

AN INVESTIGATION OF THE IMPACT OF ELECTRONIC  
COMMUNICATION ON FINANCIAL MARKETS

by

ROBERT CHARLES GIANNINI

(Under the direction of Paul Irvine)

ABSTRACT

I use a unique set of Twitter posts to analyze the impact of network communication on financial markets. In the first chapter, I show that network communication does impact firms' prices, and that an individual's influence is a function of the size of their network and their creditability. In the second chapter, I use Twitter posts and firm press releases to measure disagreement before earnings announcements. I show that disagreement with firm press releases before an earnings announcement implies future negative abnormal price drift. Finally, I use Twitter users' locations to test whether they possess a local informational advantage. I show that the differential stock return predictability between local and nonlocal Twitter posts is 19 basis points per week. Additionally, local advantage concentrates in firms without public news coverage and firms with greater information asymmetry, indicating that local advantage is attributable to individual investors' access to private information.

INDEX WORDS:     Twitter, Post earnings announcement drift, Divergence of opinions,  
                          Local information advantage, Information diffusion, Bounded  
                          rationality

AN INVESTIGATION OF THE IMPACT OF ELECTRONIC  
COMMUNICATION ON FINANCIAL MARKETS

by

ROBERT CHARLES GIANNINI

B.A., Emory University, 2000

A Dissertation Submitted to the Graduate Faculty  
of The University of Georgia in Partial Fulfillment  
of the

Requirements for the Degree

DOCTOR OF PHILOSOPHY

ATHENS, GEORGIA

2011

©2011

Robert Charles Giannini

All Rights Reserved

AN INVESTIGATION OF THE IMPACT OF ELECTRONIC  
COMMUNICATION ON FINANCIAL MARKETS

by

ROBERT CHARLES GIANNINI

Approved:

Major Professors: Paul Irvine

Committee: Paul Irvine  
Tyler Henry  
Tao Shu

Electronic Version Approved:

Maureen Grasso  
Dean of the Graduate School  
The University of Georgia  
December 2011

# **An Investigation of the Impact of Electronic Communication on Financial Markets**

Robert Charles Giannini

November 29, 2011

# Acknowledgments

"I met a gypsy and she hipped me to some life game To stimulate then activate the left and right brain Said baby boy you only funky as your last cut You focus on the past your ass'll be a has what Thats one to live by or either that one to die to I try to just throw it at you determine your own adventure"

—André Lauren Benjamin—

# Contents

|          |                                                                                    |           |
|----------|------------------------------------------------------------------------------------|-----------|
| <b>1</b> | <b>Twitter: An Investigation of the Impact of Network Communication</b>            | <b>8</b>  |
| 1        | Introduction . . . . .                                                             | 9         |
| 2        | Previous Literature . . . . .                                                      | 11        |
| 3        | Data Collection and Variable Construction . . . . .                                | 19        |
| 4        | Relating the Level of Network Communication to Stock Returns . . . . .             | 27        |
| 5        | Posting Frequency and Sentiment . . . . .                                          | 31        |
| 6        | Predicting Information Flow through Networks . . . . .                             | 35        |
| 7        | Network Communication: A Direct Measure of Attention . . . . .                     | 38        |
| 8        | Conclusion . . . . .                                                               | 40        |
| 9        | References . . . . .                                                               | 43        |
| <b>2</b> | <b>The Impact of Divergence of Opinions about Earnings within a Social Network</b> | <b>55</b> |
| 1        | Introduction . . . . .                                                             | 56        |
| 2        | Theory and Hypothesis Development . . . . .                                        | 61        |
| 3        | Data Collection and Variable Construction . . . . .                                | 64        |
| 4        | Results . . . . .                                                                  | 78        |
| 5        | What causes divergence of opinion? . . . . .                                       | 86        |
| 6        | Conclusion . . . . .                                                               | 89        |

|          |                                                                                                   |            |
|----------|---------------------------------------------------------------------------------------------------|------------|
| <b>3</b> | <b>Do Local Investors Know More? A Direct Examination of Individual Investors Information Set</b> | <b>112</b> |
| 1        | Introduction . . . . .                                                                            | 113        |
| 2        | Data and Sample Selection . . . . .                                                               | 119        |
| 3        | Do Individual Investors Have Local Advantage? . . . . .                                           | 127        |
| 4        | Is Local Advantage Private Information or Investor Sophistication? . . . . .                      | 134        |
| 5        | Conclusion . . . . .                                                                              | 140        |
| 6        | References . . . . .                                                                              | 141        |

# List of Figures

|     |                                                                          |     |
|-----|--------------------------------------------------------------------------|-----|
| 1.1 | www.Stocktwits.com . . . . .                                             | 46  |
| 1.2 | Information Flow about IBM . . . . .                                     | 47  |
| 2.1 | www.Stocktwits.com . . . . .                                             | 94  |
| 2.2 | Divergence of Opinions Event Study . . . . .                             | 95  |
| 2.3 | Divergence of Opinions Event Study top and bottom SUE terciles . . . . . | 96  |
| 2.4 | Divergence of Opinions Event Study top and bottom SUE terciles . . . . . | 97  |
| 2.5 | Divergence of Opinions Event Study . . . . .                             | 98  |
| 2.6 | <b>Predicted signs of CARs conditional on opinions</b> . . . . .         | 99  |
| 2.7 | Changes of opinions around earnings . . . . .                            | 100 |
| 2.8 | Changes of opinions around earnings: Volume Test . . . . .               | 101 |
| 3.1 | <b>Example of Twitter Stream</b> . . . . .                               | 147 |
| 3.2 | <b>Summary Statistics</b> . . . . .                                      | 147 |
| 3.3 | <b>Geographical Distribution of Sample Twitter Users</b> . . . . .       | 148 |

# List of Tables

|     |                                                                   |     |
|-----|-------------------------------------------------------------------|-----|
| 1.1 | Summary Statistics for Sample Firms and Communication Variables   | 48  |
| 1.2 | Contemporaneous and Lagged Communication . . . . .                | 49  |
| 1.3 | Contemporaneous and Lagged <i>Sentiment_Impact</i> . . . . .      | 50  |
| 1.4 | Returns <i>Senti_Impact</i> Across Size Groups . . . . .          | 51  |
| 1.5 | Returns to Local & Non-Local Senti(SP500 Members) . . . . .       | 52  |
| 1.6 | Information Flow Through a Network . . . . .                      | 53  |
| 1.7 | Factor Regression for Impact and Sentiment Impact . . . . .       | 54  |
| 1.8 | Attention vs No Attention . . . . .                               | 54  |
| 2.1 | Summary Statistics . . . . .                                      | 102 |
| 2.2 | Correlation Table . . . . .                                       | 104 |
| 2.3 | Divergence of Opinion and Returns . . . . .                       | 105 |
| 2.4 | Divergence of Opinion and Returns (Berkman et al. Controls) . .   | 106 |
| 2.5 | Divergence of Opinion and Returns (Garfinkel et al. Controls) . . | 107 |
| 2.6 | Divergence of Opinions Event Study . . . . .                      | 108 |
| 2.7 | Divergence of Opinion Event Study (SUE Ranks) . . . . .           | 109 |
| 2.8 | Changes in Opinions . . . . .                                     | 110 |
| 2.9 | Linguistic Characteristics of News . . . . .                      | 111 |
| 3.1 | Summary Statistics . . . . .                                      | 149 |

|      |                                                                                                      |     |
|------|------------------------------------------------------------------------------------------------------|-----|
| 3.2  | Panel Regressions of Abnormal Stock Returns on Investor Evaluations                                  | 150 |
| 3.3  | Panel Regressions of Stock Returns on Local and Nonlocal Evaluations                                 | 151 |
| 3.4  | Panel Regressions of Stock Returns: Alternative Construction/ Sample Selection . . . . .             | 152 |
| 3.5  | Panel Regressions of Stock Returns: State Adjusted Returns . . .                                     | 154 |
| 3.6  | Panel Regressions of Stock Returns: Alternative Measurements of Evaluations or Returns . . . . .     | 155 |
| 3.7  | Panel Regressions of Trading Volumes . . . . .                                                       | 156 |
| 3.8  | Panel Regressions of Stock Returns: Stocks without Public News vs. Stocks with Public News . . . . . | 157 |
| 3.9  | Local Advantage across Proxies of Information Asymmetry . . . .                                      | 158 |
| 3.10 | Daily Abnormal Profits (%) of Rolling Long-Short Strategies Based on Local Advantage . . . . .       | 160 |

# Chapter 1

## Twitter: An Investigation of the Impact of Network Communication

### **Abstract**

This paper uses the popular social networking site Twitter to test recent theories on the importance of communication through networks. I create a unique data set that precisely measures the network size and credibility of each contributor to the site. I test the theory of DeMarzo et al (2003) that an agent's influence is a function of his credibility and the size of his network, as well as the Colla and Mele (2010) theory of correlated trading among social groups. I show that influence is a function of network credibility and network size. I also find that there is positively correlated trading among agents in a social network and negatively correlated trading with agents outside of the network. Finally, using portfolios sorted on level of attention, I find support for the Merton (1987) theory of investor recognition which predicts that low attention stocks should exhibit positive risk adjusted returns.

# 1 Introduction

Today, electronic social networks are a central part of most people's lives, and interaction through these networks is becoming increasingly important to peoples thoughts and opinions. Although there has been a number of papers written on communication, attention and social networks, this research has only used proxies for the level of communication through a network, such as the number of news stories published about a company, the frequency of chat board posts, or the level of Google search activity about a company. However, these proxies cannot directly measure the flow of information between agents within a social network. This paper uses the website Stocktwits.com to compile a unique set of firm specific communication and network data that enables the precise measurement of how communication through a social network affects individual firms.

