-10.9866****	-11.0519****
			(-3.08)	(-3.10)
DNYC	3.3866		3.4395	
	(1.18)		(1.23)	
DNY		3.1938		2.7201
		(0.97)		(0.85)
DNY_LOC		3.8075		5.0212
		(0.91)		(1.21)
Observations	4,985	4,985	4,985	4,985
Adjusted R-squared	0.016	0.015	0.018	0.018

Robust t-statistics in parentheses
**** $p<0.01$, *** $p<0.05$, ** $p<0.10$, * $p<0.15$

Low-coverage, Mid-coverage and High-coverage analysts. After computing the number of stocks followed by each analyst in my sample, analysts who follow less firms than the 33rd percentile value are Low-coverage. Analysts who cover more than the 66th percentile value of stocks followed, are classified as High-coverage analysts. The remaining analysts fall in the Mid-coverage category. Finally I also separate the sample into forecasts by All-Star analysts (STAR) and non-All-Star analysts (NON), as well as forecasts by specialised analysts (SPEC) and analysts covering stocks in more than one major industry group (NON) otherwise.

Table 5 shows the results of regressing the expanded model from Eq. (2) on the above described subsamples, though only the New York dummy results are reported. With each sample I run two regressions with either the one combined New York dummy or the two separate dummies. Panel A displays the results of grouping together observations by the signal of the forecast and analyst status. The coefficients in the signal breakdown are not that interesting. None are significant and all but one are positive, though the magnitudes do differ somewhat between the subsets of analysts. Only the result for analysts making unambiguous downward revisions on a New York stocks deviates from the pattern of positive coefficients, while its t-statistic is equally small though. The results of the breakdown by analyst status are interesting, however. They

Table 5: New York Dummy Coefficients after Breaking Down the Sample by Characteristics

Pooled cross-sectional regressions of proportional mean absolute forecast error (PMAFE) on analyst characteristics are run, following the expanded model (Eq. 2). The sample is broken down by a number of characteristics and only the coefficients and t-statistics of the location dummies are reported. As such, for each subsample two regressions are run; one with the single New York dummy (NYC) and also one with NY and NY_LOC as location dummies. PMAFE, NYC, NY, NY_LOC are defined as in Table 2. In panel A SIGNAL = 1 if a forecast is above analyst *i*'s prior forecast for that firm period and also above the prior consensus. SIGNAL = -1 if analyst *i* revises the forecast both below his prior forecast and the prior consensus. All other forecast fall into the third category (SIGNAL = 0). In the STAR column only All-Star analysts are included, while the NON column excludes those. In Panel B Low-coverage analysts are analysts who cover less firms than the 33rd percentile of firms covered by the analysts in the sample. Analysts who are above 66th percentile are classified as High-coverage analysts. The remaining analysts fall in the Mid-coverage category. In the SPEC column only specialised analysts covering stocks in only one major industry group are included, while the NON column excludes those. All variables are mean adjusted by their respective forecast period means (indicated by D for de-meaned). All coefficients report the expected change of the dependent variable in percentages. The robust t-statistics are in parentheses.

Panel A: Breakdown by Signal and Analyst Status

		SIGNAL			Analyst Status	
Dep. Var.	Indep. Var.	1	-1	0	STAR	NON
PMAFE	DNYC	3.2929	1.3128	4.3354	-3.3297	7.5050***
		(0.46)	(0.22)	(1.43)	(-0.79)	(2.32)
PMAFE	DNY	3.7071	2.6380	3.0061	-8.7838*	8.3845***
		(0.47)	(0.40)	(0.83)	(-1.52)	(2.23)
PMAFE	DNY_LOC	2.2061	-3.6528	6.9654*	5.7484	5.3091
		(0.23)	(-0.27)	(1.50)	(1.14)	(1.00)
Obs.		1,343	616	2,731	1,635	2,871

Panel B: Breakdown by Coverage and Brokerage Status

		Coverage			Specialisation	
Dep. Var.	Indep. Var.	Low	Mid	High	SPEC	NON
PMAFE	DNYC	21.8655****	6.7192**	-14.3147****	-11.4380*	2.7510
		(4.05)	(1.67)	(-3.02)	(-1.48)	(0.99)
PMAFE	DNY	19.0585****	5.1649	-15.5077****	-10.3628	3.9183
		(2.89)	(0.97)	(-3.01)	(-1.25)	(1.20)
PMAFE	DNY_LOC	25.8429****	8.8618*	-10.1648	-15.9400	0.2387
		(3.62)	(1.51)	(-1.20)	(-1.25)	(0.06)
Obs.		1,687	1,074	1,228	772	3,992

Robust t-statistics in parentheses
**** p<0.01, *** p<0.05, ** p<0.10, * p<0.15

indicate that among non-All-Star analysts New York analysts are significantly less accurate by 7.5% (t = 2.32). This translates into a still very small \$0.007 per share accuracy disadvantage on average for New York analysts.[35] Among All-Star analysts the effect appears to be opposite, although it is smaller in magnitude (-3.3%) and statistically insignificant. These effects appear to be concentrated in the dummy that singles out New York analysts covering non-New York stocks (DNY). When it concerns New York stocks (DNY_LOC) both coefficients are positive and insignificant.