Stocktwits.com is an offshoot of Twitter, and is a social network for people who want to share information about financial markets. The website's slogan is "Real Investors. Real Ideas. Real Time". The output of the site is an aggregated feed of market commentary from all Twitter users in the network. In essence, it is a real time stream of investors' thoughts and opinions. The updates that investors post on the site are by definition constrained to be 140 characters and they usually contain a tag of the specific ticker that is being referenced in the post. These updates differ from the posts found on stock message boards because they highlight the "hot" news from the markets rather than the news of one particular company.

The global popularity of Twitter shows that people do use the site for communication, but is the information that they gather timely, reliable, and useful? Can popular authors in a social network affect the opinions of their peers? Does the size of a person's social network determine the impact of information that they disseminate? Is the frequency or sentiment of posts from influential agents related to past, current or future stock prices? Can knowledge about information flow through social networks be used to reliably predict

stock returns? I answer these questions by using computational linguistics (CL) to create a set of explanatory variables which summarize the intensity of information flow through a social network. I use these variables to test whether communication by influential agents affects prices, and whether the impact that an agent has on a network is a function of the size of his following, as well as the accuracy of his information.

In this paper I formally test the theories of Colla and Mele (2010), DeMarzo et al (2003) and Merton (1987). The DeMarzo et al (2003) theory of networking predicts that communication by influential agents in a network should influence the actions of others members of their network. I find support for this theory by showing that a one standard deviation increase in social network communication implies 8-26 bps of contemporaneous abnormal returns. Colla and Mele(2010) predict that in the short run socially close agents will execute positively correlated trades, but future trades by other market participants will be negatively correlated to these initial trades. I find support for the Colla and Mele(2010) theory of correlated trading by showing that a positive shock to the sentiment of information flow predicts positive contemporaneous abnormal returns and negative future abnormal returns over the following five to twenty trading days. Specifically, a one standard deviation increase in network sentiment over the (-1,-20) event window causes -46 bps of abnormal returns over the next twenty trading days. Thus, as predicted by the Colla and Mele model, future abnormal returns are negatively related to day zero sentiment and abnormal returns. Furthermore, these results are stronger among firms with higher information asymmetry where network information flow should be more valuable. The impact of network communication on small firms is 2.1 % greater than the impact on large firms in the ten trading days after communication.

The Colla and Mele (2010) theory also predicts that the sentiment of communication through a network should be weakly persistent in the shortrun. I investigate the time series dynamics of communication and show that a one standard deviation increase in the sentiment

of communication will increase the following day's sentiment by .08 standard deviations. The impact of lagged sentiment is only significant over the one day horizon, and longer lags of sentiment are not predictive of future network sentiment. Finally, Merton (1987) predicts that firms that have lower levels of investor attention should have higher expected future returns. I find support for this prediction by forming portfolios based on network communication. I show that firms that have no network communication exhibit 7 bps of abnormal returns the day after the portfolio formation.

The rest of the paper is organized as follows. Section 2 is a review of previous empirical literature and lays out the paper's theoretical motivation. Section 3 discusses the nature of the Twitter chat board, how the data is collected from the internet, and how it is organized for further use. Section 3 also explains the sentiment classification procedure and the construction of all explanatory and dependent variables. Sections 4 and 5 test the predictions of two theories of communication. Specifically, I test whether the level and sentiment of communication in a network is related to contemporaneous or future abnormal returns. Section 6 analyzes the time series properties of communication and tests how opinions evolve within a network. Section 7 uses the network communication data to test the predictions of two theories of attention. In Section 8, I conclude.

## 2 Previous Literature

There are a small number of empirical papers that analyze the information contained in informal communication over the internet; all of these papers have focused on message boards such as Yahoo! Finance and Raging Bull. The first paper to systematically analyze internet posting was Wysocki (1999). He shows that in the cross-section firms with higher levels of message board posting have higher analyst coverage, higher valuations relative to fundamentals, and higher volume of trading. He also looks at a time series of the 50 firms with the

highest posting on volume for an 8 month period from January to August 1998 and shows that chat boards do have weak predictive power for future abnormal returns and volume.

Chen and Das (2007) use computational linguistics to analyze the sentiment and frequency of posting about a small subset of technology firms for the period between July-August 2001. They show that the sentiment and frequency of posting can be used to predict future volume and volatility, but cannot be used to reliably predict future returns.

The most complete study of internet message boards is Antweiler and Frank (2004), who analyze the information content of Yahoo! Finance and Raging Bull's internet message boards. They determine that message boards do contain some relevant financial information. They show that a positive shock to the number of messages posted does predict negative stock returns on the next day. Specifically, a doubling of message board posting predicts a -.2% decrease in stock price the next day. This result is statistically significant but not economically relevant. They also show that a day  $t$  increase in board posting coincides with an increase in volume on day  $t$ . However, on day  $t + 1$  the opposite relationship is found, and positive shocks to messages board posting levels actually predict a decrease in volume the next day.

Bollen Mao and Zeng (2011) is the first paper to analyze Twitter as a predictor of stock market fluctuations. This paper measures the average mood of all users of Twitter and relates overall national mood to the performance of the DJIA. They show that a "calm" mood among Twitter users is positively related to future DJIA returns. Furthermore, they use a self-organizing Fuzzy Neural Network to show that some other moods such as happiness can impact future index return. This paper shows that even a noisy measure such as national sentiment does have some predictive power for stock market activity.

Beyond the information found on chat boards, a number of papers have studied the impact of professional news on stock markets. Mitchell and Mulherin (1994) is an early paper which analyzes the affect that public information has on markets. This paper looks

at the frequency of New York Times articles and Dow Jones announcements and shows that there is a weak link between media activity and market volatility and volume. Tetlock (2007) analyzes how the sentiment of news affects the market by analyzing the content of the popular Wall Street Journal column "Abreast of the Market". Tetlock (2007) quantifies the sentiment of the WSJ column by constructing a simple measure of pessimism, and then uses the pessimism factor in vector autoregressions to determine the intertemporal relationship between pessimism and the stock market. Through a series of intertemporal tests, he shows that media content can predict movements in indicators of stock market activity such as volume, returns and asset pricing factors.

More recently, Fang and Peress (2010) show that media coverage does affect the cross section of stock returns. They show that even after controlling for all well known risk factors, stocks with no media coverage earn higher expected returns than stocks with high media coverage. They argue that media coverage may decrease informational problems even if it does not break any genuine news and that the media's affect on stock prices is related to its ability to disseminate information broadly, rather than its ability to form a consensus among investors.

There are a few recent papers that specifically link attention measures to the trading habits of retail investors. Barber and Odean (2008) use abnormal returns, abnormal volume and media coverage as proxies for attention and show that retail investors are net buyers of attention grabbing stocks, but institutional investors are not affected by changes in attention proxies. Da, Engelburg and Gao (working) analyze retail investor attention using the volume of Google searching that occurs during a week. They show that search volume is positively related to returns in the shortrun, and that a price momentum strategy performs significantly better among stocks that also have a high level of investor attention.

## 2.1 Theory and Motivation

In recent years, social networking sites such as Facebook, Myspace and Twitter have changed the way that we interact with each other, exchange information and ideas and keep up to date with news. Even though the new social media phenomenon is a large part of peoples' everyday lives, research on these new forms of information is limited. In this paper I focus on Twitter, one of the most popular forms of digital social interaction. Twitter is a free social networking and micro-blogging service that allows a network of users to send short (140 characters or less) updates called tweets to all of the friends in their network. The power of this service is that one can not only publish one's own thoughts, but one can also read a real time feed of the thoughts of his entire social network. As of January of 2010, Alexa.com, a website that provides statistics on web traffic, rated Twitter as the 12th most popular website on the Internet based on the number of unique users and the number of page hits. On average, approximately 3.7% of all global internet users visit Twitter each day. These statistics clearly highlight that Twitter is a powerful source of popular news and opinions, and that this type of social media is an integral part of many people's lives. But the question still remains, why do people communicate their ideas within a social network?

In the field of finance there are very few theoretical models that predict that people should communicate within social networks. One model of social network communication which answers this question is developed in DeMarzo, Vayanos, and Zwiebel (2003). They show that members of a network overweight the opinions of the people that they speak with regularly. In other words, boundedly rational agents fail to correctly update their expectations because they fail to account for repetition in the information that they receive. Thus, an agent's power of persuasion is a function of the accuracy of his information and the size of his social network. Thus, it can be profitable for an agent to increase the size of his social network by continually communicating within the community. For this reason, it is rational to follow the opinions of well connected individuals because they can persuade a large

number of people to subscribe to their beliefs. This theory provides a rational explanation for why people participate in social communication, and how this communication might be beneficial to them. Colla and Mele (2010) extend the theoretical analysis of DeMarzo et al (2003) to a financial setting and show that communication through financial networks does affect how agents act in the market.

The focus of this paper is to synthesize previous literature on the topics of the information contained in chat boards, the affect that investor attention has on trading behavior, and the importance of social networks for communication and information dissemination. Specifically, I analyze the content of Stocktwits.com, an online forum for communication, to directly measure the cross sectional properties of attention and communication. Stocktwits.com provides specific information about the size of each contributor's social network, which allows me to directly measure the social network impact of any information that is disseminated through the site.

The ability to measure the social network impact of each contributor allows me to test the theoretical predictions of the DeMarzo et al.(2003) and Colla and Mele (2010) models. Specifically, DeMarzo et al. (2003) predict that if agents are boundedly rational then they fail to account for repetition of information that they encounter, and thus, their opinions will be subject to a persuasion bias. The presence of persuasion bias implies that individuals within a network can become socially influential, meaning that an agent's influence on the group depends on the size of the agent's networks, as well as the accuracy of his or her information. This theory gives rise to my first testable hypotheses:

**Hypothesis 1a: Firm specific communication through a network by an influential contributor should impact a firm’s contemporaneous and future abnormal returns.**

**Hypothesis 1b: The magnitude of an agent’s impact within a network is a function of the size of his network as well as the accuracy of his information.**

In this empirical setting, Hypothesis 1a implies that the frequency of network communication will impact current and future firm returns. Hypothesis 1b implies that the size of an agent’s Twitter network, as well as his credibility, will impact current and future abnormal returns.

Furthermore, Colla and Mele (2010) predict that social information linkages cause closely linked traders to trade in the same direction as their socially close peers, and that socially distant peers will trade in the opposite direction. Therefore, “neighbor” traders that are tightly connected in a network will have positively correlated trades, and “distant” traders that are not closely connected to the network will execute trades that are negatively correlated to the trades of the social network. Thus, if we assume that the market as a whole is the set of “distant” traders, future abnormal returns should be negatively correlated to day zero abnormal returns associated with the initial information flow of the set of socially connected traders. The negative relationship between current and future trading implies the second testable hypothesis.

**Hypothesis 2a: The abnormal returns on days with Twitter communication should be negatively correlated with future abnormal returns.**

Next, combining the implications of both Colla and Mele (2010) and DeMarzo et al. (2003) we can make the stronger prediction that an influential agent should be able to influence price movement in the same direction as his personal sentiment. Thus, when an agent communicates through his network with positive sentiment we should see positive

abnormal contemporaneous price movements. In this scenario, Colla and Mele predict that these positive contemporaneous price movements will be followed by negative abnormal returns generated by “distant” market traders.

**Hypothesis 2b: Contemporaneous abnormal returns should be of the same sign as the sentiment of network communication. Future abnormal returns will be negatively correlated to the initial network sentiment.**

In this empirical setting, Hypothesis 2b implies that a positive shock to network sentiment should be associated with positive contemporaneous abnormal returns and negative future abnormal returns.

Next, Colla and Mele (2010) predict that sentiment within a network is only weakly persistent and that the correlation among traders’ behaviors should decrease rather than increase over subsequent trading periods. This prediction is in opposition to theories of herding behavior that predict that correlation of opinions should increase over time essentially gaining momentum from one period to the next. This theory gives rise to Hypothesis 3.

**Hypothesis 3: Network sentiment should be positively autocorrelated in the short run but the impact of past sentiment should become insignificant over longer lags.**

Hypothesis 3 implies that in a time series setting social network sentiment should be positively autocorrelated in the short run. However, the autocorrelation coefficients should decrease rather than increase over time. Finally, long run autocorrelation terms should be insignificant in an autoregression of network sentiment.

The network communication variables developed in this paper are also useful for testing the predictions of theories of investor attention. The first theory that I consider is one of

investor recognition. Merton (1987) develops a rational agent model where investors have incomplete information about the entire universe of stocks available, and they are only informed about a subset of assets. Thus, under this framework, investors are more likely to buy assets of which they have knowledge. Furthermore, stocks with low recognition must offer a return premium to compensate their investors for holding an asset that is only recognized by a subset of investors. This phenomenon is known as the investor recognition hypothesis. The testable prediction of this theory is that low recognition stocks should exhibit positive returns even after controlling for other well known risk factors. The Merton (1987) theory leads to hypothesis 4:

**Hypothesis 4a: Stocks that are not talked about within the social network will exhibit positive future abnormal returns.**

Barber and Odean (2008) develop a model of attention driven buying which predicts that individual investors are net buyers of attention grabbing stocks. They measure attention through proxies such as news, abnormal returns or abnormal volume. Furthermore, constraints to short sales force individual investors to only sell stocks that they already own, meaning that they do not look outside their initial set of stocks for short sales. Thus, attention determines the subset of stocks upon which individual investors will make their trading decisions. This model of attention driven buying predicts that stocks with high levels of attention will experience short term buying pressure that will push their values away from fundamentals. Thus, in subsequent trading periods we should see that high attention stocks exhibit negative returns even after adjusting for well known risk factors. The Barber and Odean (2008) theory leads to the following testable hypothesis:

**Hypothesis 4b: Stocks that are talked about within the social network will exhibit negative future abnormal returns.**

I test Hypothesis 4a and 4b by analyzing the returns to portfolios formed based on the level of attention a firms receives within the social network.

### 3 Data Collection and Variable Construction

In this section I describe the data collection procedure, as well as how I store and manipulate the data once it is read from the internet. First, Stocktwits provides a list of contributors that are marked as "suggested" posters. Although the website does not use a quantitative rule for selecting the suggested contributor, in general, the suggested contributors have a large following, have been posting for a long time, and post meaningful or interesting comments. I select the sample of Twitter users from the set of recommended users as of October 1, 2008. I require that the Twitter users have at least 100 followers and 300 posts at the beginning of data collection process. This set of criteria produces a sample of 74 Twitter contributors. Then, all of their posts from October 1, 2008 to December 31, 2009 are downloaded using web scraping software that I specially designed for this particular website.

Next, the software compiles and stores each message in a text file and assigns a unique name based on its author, date and time. This process yields 143,640 individual posts that reference 1,416 different companies. I also add all of the S&P 500 members from 2009 to this dataset. This increases the sample to 1758 companies because 289 tickers which are in the S&P 500 were not tweeted about during the year. Next, I eliminate all ETFs, companies that are smaller than 100 million in market capitalization, and companies with trading price lower than two dollars. We are left with 1416 companies in the sample. The fact that such a large number of companies were mentioned in the sample illuminates the fact that micro bloggers are interested in a wide variety of companies.

The average size of a firm in the sample is 7.9 billion dollars with a book to market ratio of .4 and an average analyst coverage of 9. Thus the firms in our sample tend to be larger

growth firms that have high levels of analyst coverage. Finally, I collect specific information about the social network of each of the 74 individual contributors in my sample. Specifically, at the end of the sample period, I collect data from each contributor's Twitter account about the number of followers that they have, the number of people that they follow, and the number of posts that they have made. Table 1.1 reports firm level summary statistics as well as descriptive statistics of all communication variables.

To better understand the nature of a "tweet", it is helpful to look at some examples from Stocktwits. Figure 2.1 is the output of the Stocktwits website. It is a real time aggregated list of messages from all members of the Stocktwits network. Each message has the author, date, and time listed at the very top of the post, as well as a picture that the author provides for his online profile. The main body of each post is a short statement that can be no longer than 140 characters in length. These posts usually contain a declarative statement about a company or the economy. Finally, at the bottom of each post we can see a hyperlink for each of the tickers that are mentioned in the post. If the user clicks on a linked ticker, a daily price chart for that company pops up on the screen. As we can see from the figure, the posts on the website are different from the posts that can be found on Yahoo! chat boards because they are not constrained to one company or one line of thought. A viewer of the site sees a continual stream of financial topics that are the most interesting to the financial micro blogging community.

When a blogger wants to post about a particular company on the stocktwits website, they tag the company's ticker symbol with a "\$". For example, if I wanted to talk about Microsoft I would say, "\$MSFT, you should buy!".

[Figure 2.1 here]

Parsing each post for company ticker symbols is highly accurate because each ticker is pre tagged with the "\$". The classification of posts by company ticker symbol is done using a regular expression algorithm that searches the text for a "\$" followed by 1-4 letters that

constitute the ticker symbol. In the case where a blogger talks about multiple companies in one post, like in our example above, all references are counted as a unique comment. The final output of the text parsing procedure is a list which associates each post to a unique ticker. Next, all the references to a particular ticker are summed up over a one day period, giving us the total number of times that each ticker was mentioned during a given day. Thus, the output of this process is that each day every ticker is associated with an integer number from  $0 - N$  which corresponds to the number of posts that reference that particular company. Because posts on the weekend or on non-trading days cannot be contemporaneously compared to market data, only posts during the trading week are included in my analysis .

Next, I collect the network specific data that Twitter provides about each of the individual contributors in my sample. The author specific information that Twitter provides is the number of followers that an author has, the number of people that he or she follows, and the total number of tweets that the author has published. I then use this network information to create a variable that measures the specific influence that each contributor has when he or she posts to the site (details of this variable are provided in section 3.2). Finally, I collect daily stock return data from CRSP. Other firm level descriptive data comes from the Compustat annually updated database.

### **3.1 Sentiment Classification**

In this section, I describe the methodology used to classify the sentiment of Twitter posts. Computational linguistics (CL) is widely used to classify documents into categories such as genre, topic or author. Sentiment classification is one of the hardest classification tasks for CL to handle because levels of sentiments can not always be determined through key words alone. For example, in the phrase, “ you would be nuts to buy \$XYZ,” there is no single

word that is obviously negative, but the combined tone of the statement is unquestionably negative.

Nave Bayesian classification, which is the most popular technique for document classification, often has trouble with sentiment classification because it assumes the conditional independence of words in a document. As we saw from our example above, the sentiment of a sentence can not easily be determined if we look at each word independently. Maximum Entropy (ME) classification does not make any assumptions about the conditional independence of words, which is why ME tends to perform better when conditional independence does not hold. Pang, Lee and Vaithyanathan (2002) show that ME classification is the most robust for sentiment classification. Thus, I use ME classification throughout my study.

The general idea that drives ME classification is that when nothing is known about a distribution, the distribution should be as uniform as possible, i.e., have maximum entropy. In the case of sentiment classification, the ME classifier estimates the conditional distribution of the sentiment classes given a particular document, where all documents are represented by the frequency of words contained in them. Training data, which is pre-categorized data, is used to estimate the expected values of word counts on a class by class basis. Finally, the classifier will determine the probability that a document falls into a particular class based on its word content and the expected values of word counts conditional on categories.