Panel B of Table 5 provides the breakdown by the size of the analyst's portfolio and analyst specialisation. The results after categorising the forecasts using the coverage terciles are striking.

[35] Dollar amount = mean absolute forecast error for non-All-Stars (\$0.098) multiplied by 7.5%.

Among low-coverage analysts New Yorkers are 21.9% less accurate than others (t = 4.05), while among high-coverage analysts they are actually 14.3% more accurate (t = -3.02). This amounts to a more economically significant $0.022 per share accuracy disadvantage and a $0.014 per share accuracy advantage, respectively. Where the former result is significant (and of roughly the same magnitude) for both New York stocks and other stocks, the latter appears to be concentrated in non-New York stocks (DNY = -15.5%, t = -3.01) and is insignificant for New York stocks (t = -1.20). The results suggest that only among relatively active analysts New Yorkers have an inherent information advantage over other (also active) analysts. Mid- and low-coverage New York analysts, however, are not able to benefit from their geographic location. As a matter of fact New Yorkers in the lowest coverage tercile even appear to suffer from an inherent informational disadvantage compared to their distant counterparts.

The right-hand side of Panel B displays the breakdown by analyst specialisation. It appears that among specialised analysts New York analysts are more accurate by 11.4%, although this results is only marginally significant at the 15% level. For non-specialised analysts the coefficients for the New York effect are, again, all positive and insignificant.

4.5.1 Consequences of the Sample Breakdown Results

It seems there is not one unique effect at play that can explain both New York informational advantages and disadvantages in certain subsamples. Moreover, they have cancelled out one another in my regressions using the entire sample, where I did not report any significant results. Apart from producing significant results, the subsamples' composition can offer some insight in what drives certain inherent information asymmetries. The superior performance among high coverage analysts is indicative of better performance as a result of private information rather than reduced cost of information gathering. Where superior (private) information is only accessible to active analysts who benefit from absorbing the local culture and maintaining contact with local financial and business leaders.[36] A similar narrative fits the superior performance of New York analysts among specialised analysts. Where specialised New York analysts, by focusing on a single industry, are more immersed in the relevant New York scene than non-specialised New Yorkers.

However this does not explain the poor performance of New York analysts among lower coverage analysts and non-All-Stars. These results certainly do not coincide with reduced cost of information gathering in New York, in fact they rather imply the opposite. I conjecture that this increased cost of information gathering in New York results from an heightened difficulty of

[36] Malloy (2005) also finds that the locality effect is stronger for more active analysts, especially when the share of local stocks in their portfolio is greater. For future research it may be interesting to test whether a larger share of New York stocks in an New York analyst's portfolio also improves their performance on non-New York stocks. I am not able to test this as my sample only consists of a small sample of firms.

assessing the quality of information faced by local analysts. By this logic the abundance of information sources in New York puts New York analysts at a disadvantage. This story also fits the results of the breakdown by analyst status. I propose it is not a stretch to assume that All-Star analyst are better equipped to judge sources of information. The fact that among All-Star analysts the New York effect is insignificant (and the coefficients are even negative for two out of three dummies), whereas among non-All-Star analysts New Yorkers perform significantly worse than others analysts, supports this reasoning. Additionally, one could assume that All-Star analyst are probably also more immersed in the local culture and relevant social circles.

In other words the results of the sample breakdown suggests that geographic proximity to New York is a double-edged sword for analysts. It seems that while some New York analysts are able to benefit from superior information, others are at a disadvantage as a result of high cost of information gathering relative to analysts outside of New York.

4.6 Sorting by Forecast Periods

Finally, I utilize the design of the different forecast periods to investigate if there is a timing factor to the New York effect. The results of running the expanded model for the five forecast periods separately are presented in Table 6. Much like the results for the full sample, the coefficients for the New York dummies are predominantly positive and statistically insignificant for the forecast period subsamples. The exception being that in the early forecast period New York analysts covering local stocks perform significantly worse than other analysts covering the same stocks (DNY_LOC = 18.9%, t = 2.47).

Table 6: New York Effect sorted by Forecast Periods

This table reports the results of cross-sectional regressions of proportional mean absolute forecast error (PMAFE) on analyst characteristics, following the expanded model (Eq. 2). In this case the regressions are performed separately on forecasts from the five established forecast periods. Apart from that, these are the same regressions as the expanded model in Table 4, however only reporting the coefficients of the location dummies NYC, NY and NY_LOC. All variables are defined as before and mean adjusted by their respective forecast period means (indicated by D for de-meaned). All coefficients report the expected change of the dependent variable in percentages. The robust t-statistics are in parentheses.