Consider the example of trying to classify a document as positive, negative or neutral, where we are only told that 50% of documents that contain the word “buy” are considered positive. Intuition tells us that if the document has the word “buy” in it then there is a 50% chance that it is a positive post and a 25% chance of being either negative or neutral. If our document did not have the word “buy” in it then we would just assume an equal distribution of a 33% chance that the document falls into each category. Thus, when we know nothing about our document, we make the distribution uniform. This is the essence of ME classification. In practice, this process is constrained by many features, and the

calculations for conditional probabilities become complex, but the logic is still the same as our simple example.

To formally describe the ME procedure, I define the following set of terms. Let  $f_1, .f_m$  be a set of predefined features that can appear in a post. From our previous example, the word “buy” would be a feature, and the bi-gram ”short sell” could be a feature. Let  $n_i(d)$  be the number of times that the feature  $f_i$  occurs in a document  $d$ . Thus, each document is represented by a document vector that takes the form:  $\vec{d} = (n_1(d), n_2(d), , n_m(d), )$ . Lastly, let  $c$  be a document category of either positive, negative or neutral. Given this set of variables, the estimate of  $P(c|d)$  is as follows:

$$P_{ME}(c|d) = \frac{1}{Z(d)} (\sum_i \lambda_{i,c} F_{i,c}(d, c))$$

Where  $Z(d)$  is a normalization function, and  $F_{i,c}$  is a feature class function for the feature  $i$  in the class  $c$ .

$$F_{I,c}(d, c) = 1, n_i(d) > 0, \text{ and } c_i = c_0, 0 \text{ otherwise}$$

For example, this feature class function could only return a value of one if the post contained the bigram “short sell” and the post is hypothesized to be of negative sentiment. Finally,  $\lambda_{i,c}$  is a weighting parameter that determines the relative value of each of the features  $f_i$  contained in a document. If the value of  $\lambda_{i,c}$  is very large then the feature  $f_i$  is considered to be very strong for the document class  $c$ . Finally, the conditional probabilities  $P_{ME}(c|d)$  are estimated by maximizing the entropy across the three different groups while satisfying the constraint that the expected values of the feature class functions  $F_{i,c}$  are equal to their training data expected values.

To implement this methodology to my Stocktwits dataset, I start by hand classifying an out of sample set of 1,000 posts. This set of 1,000 hand classified posts is defined as my training set. From this training set, the expected values of the  $F_{i,c}$  are calculated. Then these relationships that are documented in the training data are used to calculate  $P_{ME}(c|d)$ , in my case, the probability that a message is positive, negative or neutral. The ME classifier is run on each post, and it produces a conditional probability for each class (negative, neutral, positive). Each post is assigned the sentiment of the class with the highest conditional probability. I assign the values (-1,0,1) to the sentiments (negative, neutral, positive). I test the accuracy of this procedure by running the ME classifier on a set of 100 posts that are hand classified. The ME classifier worked well in this out of sample test, and it was able to correctly classify 67% of all posts in the test sample. This accuracy rate is similar to the accuracy level that is achieved in other sentiment classification studies, such as Lee and Vaithyanathan (2002).

### 3.2 Variable Construction

The following explanatory variables are used throughout the rest of the paper. The first and most basic variable is *count*, which is defined as the daily frequency of posts for a given company. Next, I create a variable that measures the social network impact of each post as a function of the number of followers an author has, how many people the author follows, and the total number of tweets they have written. I call the social network impact variable *Impact*, and it is defined as follows:

$$Impact = \left( \frac{1 + Followers}{1 + Following} \right) * \ln( Number\ of\ Posts)$$

*Impact* incorporates two important features of the DeMarzo et al. (2003) theory of social network communication. First, the ratio of followers to following not only measures the size of an author’s network, but it also measures his ability to attract people outside of his network. Second, the number of posts an author has is a proxy for credibility. Each additional post that an author publishes adds to his historical track record, thus allowing a potential follower to analyze the previous view points and opinions that the author has posted to determine their accuracy. Finally, I take the natural log of the number of posts to minimize the impact of outliers, and then multiple the two elements together. Thus, *Impact* measures the magnitude of the social network impact that an author has when he comments about a stock.

Next, I create a set of variables that utilize the sentiment measures that are generated by the CL process. Once the CL algorithm assigns a sentiment to each post of either positive, neutral or negative (+1, 0, -1), I take the daily average of the sentiment of all posts for a given company to create the variable *senti*. By definition, this variable is continuous and constrained between -1 and 1. The construction of *senti* lends itself to very intuitive interpretation because large positive values denote strong positive sentiment, and large negative values denote strong negative sentiment. In order to determine the differential affect of positive or negative network communication, I interact *Impact* with *senti* to create the variable *Senti\_Impact*. This variable measures both the reach that information has within a network as well as its sentiment. Large positive values of *Senti\_Impact* denote a broad dissemination of positive sentiment, and large negative values denote a broad dissemination of negative sentiment.

In order to include longer lags of the communication variables I adopt a standard naming convention for the cumulation of network communication over a twenty day horizon. I define *Count20*, *Impact20* and *Senti\_Impact20* as the summations of *Count*, *Impact* and

*Senti\_Impact* over the  $(-1,-20)$ <sup>1</sup> horizon respectively. The formal representation of these variables is as follows:

$$\begin{aligned} Count20jt &= \sum_{t=-1}^{t=-20} Count_{jt} \\ Impact20jt &= \sum_{t=-1}^{t=-20} Impact_{jt} \\ Senti_Impact20jt &= \sum_{t=-1}^{t=-20} Senti_Impact_{jt} \end{aligned}$$

I define  $L$  as a daily lag operator. Thus,  $L, L2, L3....Ln$  create the  $1, ..., n$  daily lags of any variable with which they are interacted. Finally, in order to aid in interpretation, all explanatory variables are converted to  $z - scores$ .

The variable  $R_{mt}$  is the value-weight return on all NYSE, AMEX, and NASDAQ stocks (from CRSP) minus the one-month Treasury bill rate (from Ibbotson Associates).  $Ret$  is defined as the close to close total return to an asset  $i$ .  $ar$  is the abnormal return from a Fama-French three factor model with rolling beta estimation. For each day in my dataset, I estimate the model  $ret = \alpha + \beta * mrkt - rf + \gamma_1 HML + \gamma_2 SMB + \gamma_3 UMD$  using 150 days of lagged data to estimate the parameters  $\alpha$  and  $\beta$ ,  $\gamma_1$ ,  $\gamma_2$  and  $\gamma_3$ . The variable  $ar$  is the daily residual from this regression and is estimated as follows:

$$AR_{jt} = R_{jt} - (\hat{\alpha}_{jt} + \hat{\beta}_{jt}R_{mt} + \gamma_{1jt}HML_t + \gamma_{2jt}SMB_t + \gamma_{3jt}UMD_t)$$

Where  $R_{mt}$ ,  $HML_t$  and  $SMB_t$ ,  $UMD_t$  are the returns to the market, book value, size and momentum factors respectively on day  $t$ , and  $R_{jt}$  is the return to the  $j^{th}$  firm on day  $t$ .

Finally, the abnormal or raw returns are cumulated over the (0,1), (0,4), (0,9), (0,19) event windows to create the variables  $CAR_2, CAR_5, CAR_{10}$ , and  $CAR_{20}$ . The specific

---

<sup>1</sup>The results of the paper are robust to alternative lag lengths such as five or ten days.

form of these variables is as follows:

$$\begin{aligned}
CAR2_{jt} &= \sum_{t=0}^{t=1} AR_{jt} \\
CAR5_{jt} &= \sum_{t=0}^{t=4} AR_{jt} \\
CAR10_{jt} &= \sum_{t=0}^{t=9} AR_{jt} \\
CAR20_{jt} &= \sum_{t=0}^{t=19} AR_{jt}
\end{aligned}$$

The following is an outline of the Stocktwits data collection methodology:

1. Parse all tweets and associate each post with a particular ticker (or tickers)
2. Classify the sentiment of the post (Positive = 1, Negative = -1, neutral = 0)
3. Sum all stock references up over daily horizon.
4. The output is a list of tickers for each day with the frequency and sentiment of posts
5. From this raw data, I create social impact variables utilizing the information obtained about each individual poster.

## 4 Relating the Level of Network Communication to Stock Returns

In this section I explicitly test Hypotheses 1a, 1b and 2a. The assertions of these hypotheses are directly related to the theories of DeMarzo et al. (2003) and Colla and Mele (2010).

DeMarzo et al. (2003) assert that an agent’s influence is a function of his creditability and the size of his network. Colla and Mele (2010) assert that socially close traders will execute positively correlated trades and socially distant will have negatively correlated trading behavior.

The variables *count*, *Impact* are both measures of network communication, thus, systematic relations between the network variables and the market will support Hypothesis 1a that network communication should impact contemporaneous and future returns. Hypothesis 1b that an agent’s influence is a function of his credibility and network size is tested with the variable *Impact* which proxies for these characteristics. Finally, if we assume the market as a whole is a “distant” trader and the agents in the Stocktwits network are socially close traders, the correlated trading theory of Colla and Mele (2010) predicts that in the short run we should see positively correlated trading from closely related individuals in the social network. Subsequently, we should see trading by the market as a whole that is negatively correlated to the initial trades made from the social network. This theory gives rise to Hypothesis 2a that contemporaneous and future abnormal returns should be negatively correlated after network communication.

Panel A of Table 1.2 presents the results from a set of fixed effects panel data regressions that utilize abnormal returns as the dependant variable and lags of the variable *count* as the explanatory variables. All independent variables have been standardized to a z-score. Ten lags of *ar* are added as controls, and Driscoll and Kraay (1998) standard errors are used to control for cross-sectional and time series correlation. The variable *count* is a measure of the frequency of posting. *Count* implicitly assumes that all posts are of equal importance, and each contributor is assumed to have equal influence within the network. This simplifying assumption allows for a base line test to measure the robustness of future network variables. It is also important to note that the sample of contributor is only taken for the most “influential” and high profile contributors in the network. Since the sample has already been

conditioned on the network influence of the authors, the frequency of posts alone should be a good proxy for the total amount of information flow in the social network.

Column 1 of Panel A presents the contemporaneous regression of one day abnormal returns on contemporaneous and lagged values of *count*. The estimated coefficient on *count* shows that there is a strong positive contemporaneous relationship between information flow and abnormal returns. Specially, a one standard deviation increase in *count* implies a 26 bps increase in contemporaneous daily abnormal returns. This result provides initial support for Hypothesis 1a.

Although the contemporaneous relationship between returns and posting frequency is interesting, the more important question is whether posting frequency is a leading indicator of returns. Columns 2-5 in Panel A present time sequenced regressions that analyze the relationship between lagged information and future abnormal returns. Contrary to the time zero values of *count*, all of the lagged values of *count* have a significant negative relationship with future abnormal returns. Column 2 shows that a one standard deviation increase in posting activity over the (-1,-20) event window implies -8 bps over the following 2 days. This relationship is most pronounced over the longer horizon event windows of 5-20 days after posting occurs. The estimated coefficient for *Count20* in Column 3 implies that a one standard deviation increase in the posting activity over the (-1,-20) window will decrease abnormal returns by 21 bps over the next five days of trading, and the magnitude of this estimated coefficient increases to -39 bps and -74 bps over the (0,9) and (0,19) event windows respectively. These results support Hypothesis 1a that network communication should impact current and future returns. Furthermore, the positive 23 bps of abnormal returns on day zero followed by -101 bps of abnormal returns over the (0,19) event window supports Hypothesis 2a that current and future returns should be negatively correlated.

Panel B of Table 1.2 presents contemporaneous and lagged regressions of abnormal returns on contemporaneous and lagged values of the social network impact variable *Impact*.

Ten lags of  $ar$  are added as controls. *Impact* includes a measure of an agent’s credibility as well as the size of his network following. Thus, these regressions of *Impact* are direct tests of the Hypothesis 1b, that an agent’s influence within a social network is a function of his credibility and the size of his network.

Column 1 in Panel B presents the contemporaneous regression of *Impact* on a one day abnormal return. The same positive relationship that we saw in the regression of *count* exists in this model specification. Specifically, a one standard deviation increase in *Impact* implies a 15 bps increase in day zero abnormal returns. Columns 2 - 5 present the time sequenced regressions of *Impact* on abnormal returns from one to twenty days in the future. The estimated coefficient on *Impact*20 in column 2 shows that a one standard deviation increased network communication in the preceding twenty days will decrease abnormal returns by 8 bps over the (0,1) event window.

This result is robust over longer event windows as well. The estimated coefficients on *Impact*20 in Columns 3-5 imply that a one standard deviation increase in information flow over the (-1,-20) window will decrease abnormal returns by 19 bps over the following five trading days, 35 bps over ten trading days and 62 bps over twenty trading days. The significant affect that *Impact* has on contemporaneous and future abnormal returns provides strong support for Hypothesis 1b that an agent’s influence is a function of his credibility and network size. The negative relationship between day zero returns and future returns also provides support for Hypothesis 2a.

The contemporaneous and time sequenced results from this set of regressions provide support for the DeMarzo et al. (2003) and Colla and Mele (2010) theories of social influence. The contemporaneous results show that the level of communication is contemporaneously positively related to abnormal price movements, and that spikes in social persuasion can affect market prices. Furthermore, the negative relationship that is seen in the time sequenced regressions provides support for the Colla and Mele (2010) theory of correlated

trading. Specifically, I document a positive move in prices on day zero that is reversed over the subsequent one to four weeks of trading. This pattern of negative correlation is seen with both *count* and *Impact*.

## 5 Posting Frequency and Sentiment

In this section I test the robustness of the results found in the previous section by taking into account the sentiment as well as the frequency of network communication. The variable *Senti\_impact* includes a measure of an agent’s influence and credibility as well as a measure of the sentiment of the information that the agent disseminates through the social network. The relationship of sentiment to abnormal returns is a powerful test of the theories of DeMarzo et al. (2003) and Colla and Mele (2010). The variable *Senti\_impact* allows me to directly test the prediction of Hypothesis 2b that day zero abnormal returns should be positively correlated to network sentiment, and future abnormal returns should be negatively correlated to network sentiment.

To test Hypothesis 2b I use the variable *Senti\_impact* to exploit the information in both posting frequency and posting sentiment. The variable *Senti\_impact* is the interaction of the sentiment indicator, *senti*, with *Impact*. Therefore, the magnitude and the sign of *Senti\_impact* is important. Table 1.3 summarizes the contemporaneous and lagged relationship between abnormal returns and the level and sentiment of network communication. I include the variable *Senti\_Impact20*, which is the cumulation of *Senti\_Impact* over the (-1,-20) event window in the regressions, as well as ten days of lagged abnormal returns. *Senti\_impact* exhibits the same patterns as *Impact*. The estimated coefficient on the contemporaneous value of *Senti\_impact* in Column 1 of Table 1.3 implies a 9 bps increase in abnormal returns on day zero given a one standard deviation increase in the sentiment

of total information flow within the network. Thus, sentiment is positively correlated to contemporaneous returns as predicted by Hypothesis 2b.

Column 2 reveals that *Senti\_Impact*20 is negatively correlated to future abnormal returns. Specifically, a one standard deviation increase in *Senti\_Impact*20 implies -7 bps of abnormal returns over the (0,1) event window. In columns 3-5 we see that this negative relationship also continues over the (0,4), (0,9) and (0,19) event windows with abnormal returns of -14 bps, -27 bps and -46 bps respectively. The positive relationship between contemporaneous abnormal returns, and the negative relationship between lagged sentiment and future returns provide clear support for Hypothesis 2b. This shows that sentiment should be positively correlated to contemporaneous abnormal returns, and negatively correlated to future abnormal returns. This pattern of abnormal returns also provides further support for Hypotheses 1a, 1b and 2a that a network agent should have an impact on his social connections that is increasing function of his importance within the network.

The addition of sentiment to the analysis confirms the support for the theories of DeMarzo et al. (2003) and Colla and Mele (2010) that was found in the previous section. First, the variable *senti\_impact*, which interacts a sentiment indicator with *Impact*, is significantly related to market activity. This supports the prediction that influential agents should be able to impact the opinions of their socially close peers. Secondly, negative correlation between contemporaneous and lagged returns supports the first prediction of Colla and Mele (2010) that socially close traders should have positively correlated trading behavior and socially distant traders should have negatively correlated trading behavior. Thus, when communication occurs on day zero, socially close traders trade in the same direction, and this is associated with positive abnormal returns. Then the market as a whole, acting as the socially distant trading group, execute trades that are negatively correlated to the day zero trades. This causes negative abnormal returns in the subsequent weeks.

## 5.1 Cross Sectional Analysis of Sentiment

This section provides additional support for the main results of this paper by testing the differential impact of *senti\_impact* across groups of stocks with high and low levels of information asymmetry. Since network communication facilitates the flow of information across sets of individuals, we should see a larger impact from network communication in stock that have high levels of information asymmetry. Simply, if the information in the market about a firm is not reliable or is hard to extract, then new flows of information should have a larger impact on abnormal stock returns. We test this idea by sorting firms based on size and S&P500 membership, two common proxies for information asymmetry. Size is defined as fiscal year-end market capitalization taken from Compustat. For each day in the sample, firms are sorted into size terciles. Then separate panel regressions are run on the top and bottom terciles. S&P500 membership is determined on a monthly basis throughout the sample, and separate regressions are performed on non-S&P500 firms and S&P500 firms. Finally, I test the difference between the estimated coefficients for *Senti\_Impact*<sub>20</sub> across high and low information asymmetry groups.

Table 1.4 presents the results of sorting firms based on size. The literature widely agrees that small stocks have a higher degree of information asymmetry than large stocks, see Hong et al.. (2000). Thus, the predictions of Hypothesis 2a and 2b should be stronger for small stocks than for large stocks. Panel A of Table 1.4 shows the contemporaneous and time sequenced regression for small stock in our sample, and Panel B presents the results for large stocks. Column 1 of Panels A and B present the contemporaneous relationship of *Senti\_Impact* and abnormal returns. *Senti\_Impact* is positively related contemporaneous abnormal returns for both small and large firms. Specifically, a one standard deviation increase in *Senti\_Impact* implies 47 bps and 3 bps of abnormal returns for small and large firms respectively. The field (1)-(3) at the bottom of Panel B tests the difference between these estimated coefficients. The difference of 43 bps between small and large stocks shows

contemporaneous network communication has a much larger positive impact on higher information asymmetry firms. This result supports the theory that information flow should be more valuable for firms that investors are relatively less informed about. Columns 2-4 of Panels A and B confirm that abnormal returns for two to twenty days in the future are negatively related to the sentiment of network communication. This negative relationship is also stronger for small firms. Specifically, small firms exhibit -48 bps, -2.24% and -3.81% of abnormal returns over the (0,1), (0,9) and (0,19) windows, whereas large stocks exhibit only -2 bps, -10 bps and -19 bps over the same period. The bottom of Panel B tests the difference between the estimated coefficients of the groups of large and small firms. We can see in row (2)-(4) that small firms have a larger reaction to network information flow. Over two, ten and twenty trading days, small firms return 45 bps, 2.13 % and 3.61 % less than large firms. This set of contemporaneous and time sequenced results support the predictions of Hypotheses 2a and 2b by showing that the network sentiment has a larger impact on firms in poor information environments.

Table 1.5 presents the cross sectional results from sorting firms based on S&P500 membership. Following Seasholes and Zhu (2010), I use S&P500 membership as a proxy for information asymmetry because the index consists of large firms which constitute nearly 80% of the market capitalization of US equities. Members of the S&P500 are also generally followed by a large number of analysts as well as having high levels of institutional ownership. On the contrary, Non-S&P500 firms are generally smaller, less well covered by analysts and have lower institutional ownership. Therefore, Non-S&P500 should have higher levels of information asymmetry. Network communication should affect non-S&P500 firms more than S&P500 firms as a result of the differential in information asymmetry across the two groups.