Dep. Var.	Indep. Var.	Early Forecast Period	1st Revision Period	2nd Revision Period	3rd Revision Period	Late Forecast Period
PMAFE	DNYC	7.1542	-0.9068	8.5779*	3.4211	2.3097
		(1.03)	(-0.17)	(1.48)	(0.64)	(0.18)
PMAFE	DNY	1.0928	-1.9696	8.2554	2.7688	10.9079
		(0.12)	(-0.30)	(1.27)	(0.45)	(0.85)
PMAFE	DNY_LOC	18.9465***	1.2660	9.3586	4.7765	-27.6426
		(2.47)	(0.17)	(1.17)	(0.62)	(-0.83)
Observations		582	1,137	1,281	1,520	465

Robust t-statistics in parentheses
**** p<0.01, *** p<0.05, ** p<0.10, * p<0.15

5. Robustness Tests

In this section I examine whether the results presented thus far are sensitive to a number of assumptions and choices I made. This concerns choices that influence the sample size and the distribution of forecasts among the New York categories. In this section I will outline a number of alterations I make to my research so far and describe why they matter as a robustness check. Thereafter I will briefly discuss the results. Detailed information of the regression outputs is provided in the appendix.

5.1 Alternative Samples and Alternative Distance Thresholds

My goal in this section is to find out whether my results may be influenced by assumptions and choices made in creating the final sample. To this purpose I rerun several regression with altered samples and conditions.

First of all I consider the influence outliers may have on my results. In my main sample I do not control for any 'extreme' forecast values. Therefore there are some relatively large values of the dependent (PMAFE) that may have a considerable impact on the reported findings. Nevertheless, it has been a conscious decision to keep all observations in the main sample after examining my sample. Mainly by taking a closer look at the forecasts that rendered those high values of PMAFE. In the end the most logical reason was to include those forecasts too, for the three following reasons. Firstly, the high-PMAFE observations appear in a limited number of firm periods, where multiple analysts have issued forecasts with large forecast errors. Secondly, a small number of 'bad' analysts show up repeatedly in the 1-2% largest relative forecast errors. The two former reasons, along with the absence of strange values of the predicted EPS, supports me in the conviction that the 'extreme' values of PMAFE are not a result of data errors. Hence, the choice to keep all forecasts in the main sample, as they hold relevant information about 'bad' analysts. However, that does not take away the need to test whether it are exactly these observations that have driven the results reported thus far. Simply winsorising the PMAFE variable, however, is not an option as it would render its mean non-zero. Alternatively, I use a trimmed sample, in which for each firm period the 10% lowest forecasts as well as the 10% of forecasts with the highest estimated EPS are dropped. After which I de-mean all the variables by their new forecast period means, so that the regression properties are kept. As a result 800 'extreme' forecasts are dropped from the main sample to create this new sample.

Besides this sample where the top and bottom observations are dropped, I also use a sample without forecast periods. In my main sample I introduce a specific set of forecast periods that suits my sample well. However, in similar research it is most common to merely use an analyst's latest forecast per firm period. Although my setup of forecast periods does not suggest there to be any issues with the reported results[37], I also rerun the regression of the expanded model using a smaller sample without those forecast periods. This implies that from the main sample all but the analyst's latest forecast per firm period are dropped.

I also consider two alternative distance thresholds for the New York variables. The analysts in the main sample are located at 33 different locations according to their area codes. In using area codes to pinpoint the analysts' locations, there obviously is a margin of error in the estimation of the exact location of the analyst's office. Although it does give a reliable approximation, I want to examine whether my results are sensitive to the distance requirement of the New York dummies. In the main sample, I have chosen to stick with the same distance threshold of 100 kilometres that is used for the locality variable. This threshold is well established in previous research and heavily depends on the notion of the analysts' increased ability to make house calls to the firms they visit. While the idea for the New York effect is very similar, there is a big difference between proximity to a company's headquarters and to a financial centre. It may well be that the New York effect is sensitive to the distance threshold I set. For instance, informational advantages may be reserved to those analysts that are actually located in the Manhattan area, rather than those in surrounding areas. Table 7 (in the appendix) gives an overview of the eight analyst locations closest to the financial district in Downtown Manhattan. Based on those figures I introduce two alternative distance thresholds for the New York variables. Namely zero kilometres, solely focussing on those analysts located in the city. And 50 kilometres as an intermediate. The threshold for the locality dummy (with respect to the covered firm) remains at 100 kilometres.

5.2 Results

To test whether the results in the main text are sensitive to any of the concerns brought up in the previous section, I rerun regressions of the expanded model from Eq. (2). For brevity, I only discuss whether and how the measured New York effect is affected. The detailed regression results are depicted in Table 8 through 11 (see appendix).

When rerunning the regression of the expanded model from Table 4, the conclusions are unchanged when using the trimmed sample. All three New York dummy variables still produce

[37] It merely allows for a larger number of forecasts to be evaluated, while additionally adding a new dimension by comparing forecasts issued within the same timeframe through the use of forecast period means.

insignificant coefficients, although they are negative in sign now. When using the sample without forecast periods the three dummy coefficients remain positive and insignificant.[38]

My results seem to be somewhat more sensitive to the distance threshold I set. In the 50 km sample, the New York effect is now significant at the 5% level, indicating that those analysts are 5.9% less accurate than others. When only Manhattan-based analysts are counted as New Yorkers (0 km), the New York variable's coefficient becomes negative in sign, though still statistically insignificant. For New York analysts covering non-local stocks, however, this result is statistically significant (DNY = -5.7%, t = -1.97).