Support for this prediction is clearly seen in the difference between the non-S&P500 firms in Panel A and the S&P500 firms in Panel B. *Senti\_Impact* has a large affect of 15

bps on the contemporaneous abnormal returns of non-S&P500 firms, but a small affect of 4 bps on S&P500 firms. The difference in estimated coefficients across the groups is 10 bps, which is economically and statistically significant. *Senti\_Impact* also has a larger impact on future abnormal returns for the non-S&P500 group. The estimated coefficients for *CAR2*, *CAR10* and *CAR20* are -15 bps, -65 bps and -1.03 % for the non-S&P500 but the estimated coefficients S&P500 firms are of a much smaller magnitude of -3 bps, -11 bps and -17 bps. The difference between the estimated coefficients of the two groups in row (2)-(4) at the bottom of Panel B shows that non-S&P500 firms have negative abnormal returns in excess of S&P500 firms by -14 bps, -55 bps and -83 bps over two, ten and twenty trading days respectively. This large and significant difference provides additional support for the fact that network communication is more influential for firms with high informational asymmetry

## 6 Predicting Information Flow through Networks

This section tests Hypothesis 3 and the prediction of Colla and Mele (2010) that the correlation among beliefs inside a social network should decay over time rather than get stronger. This prediction is contrary to theories of herding behavior which predict that the correlations among beliefs will become stronger over time. Empirically, this theory implies that the flow of information and opinion through a social network should follow a positive autoregressive process. Furthermore, the autoregressive properties should decay to zero over a longer horizon.

I test the prediction that close peers should have positively correlated opinions in a set of autoregressive panel data models. Table 1.6 presents the results of the fixed effects panel data regressions that use lagged information in order to predict the variable *Senti\_Impact* on day zero. In order to measure the information flow over various horizons, I define a set of lags of *Senti\_Impact*. The variables *Week1Senti\_Impact*, *Week2Senti\_Impact*, *Week3Senti\_Impact*

and *Week4Senti\_Impact* are defined as the summation of *Senti\_Impact* over the (-1,-5),(-6,-10),(-11,-15) and (-15,-20) event windows respectively. All regressions are firm fixed effects with robust standard errors. Column 1 of Table 1.6 is an autoregressive model of *Senti\_Impact* which includes *Week1Senti\_Impact* and *Week2Senti\_Impact*.

The time series of autoregressive coefficients supports the predictions of Colla and Mele (2010) because we see strong positive correlation between communication in the short term. Specifically, in column 1 of Table 1.6, the estimated coefficients on *Week1Senti\_Impact* and *Week2Senti\_Impact* imply that a one standard deviation increase network sentiment over the (-1,-5) and (-6,-10) horizon will cause a .078 and .049 standard deviation increase in *Senti\_Impact* on day zero. These estimated coefficients are significant, but they are not economically very large. This shows that sentiment has weak positive autoregressive characteristics over a short horizon.

Column 2 of Table 1.6 tests longer time horizons by adding the variables *Week3Senti\_Impact* and *Week4Senti\_Impact* which lag the dependant variable over the (-11,-15) and (-15,-20) event windows respectively. Column 2 shows that the autocorrelation displayed over the (-1,-10) horizon becomes insignificant over the longer (-11,-15) and (-15,-20) windows. This shows that current sentiment has influence over future sentiment over about a ten day period. However, after ten days current sentiment cannot be used to predict future sentiment. This shows that correlation among the opinions of individuals in a network decays to zero beyond a short time horizon. Finally, Column 3 also includes the variables *CARLAG5* and *CARLAG10* to the regression model which are defined as the summation of abnormal returns over the (-1,-5),(-6,-10) horizon respectively. The estimated coefficient on *CARLAG5* implies that positive abnormal returns over the (-1,-5) horizon will have a small positive affect on future sentiment. However, the insignificant estimated coefficient on *CARLAG10* shows us that abnormal only impact future sentiment over a short horizon of five days.

This set of lagged information models provides support for Hypothesis 3, and the Colla and Mele (2010) theory on the dynamics of network communication. First, the autoregressive models provide support for the prediction that opinion should be positively correlated among socially close peers. Furthermore, the long-run dynamics of the autoregressive models provide support for the Colla and Mele (2010) model over competing models of herding because the positive correlation among opinions decays to zero over a ten day time period, rather than increasing in strength over a longer time horizon.

## 6.1 IBM Case Study

In this section, I present a simple case study of IBM to help the reader visualize the time series properties of network communication. In Figure 1.2, I plot the time series of *Senti\_Impact* for International Business Machines (IBM). The time series of data shows that there are many days where there is no information flow about IBM at all. However, the low level of information flow that is common in the data is periodically interrupted by large spikes in the level of information flow that usually persist for multiple days at a time. For example, the large clustering of spikes in the middle of the graph lasts for 15 days before it dies off. These spikes in information flow that periodically occur and then subside are anecdotal evidence that communication in social networks acts as a conduit for analysis and dissemination of pertinent company specific information. Furthermore, the graph has been annotated with the headlines that correspond to the larger spikes in information flow. The annotations show that increased information flow corresponds to significant events such as takeover talk, initiations of new contacts, dividend policy, and earnings guidance.<sup>2</sup>

[Figure 1.2 here]

---

<sup>2</sup>Further research has shown that network communication often disagrees with the sentiment of popular media. Giannini and Irvine (working) show that network agents often disagree with media sentiment around earnings announcements.

## 7 Network Communication: A Direct Measure of Attention

The communication variables that have been used throughout this paper are also useful in testing theories on investor attention. In this section, I consider Merton’s investor recognition hypothesis, as well as the behavioral model of Barber and Odean (2008). These two theories have similar motivations but make two separate predictions about the impact of attention. The Merton (1987) investor recognition hypothesis predicts that low attention stocks should exhibit a positive premium; this prediction is stated in Hypothesis 4a. The Barber and Odean (2008) theory predicts that high attention stocks should exhibit negative risk adjusted returns. This prediction is stated in Hypothesis 4b. I am able to test both of these theories simultaneously using a time series factor mimicking approach. The variables *Count*, *Impact* and *Senti\_Impact* are direct measures of how much attention is being paid to certain stocks within a network of traders. Therefore, rather than using proxies for attention, such as abnormal returns, volume and news, these variables allow for the direct measurement of the impact of attention on asset prices.

The methodology of this test is as follows. At the end of each day, I sort stocks based on the level of communication that the company received over the previous week. Stocks are sorted into three bins based on the values of the variables *count*, *Impact* and *Senti\_Impact*. I separate stocks that have not experienced any communication and place them in a bin named no-talk. The remaining stocks that have been talked about within the network are sorted into two groups, high-talk and medium-talk. Each of the groups is then combined into equally weighted portfolios to generate a daily portfolio return. This process is repeated over the sample to create a time series of portfolio returns for each of the three levels of communication. Finally, I create zero investment portfolios that go long the no-talk stocks and short the high-talk stocks. I regress the time series of returns generated from these

twelve portfolios on the Fama-French three factor model to determine if their returns can be fully explained by well known risk factors. If the level of communication that a stock has received is not priced by these factors then the portfolios should exhibit alphas that are significantly different from zero.

Table 1.7 presents the results of regressing the time series of the twelve portfolios created based on level of communication on the Fama French three factor model. The intercepts from the three factor regressions are reported along with robust standard errors. The portfolios presented in Table 1.7 are formed based on one week of lagged communication information and subsequently held for one day. The pattern that emerges is that companies that do not have any communication about them in the previous week produce positive risk adjusted returns over the subsequent day of 7 bps. This pattern is significant and persists for all of the measures of communication *count*, *Impact* and *Senti\_Impact*. This positive performance supports Hypothesis 4a and the investor recognition hypothesis.

Although the notalk portfolios exhibit abnormal performance, none of the high or medium attention portfolios exhibit any abnormal risk adjusted performance. Furthermore, none of the zero investment portfolios exhibit any abnormal performance. The insignificant alphas on the high talk portfolios do not provide support for Hypothesis 4b and the Barber and Odean theory. This set of results shows that only the subset of assets that receive no investor attention exhibit positive abnormal returns.

I further test the relationship between attention and returns with a set of fixed effects panel models which include a “No Attention” dummy variable. The dummy variable *None* is equal to one if there is no network communication about a firm in the previous twenty trading days. Merton’s model predicts that *None* should be positively related to future returns. Table 1.8 presents the results of this regression framework. Columns 2-5 show that no attention stock have positive risk adjusted returns over the next two-twenty trading days. For example, over the next five trading days a firm with no communication will have 31 bps of

positive abnormal returns. This result adds further confirmation to the investor recognition hypothesis. Furthermore, Barber and Odean (2008) predicts that increased attention should cause negative returns. This prediction is supported by the negative estimated coefficient on *Impact*<sub>20</sub>, which is a direct measure of attention. Columns, 2-5 show that an increase in attention predicts negative abnormal returns over the subsequent two-twenty trading days. For example, a one standard deviation increase in attention implies -18 bps of abnormal returns over the next five trading days.

The results of this section support the Merton (1987) theory of investor recognition and Hypothesis 4a. As predicted by the theory, there is a positive return premium for stocks that have a low level of investor attention. Similar empirical support for this theory has been found by Fang and Peress (2010) using news as a proxy for attention. I find support for the Barber and Odean (2008) theory that high attention stocks should exhibit negative risk adjusted returns in the in the panel regression framework but not in portfolio tests. These results lend weak support for the predictions of Barber and Odean (2008).