5.2.1 Breakdown Results

Given previous results I want to extend my analysis to also include the breakdown of the sample by characteristics. This means I also rerun the breakdown regressions from Table 5 using the trimmed sample and the sample where the distance threshold is set to zero kilometre. Besides in the main text I found the strongest results after breaking down the sample by the analyst's coverage and status, it is interesting to see whether the same patterns emerge in these samples. Patterns that saw a positive New York effect present among more active, specialised and high status analysts.

Although the results using the trimmed sample are less strong than in the main sample, a similar patterns arises. Only among high-coverage analysts there is a significant accuracy advantage for New York analysts (DNYC = -7.8%, t = 1.74). Also when breaking down the sample by analyst specialisation, a pattern similar to the main sample emerges. The breakdown by analyst status does not render any significant results or pattern. Finally, the breakdown of the trimmed sample by forecast signal indicates that New York analysts are 12.6% more accurate when making unambiguously downward forecast revisions (t = -1.86). A result that was not found in the main sample.

The results of the same breakdown using the zero kilometre sample renders similar results. Especially the pattern of the New York effect coefficients in the coverage breakdown and the breakdown by specialised versus others is clear. The pattern of the New York effect broken down by analyst status, again, is not as prominent as it was in the main sample.

[38] In this sample the coefficient of the age variable (DAGE) becomes strongly significant (and positive). A result that was to be expected as in this sample without forecast periods, DAGE now captures the trend of forecast errors over the entire range of forecast horizons up to one year. Whereas in the main sample DAGE only captures the trend over time within the forecasts periods, with the forecast periods themselves in a way taking away the need to control for the trend from forecast periods to forecast period.

5.3 Implications

On a whole, the results of using the trimmed sample and zero kilometre sample show that my results are in fact sensitive to these changes. In the regular regressions of the expanded model the DNYC coefficient becomes negative in both cases, rather than positive in the main sample. However the coefficients remain insignificant. The results after breaking down the sample, however, are more interesting. First of all, the results using the two alternative samples do confirm the notion of a New York information advantage driven by private information. Namely, I find that in every sample, among high coverage analysts, New Yorkers are more accurate. Similarly, the results after breaking the sample down by specialisation are unchanged, and even somewhat stronger in the trimmed sample.

However, this does not hold for the results in the main sample that indicated an increased cost of information gathering in New York. In the trimmed sample I cannot replicate any significant results that indicate worse performance by New Yorkers, like I find in the main sample. In the 'Manhattan only' sample, I do find worse New York performance in the two lower coverage terciles. Although the magnitude of the effect in the lowest coverage third is notably smaller in this sample, when compared to the main sample results.

Above all, this shows that a limited number of analysts and forecasts are of influence on my results. In the case of the distance threshold this only concerns the matter of defining whether seven analysts fall in the New York category.[39] In the case of the trimmed sample, it proves that 'extreme forecasts' can distort test outcomes. The question, then, remains how to handle these observations. As I explained earlier, I have chosen to keep all observations in the main sample. For the simple reason that I have no reason to assume that those observations are the result of data errors. In the end it comes down to the fact that bad analysts with big forecast errors in my sample are equally representative as the better performing analysts are. The most important conclusion from this section is that the results indicative of a New York information advantage through private information is robust to alternative distance thresholds and the omission of 'extreme' forecasts.

[39] In Table 7 (see appendix) it can be seen that only 154 forecasts by seven analysts are no longer deemed to be originating from New York in the zero kilometre threshold sample.

6. Conclusion

This paper tries to establish a link between analyst performance and their residence in New York City. To this purpose I analyse the forecast accuracy of 299 analysts using close to 5000 yearly and quarterly forecasts on 27 Dow Jones firms' earnings, controlling for a number of analyst characteristics proven to be of influence on said forecast accuracy in prior academic literature. While I am able to confirm the classic locality effect from Malloy (2005), I do not find evidence of a New York effect. However I do produce some significant results after breaking down my sample by a number of characteristics. I find both significantly worse (among low-status analysts and low-coverage analysts) and better performance (among specialised analysts and analysts who cover a relatively large number of stocks) in those regressions. In the end, only the better performance of New York analysts among high-coverage analysts proves to be robust to all model and sample alterations. This result shows that among the third most active analysts New Yorkers are 14.3% more accurate than others, translating in a $0.014 per share accuracy advantage.

For the most part my analysis is analogous to the one used in Malloy (2005). To this I, obviously, add variables to analyse the New York effect. Additionally, I introduce five forecast periods based on the pattern of forecast issuance by the analysts in the sample. The forecast periods prove to be an effective manner to control for differences in forecast horizon and inherent information asymmetry. Apart from the locality effect from Malloy (2005) I am able to confirm four more characteristics as having a positive effect on an analyst's forecast accuracy, being: the size of the analyst's brokerage firm, the analyst's experience, the specialisation of the analyst in a major industry and the analyst's effort.