## 8 Conclusion

In this paper I measure the value of information flow through social networks by analyzing the content of posts to Stocktwits.com. Stocktwits provides a unique arena to test new theories on the impact of social networks on beliefs and trading behavior. I find support for the DeMarzo et al. (2003) theory that an agent's influence within a network is a function of his credibility and the size of his network. I construct a variable that is a function of an agent's network size, credibility and the sentiment of his opinions and show that an increase in the level of information dissemination within a network is positively related to contemporaneous returns.

I also find support for the Colla and Mele (2010) theory on correlated trading among groups. I show that socially “close” agents trade in the same direction when communication occurs, and in subsequent days socially “distant” agents execute trades that are negatively correlated to the trades on day zero. This empirical regularity is seen as positive abnormal returns on the day communication occurs of 26 bps followed by up to -74 bps of abnormal returns over the next twenty trading days. I also find that these results are stronger for groups of stocks that have higher levels of information asymmetry such as small firms and non-S&P500 firms. The impact of network communication on small firms is 61 bps greater than the impact on large firms in the ten trading days after communication.

Furthermore, I find support for the prediction that correlation of opinions within a group fades over time using a set of autoregressions on my measure of sentiment flow. These autoregressions show that a one standard deviation increase in sentiment over the previous five trading days will increase the next day’s sentiment by .084 standard deviations. Furthermore, at longer horizons correlation among opinions within a network decay to zero. In the final section of the paper, I use my measures of network communication to test the predictions of two theories of attention. I construct a set of portfolios that are based on the lagged level of attention that a stock has received. Portfolios formed of stocks with no network communication exhibit 7 bps of abnormal returns the day after formation. This result provides support for the Merton (1987) theory of investor recognition by showing that stocks that have no recognition within the network exhibit future positive risk adjusted returns.

The information that flows through modern electronic social networks provides us with a new source of real time information that has yet to be fully exploited. There are many aspects of social communication that have never been tested in the empirical finance literature, and it is a fruitful area to look for new avenues of study. Future research could focus on the relationship between the communication of various groups, such as professional investors and retail investors. Another area that would be interesting to investigate is the relationship

between distinct events and network communication. In conclusion, as the use of the internet as an investing tool grows, so will the power of and reliance on electronic social networks. Today, only a small subset of investors use electronic social networks, but what is a small minority today might become a majority in the future.

## 9 References

- Antweiler, Werner, and Murray Z. Frank, 2004, Is all that talk just noise? The information content of Internet stock message boards, *Journal of Finance* 59, 1259-1294
- Antweiler, Werner, and Murray Z. Frank, 2006, Do U.S. stock markets typically overreact to corporate news stories? Working paper, University of British Columbia
- Barber, B. M., Y. Lee, Y. Liu, and T. Odean, 2009, Just How Much Do Individual Investors Lose by Trading? *Review of Financial Studies* 22, 609-632
- Barber, B. M., Y. Lee, Y. Liu, T. Odean, 2005, Do Individual Day Traders Make Money? Evidence from Taiwan. <http://ssrn.com/abstract=529063>
- Barber, B. M., and T. Odean. 2000. Trading is hazardous to your wealth: The common Stock Investment Performance of Individual Investors. *Journal of Finance* 55:773806
- Barber, B. M., and T. Odean. All that Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors. *Review of Financial Studies* 21(2):785818
- Boehmer, Ekkehart , Jim Masumeci & Annette Poulsen, 1991, Event-study methodology under conditions of event-induced variance *Journal of Financial Economics*, 30(2), 253-272
- Bollen,Johan , Huina Maoa and Xiaojun Zeng, 2011, Twitter mood predicts the stock market, *Journal of Computational Science*, 2, 1-8
- Chan, Wesley, 2003. Stock price reaction to news and no-news: drift and reversal after headlines, *Journal of Financial Economics* 70, 223-260
- Cohen,Lauren, Andrea Frazzini, and Christopher Malloy, 2008. The Small World of Investing: Board Connections and Mutual Fund Returns *Journal of Political Economy* 116:5, 951-979
- Paolo Colla , and Antonio Mele, 2010. Information Linkages and Correlated Trading. *Review of Financial Studies* 23: 203-246
- Da, Zhi, Joseph Engelberg and Pengjie Gao, In Search of Attention (June 4, 2009). Available at SSRN: <http://ssrn.com/abstract=1364209>
- Das, Sanjiv, and Mike Chen, 2007, Yahoo! for Amazon: Sentiment extraction from small talk on the web, *Management Science* 53, 1375-1388
- DeMarzo, Peter, Dimitri Vayanos, Jeffrey Zwiebel,2003, Persuasion Bias, Social Influence, and Unidimensional Opinions, *Quarterly Journal of Economics* 118:3, 909-968

Driscoll, John and Aart Kraay, 1998, Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data, *Review of Economics and Statistics*, 80:4, 549-560

Fama E. Efficient Capital Markets: II, *Journal of Finance* 2010

Fang, Lily H. and Peress, Joel, Media Coverage and the Cross-Section of Stock Returns, 2009. *Journal of Finance* 64 2023-2052

Frazzini, Andrea, Christopher J. Malloy and Lauren Cohen, Sell Side School Ties, 2010, *Journal of Finance* 65, 1409-1437

French, Kenneth R., G. Willaim Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3-29

Hong, Harrison, Terence Lim and Jeremy C. Stein., 2000, Bad News Travels Slowly: Size, Analyst Coverage, And The Profitability Of Momentum Strategies, *Journal of Finance* 55, 265-295

Lavrenko, V., M. Schmill, D. Lawrie, P. Ogilvie, D. Jensen, J. Allen, 2000, Mining of concurrent text and time series. *Proc. Knowledge Discovery Data Mining, 2000 Conf. Text Mining Workshop*, 37-44

Pang, Bo and Lee, Lillian and Vaithyanathan, Shivakumar, Thumbs up?: sentiment classification using machine learning techniques EMNLP '02: *Proceedings of the ACL-02 conference on Empirical methods in natural language processing*

Li, Feng, 2006, Do stock market investors understand the risk sentiment of corporate annual reports? Working paper, University of Michigan

Merton, R., 1987, A Simple Model of Capital Market Equilibrium with Incomplete Information, *Journal of Finance* 42(3): 483-510

Mitchell, M. and Mulherin J., 1994, The Impact of Public Information on the Stock Market, *Journal of Finance* 49(3): 923-950

Nigam, K. ,J. Lafferty ,A. McCallum ,1999, Using Maximum Entropy for Text Classification IJCAI-99 Workshop on Machine Learning for Information Filtering, pages 61-67

Pastor, L. and R.F. Stambaugh, 2002, Liquidity Risk and Expected Stock Returns *Journal of Political Economy*

Seasholes, Mark S., and Ning Zhu, 2010, Individual investors and local bias, *Journal of Finance* 65, 1987-2010

Stein, Jeremy C. 2008. Conversations among Competitors. *American Economic Review*, 98(5): 2150-62

Tetlock, Paul C., 2007, Giving content to investor sentiment: The role of media in the stock market, *Journal of Finance* 62, 1139-1168

Tetlock, Paul C., Maytal Saar-Tsechansky and Sofus Macskassy, 2008, More than Words: Quantifying Language to Measure Firms' Fundamentals, *The Journal of Finance* 63, 1437-1467

Tumarkin, R., R. Whitelaw. 2001. News or noise? Internet postings and stock prices. *Financial Analysts Journal* 57(3) 41-51

Wysocki, P., 1999, Cheap talk on the web: The determinants of postings on stock message boards, University of Michigan Business School Working Paper

Zhang, Ying, Swanson, Peggy, 2010. Are day traders bias free?-evidence from internet stock message boards. *Journal of Economics and Finance* 34, 96-112