I believe my research opens up possibilities to further research. They can roughly be sorted into two categories: extensions onto this research and new ideas to take the research in new directions. The former category touches on some of the shortcomings of my research. Most notably this considers the relatively small data set which consists of (27) large high profile firms exclusively. A larger sample that also includes smaller, less covered firms would pose as an interesting testing ground to replicate my analysis. Moreover, it should be taken into account that the New York effect might be more pronounced in a dataset with smaller, less visible stocks. The reasoning behind this is that it is more likely for information asymmetries to exist when it concerns smaller, less visible stocks.[40] A larger sample also opens possibilities to explore the influence of firm characteristics on a possible New York effect. Additionally, a dataset could span for a number of years. Finally, apart from analysing forecast errors, the research could be broadened to also include

[40] To mark this point, Malloy (2005) finds that his local advantage is mainly concentrated in smaller stocks. For large stocks, which my sample solely consists of, he does not find such a local advantage. Indicating that asymmetric information is less prominent, or even absent, in large stocks.

the impact that (New York) forecasts have on stock prices to be able to draw a more definitive conclusion in the matter of any information asymmetries.

However my research also brings up some new ways to approach this subject. Something that all New York analysts in my sample have in common is that they have a relatively large number of nearby located colleagues covering the same stock. A different approach could be to compare forecasts originating from analyst-dense areas to others. This way you automatically select the areas where there is an active community of analysts following a stock. Alternatively, other financial districts could be investigated as well. Extending on this last point, it may be interesting to investigate whether the New York effect is more pronounced if a large proportion of the analysts covering a stock are situated in New York.

References

Abarbanell, J. S. and Bernard, V. L., 1992, Tests of Analysts' Overreaction/Underreaction to Earnings Information as an Explanation for Anomalous Stock Price Behavior, *Journal of Finance* 47, 1181–1207.

Agrawal, A., Chadha, S. and Chen, M. A., 2006, Who is Afraid of Reg FD? The Behavior and Performance of Sell-Side Analysts Following the SEC's Fair Disclosure Rules, *Journal of Business* 79, 2811-2834.

Bacmann, J. and Bolliger, G., 2001, Who Are the Best? Local Versus Foreign Analysts on Latin American Stock Markets, EFMA 2002 London Meetings.

Bae, K., Stulz, R. M. and Tan, H., 2008, Do local analysts know more? A cross-country study of performance of local analysts and foreign analysts, *Journal of Financial Economics* 88, 581-606.

Bartholdy, J. and Feng, T., 2013, The quality of securities firms' earnings forecasts and stock recommendations: Do informational advantages, reputation and experience matter in China?, *Pacific-Basin Finance Journal*, Volume 24, 66-88.

Beyer, A., Cohen, D. A., Lys, T. Z., & Walther, B. R., 2010, The financial reporting enviornment: Review of the recent literature, *Journal of Accounting and Economics*, 296-343.

Clement, M. B., 1998, Some considerations in measuring analysts' forecasting performance, Working paper, University of Texas.

Clement, M. B., 1999, Analyst forecast accuracy: Do ability, resources, and portfolio complexity matter? *Journal of Accounting and Economics* 27, 285-303.

Coval, J. D., and Moskowitz, T. J., 2001, The geography of investment: Informed trading and asset prices, *Journal of Political Economy* 4, 811-841.

DeBondt, W. F. M. and Thaler, R. H., 1990, Do security analysts overreact? *American Economic Review* 80, 52-57.

Easterwood, J. C. and Nutt, S. R., 1999, Inefficiency in Analysts' Earnings Forecasts: Systematic Misreaction or Systematic Optimism?, *Journal of Finance* 54, 1777–1797.

Fang, L. H. and Yasuda, A., 2014, Are Stars' Opinions Worth More? The Relation between Analyst Reputation and Recommendation Values, *Journal of Financial Services Research* 46, 235-269.

Gleason, C. A., and Lee, C. M.C., 2003, Analyst forecast revisions and market price discovery, *Accounting Review* 78, 193-225.

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

Hong, H., Kubik, J. D. and Stein, J. C., 2005, Thy Neighbor's Portfolio: Word-of-Mouth Effects in the Holdings and Trades of Money Managers. *Journal of Finance* 60, 2801–2824.

Jackson, A. R., 2005, Trade Generation, Reputation, and Sell-Side Analysts. *The Journal of Finance* 60, 673–717.

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

Klettke, T., Homburg, C. and Gell, S., 2014, How to Measure Analyst Forecast Effort, *European Accounting Review*, Forthcoming.

Leone, A. J. and Wu, J. S., 2007, What Does it Take to Become a Superstar? Evidence from Institutional Investor Rankings of Financial Analysts, Simon School of Business Working Paper.

Lin, H., and McNichols, M. F., 1998, Underwriting relationships, analysts' earnings forecasts and investment recommendations, *Journal of Accounting and Economics* 25, 101-127.

Louis, H., Sun, A. X. and Urcan, O., 2013, Do Analysts Sacrifice Forecast Accuracy for Informativeness?, *Management Science* 59:7, 1688-1708.

Malloy, C. J., 2005, The Geography of Equity Analysis, *Journal of Finance* 60, 719–755.

Milgrom, P. and Roberts, J., 1992. Economics, Organization & Management. Prentice Hall, Englewood Cliffs, NJ.

Ramnath, S., Rock, S. and Shane, P. B., 2008, The Financial Analyst Forecasting Literature: A Taxonomy with Suggestions for Further Research, *International Journal of Forecasting* 24, 34-75.

Stickel, S. E., 1990, Predicting Individual Analyst Earnings Forecasts, *Journal of Accounting Research* 28.2, 409–417.

Stickel, S. E., 1992, Reputation and performance among security analysts, *Journal of Finance* 47, 1811-1836.

7. Appendix

Variable Definitions

AFFIL	is a dummy variable equal to one if analyst *i*'s brokerage firm has an underwriting relationship with the covered firm *j*.
AGE	equals the natural logarithm of 1 plus the age of analyst *i*'s forecast, where the age is equal to the number of days between the issue date of the forecast and the date the actual earnings are released by the firm.
COV	equals the natural logarithm of 1 plus the number of unique firms analyst *i* covers in 2007.
DAFE	(de-meaned absolute forecast error) equals analyst *i*'s absolute forecast error minus the mean absolute forecast error for firm *j* in forecast period *t*. Where absolute forecast error equals the absolute value of an analyst's latest forecast minus actual company earnings, as a percentage of the stock price.
EXP	equals the natural logarithm of 1 plus the number of years in which analyst *i* has issued at least one earnings forecast, for any random stock.
FEXP	equals the natural logarithm of 1 plus the number of years in which analyst *i* has issued at least one earnings forecast for the covered firm *j*.
FREQ	equals the natural logarithm of 1 plus the mean number of forecasts analyst *i* issues per covered company in 2007.
HIGH	is a dummy variable equal to one if analyst *i* works at a high-status brokerage firm according to the rankings in *Institutional Investor*.
LOC	is a dummy variable equal to one if analyst *i* is located within 100 kilometres of the headquarters of the covered firm *j* and more than 100 kilometres away from New York City.
LOGD	equals the natural logarithm of 1 plus the physical distance (in kilometres) between analyst *i* and New York City.
NY	is a dummy variable equal to one if analyst *i* is located within 100 kilometres of New York City and more than 100 kilometres away from headquarters of the covered firm *j*.
NYC	is a dummy variable equal to one if analyst *i* is located within 100 kilometres of the financial centre in New York City.
NY_LOC	is a dummy variable equal to one if analyst *i* is located within 100 kilometres of New York City and within 100 kilometres of the headquarters of the covered firm *j*.
PCTOP	equals the percentage of firm periods for which analyst *i*'s last forecast is above actual earnings, looking at all companies followed by analyst *i* in 2007.
PMAFE	(proportional mean forecast error) equals DAFE divided by the mean absolute forecast error for firm *j* in forecast period *t*.
ROPT	equals analyst *i*'s forecast error (FE_{ijt}) minus the forecast period mean forecast error (MFE_{jt}), divided by the absolute value of MFE_{jt}.
SIZE	equals the natural logarithm of 1 plus the number of analysts issuing earnings forecasts in 2007 at analyst *i*'s brokerage firm.
SPEC	is a dummy variable equal to one if analyst *i* only covers firms in a single major industry group.
STAR	is a dummy variable equal to one if analyst *i* is an All-American analysts according to *Institutional Investor*.

Table 7: Overview of the Eight Locations Closest to the Financial District in New York

This table gives an overview of the eight areas in my sample that are located closest to the financial district in Downtown Manhattan. The reported distances are computed as the crow flies.

Area Code	Location	Distance to New York (km.)	Number of Forecasts	Number of Analysts
212/917	Manhattan (NYC), NY	0	3388	192
973	Paterson, NJ	26	13	1
914	Yonkers, NY	29	30	1
732	Rumson Borough, NJ	38	28	1
631	Babylon Village, NY	58	44	2
203	Bridgeport, CT	87	39	2
401	Providence, RI	250	33	1
410/443	Baltimore, MD	271	57	4

Table 8: Rerunning the Expanded Regression Model with Alternative Samples

In this table I report the results of regressing the expanded model (Eq. 2) using two alternative samples. The results of the expanded model on the left-hand side are the same that I reported in Table 4. In the trimmed sample the 10% lowest forecasts as well as the 10% of forecasts with the highest estimated EPS are dropped for each firm period. The other alternative sample is one without the use of forecast periods; only the analyst's last forecast per firm period is kept. All variables are defined as before and mean adjusted by their respective forecast (or firm) period means (indicated by D for de-meaned). All coefficients report the expected change of the dependent variable in percentages. The robust t-statistics are in parentheses.

	Original Sample		Trimmed Sample		No Forecast Periods	
Indep. Var.	PMAFE	PMAFE	PMAFE	PMAFE	PMAFE	PMAFE
DLOC	-18.2164****	-18.6889****	-14.6599***	-15.2610***	-21.6422***	-23.4714***
	(-3.10)	(-3.14)	(-2.30)	(-2.38)	(-2.10)	(-2.24)
DAGE	8.0628	8.0894	13.5822**	13.6169**	15.9146****	15.9390****
	(1.06)	(1.06)	(1.77)	(1.77)	(6.44)	(6.46)
DSIZE	-8.6692****	-8.6784****	1.2133	1.2029	-7.5668**	-7.6328**
	(-3.79)	(-3.79)	(1.00)	(0.99)	(-1.75)	(-1.77)
DSTAR	3.3702*	3.3850*	0.3761	0.4070	2.2688	2.3211
	(1.47)	(1.48)	(0.17)	(0.19)	(0.57)	(0.58)
DHIGH	-1.5770	-1.5447	-2.5462	-2.5150	1.5368	1.6844
	(-0.68)	(-0.67)	(-1.17)	(-1.15)	(0.37)	(0.40)
DEXP	-5.3357***	-5.2920***	-3.8705*	-3.8191*	2.6216	2.7898
	(-2.05)	(-2.03)	(-1.47)	(-1.45)	(0.53)	(0.57)
DFEXP	0.0046	0.0060	2.8828	2.8930	-6.9404*	-6.9482*
	(0.00)	(0.00)	(1.17)	(1.17)	(-1.49)	(-1.49)
DPCTOP	-0.1430	-0.1404	0.1333	0.1354	0.0748	0.0841
	(-1.07)	(-1.05)	(1.11)	(1.13)	(0.38)	(0.43)
DAFFIL	7.0074	6.8441	5.7803	5.5404	-1.8036	-2.3483
	(1.06)	(1.04)	(0.70)	(0.67)	(-0.16)	(-0.21)
DCOV	-4.1644	-4.1450	0.4894	0.5039	-2.0670	-1.9820
	(-1.29)	(-1.29)	(0.17)	(0.17)	(-0.36)	(-0.34)
DSPEC	-10.9866****	-11.0519****	-6.3556***	-6.4465***	-9.8342*	-9.9867*
	(-3.08)	(-3.10)	(-2.15)	(-2.18)	(-1.60)	(-1.63)
DFREQ	-4.7504***	-4.8638***	-5.7671***	-5.9194***	-7.8618**	-8.2966**
	(-1.98)	(-2.03)	(-2.43)	(-2.50)	(-1.67)	(-1.76)
DNYC	3.4395		-2.3200		3.1260	
	(1.23)		(-1.02)		(0.62)	
DNY		2.7201		-3.2617		0.3080
		(0.85)		(-1.27)		(0.05)
DNY_LOC		5.0212		-0.1512		9.0104
		(1.21)		(-0.04)		(1.33)
Observations	4,985	4,985	4,185	4,185	2,209	2,209
Adjusted R-squared	0.018	0.018	0.005	0.005	0.036	0.036

Robust t-statistics in parentheses
**** p<0.01, *** p<0.05, ** p<0.10, * p<0.15

Table 9: Rerunning the Expanded Regression Model with Alternative Distance Thresholds

This table reports the results of regressing the expanded model (Eq. 2) using two alternative distance thresholds for the New York dummies. Pooled cross-sectional regressions of proportional mean absolute forecast error (PMAFE) on analyst characteristics are run, following the expanded model (Eq. 2). The results of the expanded model on the left-hand side are the same that I reported in Table 4 and Table 8. The new thresholds are 50 kilometres and 0 zero kilometres and are used to indicate whether an analyst is deemed an New York analyst or not. For brevity, only the New York dummy variable coefficients are depicted. All variables are defined as before and mean adjusted by their respective forecast (or firm) period means (indicated by D for de-meaned). All coefficients report the expected change of the dependent variable in percentages. The robust t-statistics are in parentheses.

Indep. Var.	Original (100 km.)		50 kilometers		0 kilometers	
DNYC	3.4395		5.8667***		-2.3414	
	(1.23)		(2.02)		(-1.00)	
DNY		2.7201		5.4825**		-5.6990***
		(0.85)		(1.67)		(-1.97)
DNY_LOC		5.0212		6.7315*		5.4763
		(1.21)		(1.62)		(1.34)
Observations	4,985	4,985	4,985	4,985	4,985	4,985

Robust t-statistics in parentheses

**** p<0.01, *** p<0.05, ** p<0.10, * p<0.15

Table 10: Breakdown of the Trimmed Sample by Characteristics

This table reports the results of rerunning the regressions from Table 5 using the trimmed sample from Table 8 instead of the main sample. Pooled cross-sectional regressions of proportional mean absolute forecast error (PMAFE) on analyst characteristics are run, following the expanded model (Eq. 2). The sample is broken down by a number of characteristics and only the coefficients and t-statistics of the location dummies are reported. As such, for each subsample two regressions are run; one with the single New York dummy (NYC) and also one with NY and NY_LOC as location dummies. PMAFE, NYC, NY, NY_LOC are defined as in Table 2. In panel A SIGNAL = 1 if a forecast is above analyst *i*'s prior forecast for that firm period and also above the prior consensus. SIGNAL = -1 if analyst *i* revises the forecast both below his prior forecast and the prior consensus. All other forecast fall into the third category (SIGNAL = 0). In the STAR column only All-Star analysts are included, while the NON column excludes those. In Panel B Low-coverage analysts are analysts who cover less firms than the 33rd percentile of firms covered by the analysts in the sample. Analysts who are above 66th percentile are classified as High-coverage analysts. The remaining analysts fall in the Mid-coverage category. In the SPEC column only specialised analysts covering stocks in only one major industry group are included, while the NON column excludes those. All variables are mean adjusted by their respective forecast period means (indicated by D for de-meaned). All coefficients report the expected change of the dependent variable in percentages. The robust t-statistics are in parentheses.

		SIGNAL			Analyst Status	
Dep. Var.	Indep. Var.	1	-1	0	STAR	NON
PMAFE	DNYC	-2.5672	-12.5556**	-0.0308	-2.9537	-0.1974
		(-0.55)	(-1.86)	(-0.01)	(-0.70)	(-0.07)
PMAFE	DNY	-1.8336	-8.6693	-1.6564	-5.1305	-1.4336
		(-0.36)	(-1.24)	(-0.51)	(-0.87)	(-0.48)
PMAFE	DNY_LOC	-4.7251	-30.9527	3.1318	0.5775	3.0284
		(-0.55)	(-1.18)	(0.71)	(0.11)	(0.51)
Obs.		1,026	425	2,486	1,303	2,388

Panel B: Breakdown by Coverage and Brokerage Status

		Coverage			Specialisation	
Dep. Var.	Indep. Var.	Low	Mid	High	SPEC	NON
PMAFE	DNYC	3.4623	-3.9235	-7.7575**	-13.4050**	-1.6906
		(0.73)	(-0.94)	(-1.74)	(-1.78)	(-0.70)
PMAFE	DNY	7.9684	-5.6420	-10.8396***	-12.7267*	-3.1737
		(1.26)	(-1.03)	(-2.33)	(-1.62)	(-1.19)
PMAFE	DNY_LOC	-2.3754	-1.3230	5.8898	-16.6582	1.5694
		(-0.37)	(-0.20)	(0.61)	(-1.30)	(0.35)
Obs.		1,307	890	972	655	3,314

Robust t-statistics in parentheses

**** $p<0.01$, *** $p<0.05$, ** $p<0.10$, * $p<0.15$

Table 11: Breakdown of the Zero Kilometre Sample by Characteristics

This table reports the results of rerunning the regressions from Table 5, with the one difference that the New York dummy threshold is set to zero kilometre instead of the 100 kilometres used in the main text. Pooled cross-sectional regressions of proportional mean absolute forecast error (PMAFE) on analyst characteristics are run, following the expanded model (Eq. 2). The sample is broken down by a number of characteristics and only the coefficients and t-statistics of the location dummies are reported. As such, for each subsample two regressions are run; one with the single New York dummy (NYC) and also one with NY and NY_LOC as location dummies. PMAFE, NYC, NY, NY_LOC are defined as in Table 2. In panel A SIGNAL = 1 if a forecast is above analyst *i*'s prior forecast for that firm period and also above the prior consensus. SIGNAL = -1 if analyst *i* revises the forecast both below his prior forecast and the prior consensus. All other forecast fall into the third category (SIGNAL = 0). In the STAR column only All-Star analysts are included, while the NON column excludes those. In Panel B Low-coverage analysts are analysts who cover less firms than the 33rd percentile of firms covered by the analysts in the sample. Analysts who are above 66th percentile are classified as High-coverage analysts. The remaining analysts fall in the Mid-coverage category. In the SPEC column only specialised analysts covering stocks in only one major industry group are included, while the NON column excludes those. All variables are mean adjusted by their respective forecast period means (indicated by D for de-meaned). All coefficients report the expected change of the dependent variable in percentages. The robust t-statistics are in parentheses.

		SIGNAL			Analyst Status	
Dep. Var.	Indep. Var.	1	-1	0	STAR	NON
PMAFE	DNYC	-8.2326*	3.8779	1.1424	-1.1038	1.1421
		(-1.64)	(0.65)	(0.40)	(-0.26)	(0.42)
PMAFE	DNY	-12.5892**	5.0770	-1.9994	-6.9527	-1.0417
		(-1.96)	(0.77)	(-0.59)	(-1.27)	(-0.34)
PMAFE	DNY_LOC	3.4049	-1.0054	7.8308**	9.7026**	6.9100
		(0.38)	(-0.08)	(1.71)	(1.87)	(1.26)
Obs.		1,343	616	2,731	1,635	2,871

Panel B: Breakdown by Coverage and Brokerage Status

		Coverage			Specialisation	
Dep. Var.	Indep. Var.	Low	Mid	High	SPEC	NON
PMAFE	DNYC	8.1596**	7.3031**	-12.7430****	-11.4380*	-0.6325
		(1.80)	(1.80)	(-2.62)	(-1.48)	(-0.26)
PMAFE	DNY	-2.2162	7.5818*	-13.7262***	-10.3628	-0.8728
		(-0.37)	(1.45)	(-2.56)	(-1.25)	(-0.31)
PMAFE	DNY_LOC	29.0371****	6.8718	-9.4268	-15.9400	-0.0875
		(4.14)	(1.22)	(-1.11)	(-1.25)	(-0.02)
Obs.		1,687	1,074	1,228	772	3,992

Robust t-statistics in parentheses
**** p<0.01, *** p<0.05, ** p<0.10, * p<0.15