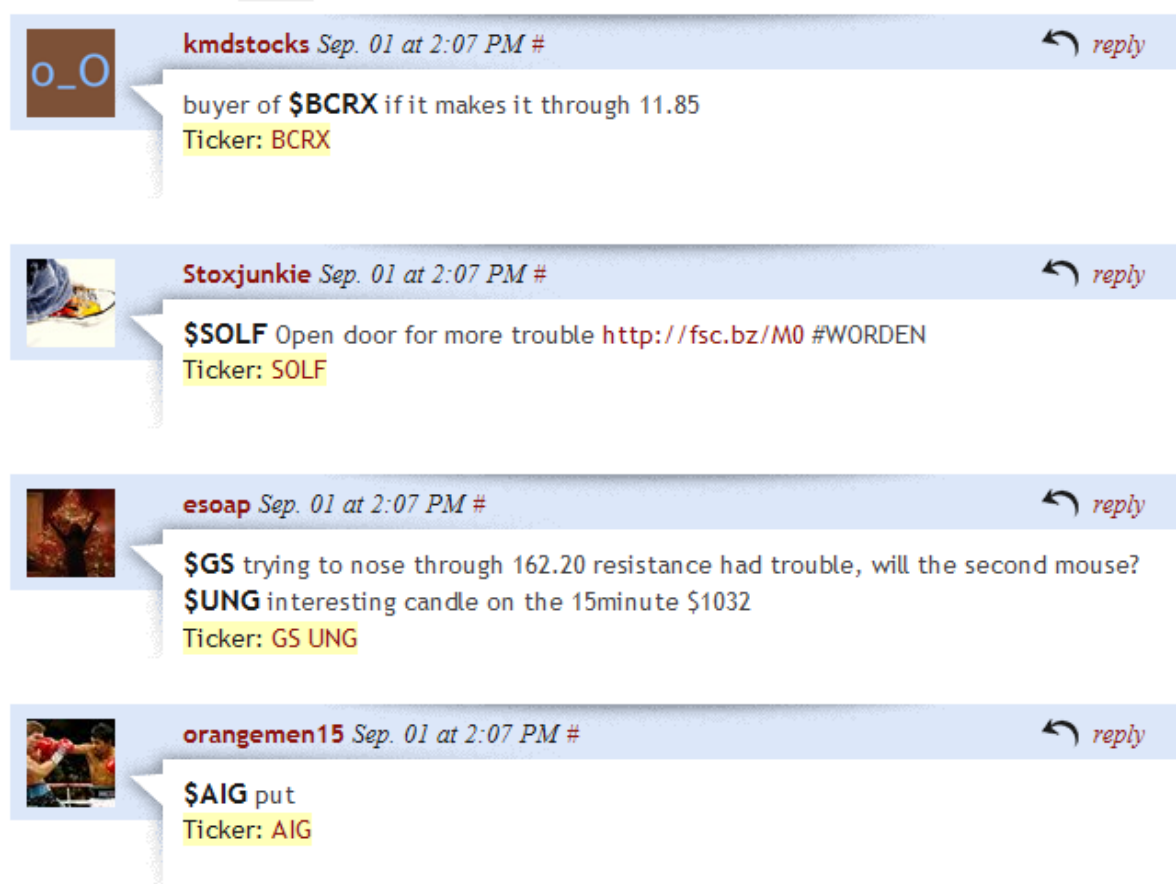


Figure 1.1: www.Stocktwits.com




Figure 1.2: Information Flow about IBM

**Table 1.1: Summary Statistics for Sample Firms and Communication Variables**

This table presents summary statistics of the 1,384 sample firms from the sample period of October 1st, 2008 to December 31st, 2009. Market capitalization is measured at the end of 2007. Book-to-market ratio is book equity of fiscal year ending in calendar year 2007 divided by market capitalization at the end of 2007. Analyst coverage is the number of analysts covering the firm as of September 2007. Market R2 is the R2 of time-series regressions of a firms excess returns on the market excess returns and *SMB*, *HML*, and *UMD* factors during our sample period. Finally, I include summary statistics of the non-standardized values of the communication variables (*Count*, *Impact*, *Senti\_Impact*) used throughout the paper.

| Variable                | Mean   | Std    | P10th | P25th | P50th | P75th | P90th  |
|-------------------------|--------|--------|-------|-------|-------|-------|--------|
| Market Value (Millions) | 7,940  | 27,630 | 245   | 465   | 1,288 | 4,499 | 16,095 |
| Book/Market             | 0.40   | 0.39   | 0.10  | 0.19  | 0.33  | 0.55  | 0.79   |
| # Analysts              | 9.11   | 6.14   | 3.00  | 5.00  | 8.00  | 12.00 | 18.00  |
| $R^2$                   | 0.46   | 0.14   | 0.27  | 0.36  | 0.46  | 0.55  | 0.65   |
| Ret (%)                 | 0.02   | 0.66   | -0.11 | 0.01  | 0.09  | 0.19  | 0.30   |
| <i>Count</i>            | 35     | 131    | 1     | 3     | 8     | 26    | 70     |
| <i>Impact</i>           | 13,078 | 78,569 | 96    | 286   | 1,147 | 6,131 | 20,684 |
| <i>Senti_impact</i>     | 1,634  | 7,820  | -37   | 0     | 206   | 992   | 4,326  |

**Table 1.2: Contemporaneous and Lagged Communication**

In this table, Panel A presents the regressions of contemporaneous and future stock returns on the frequency of network communication. Panel B presents regressions of abnormal returns on a measure of the user specific impact of network communication. *ar* is defined as the abnormal returns from a Fama-French 4 Factor model with 150 day rolling estimation. *CAR2*, *CAR5*, *CAR10*, and *CAR20* are defined as the sum of abnormal returns over the (0,1), (0,4),(0,9), (0,19) event windows. *Count* is the day zero frequency of posts. *Count20* and *Impact20* are the summations of *Count* and *Impact* over the (-1,-20) event window. Ten daily lags of *ar* are added as controls. All regressions are firm fixed effects panel data regressions. T-statistics are calculated using Driscoll-Kraay (1998) standard errors that control for cross-sectional and time-series correlation and are reported in parentheses \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

| <b>Panel A: Contemporaneous and Lagged Count</b>          |                      |                      |                      |                      |                      |
|-----------------------------------------------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| VARIABLES                                                 | (1)<br>ar            | (2)<br>CAR2          | (3)<br>CAR5          | (4)<br>CAR10         | (5)<br>CAR20         |
| Count                                                     | 0.260***<br>(4.81)   |                      |                      |                      |                      |
| Count20                                                   | -0.128***<br>(-4.42) | -0.084***<br>(-3.18) | -0.207***<br>(-3.52) | -0.385***<br>(-4.52) | -0.740***<br>(-7.22) |
| Observations                                              | 384,124              | 384,124              | 384,124              | 384,124              | 384,124              |
| Number of groups                                          | 1,416                | 1,416                | 1,416                | 1,416                | 1,416                |
| Lagged AR Control                                         | Yes                  | Yes                  | Yes                  | Yes                  | Yes                  |
| <b>Panel B: Contemporaneous and Lagged Network Impact</b> |                      |                      |                      |                      |                      |
| VARIABLES                                                 | (1)<br>ar            | (2)<br>CAR2          | (3)<br>CAR5          | (4)<br>CAR10         | (5)<br>CAR20         |
| Impact                                                    | 0.149***<br>(3.10)   |                      |                      |                      |                      |
| Impact20                                                  | -0.075***<br>(-4.10) | -0.078***<br>(-4.11) | -0.191***<br>(-4.60) | -0.352***<br>(-5.77) | -0.620***<br>(-6.48) |
| Observations                                              | 384,124              | 384,124              | 384,124              | 384,124              | 384,124              |
| Number of groups                                          | 1,416                | 1,416                | 1,416                | 1,416                | 1,416                |
| Lagged AR Control                                         | Yes                  | Yes                  | Yes                  | Yes                  | Yes                  |

Table 1.3: **Contemporaneous and Lagged *Sentiment\_Impact***

This table presents panel regressions of abnormal returns on a measure of the sentiment of network communication. *ar* is defined as the abnormal returns from a Fama-French 4 Factor model with 150 day rolling estimation. *CAR2*, *CAR5*, *CAR10*, and *CAR20* are defined as the sum of abnormal returns over the (0,1), (0,4), (0,9), (0,19) event windows. *Sentiment\_Impact* is the day zero value of the social network sentiment variable. *Senti\_Impact20* is the standardized summation of *Sentiment\_Impact* over the (-1,-20) event window. All regressions are firm fixed effects panel data regressions. Ten daily lags of *ar* are added as controls. T-statistics are calculated using Driscoll-Kraay (1998) standard errors that control for cross-sectional and time-series correlation and are reported in parentheses \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

| VARIABLES         | (1)<br>ar            | (2)<br>CAR2          | (3)<br>CAR5          | (4)<br>CAR10         | (5)<br>CAR20         |
|-------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| Senti_impact      | 0.088***<br>(3.45)   |                      |                      |                      |                      |
| Senti_Impact20    | -0.038***<br>(-4.12) | -0.060***<br>(-3.94) | -0.143***<br>(-4.39) | -0.268***<br>(-5.02) | -0.463***<br>(-5.86) |
| Observations      | 384,124              | 384,124              | 384,124              | 384,124              | 384,124              |
| Number of groups  | 1,416                | 1,416                | 1,416                | 1,416                | 1,416                |
| Lagged AR Control | Yes                  | Yes                  | Yes                  | Yes                  | Yes                  |

Table 1.4: **Returns *Senti\_Impact* Across Size Groups**

This table presents panel regressions of contemporaneous and future abnormal returns on *Senti\_Impact*. Panels A and B present the result for small and large stocks in our sample respectively. The difference in the estimated coefficients of *Senti\_Impact* and *Senti\_Impact20* across the two groups are reported in the rows marked (1)-(3) and (2)-(4) respectively. *ar* is defined as the abnormal returns from a Fama-French 4 Factor model with 150 day rolling estimation. *CAR2*, *CAR5*, *CAR10*, and *CAR20* are defined as the sum of abnormal returns over the (0,1), (0,4), (0,9), (0,19) event windows. *Sentiment\_Impact* is the day zero value of the social network sentiment variable. *Senti\_Impact20* is the standardized summation of *Sentiment\_Impact* over the (-1,-20) event window. All regressions are firm fixed effects panel data regressions. Ten daily lags of *ar* are added as controls. Driscoll and Kraay (1998) robust t-stats are reported below coefficients

| <b>Panel A: Small</b> |                      |                      |                      |                      |
|-----------------------|----------------------|----------------------|----------------------|----------------------|
| VARIABLES             | (1)<br>ar            | (2)<br>CAR2          | (3)<br>CAR10         | (4)<br>CAR20         |
| Senti_impact          | 0.466***<br>(2.92)   |                      |                      |                      |
| Senti_Impact20        | -0.264***<br>(-3.25) | -0.476***<br>(-3.13) | -2.239***<br>(-3.85) | -3.810***<br>(-4.98) |
| Observations          | 110,621              | 110,621              | 110,621              | 110,621              |
| Number of groups      | 646                  | 646                  | 646                  | 646                  |
| <b>Panel B: Large</b> |                      |                      |                      |                      |
| VARIABLES             | (1)<br>ar            | (2)<br>CAR2          | (3)<br>CAR10         | (4)<br>CAR20         |
| Senti_impact          | 0.034***<br>(2.62)   |                      |                      |                      |
| Senti_Impact20        | -0.016**<br>(-2.21)  | -0.023<br>(-1.53)    | -0.103**<br>(-2.09)  | -0.194***<br>(-3.21) |
| Observations          | 136,857              | 136,857              | 136,857              | 136,857              |
| Number of groups      | 536                  | 536                  | 536                  | 536                  |
| (1)-(3)               | 0.432***<br>(2.70)   |                      |                      |                      |
| tstat (1)-(3)         |                      |                      |                      |                      |
| (2)-(4)               | -0.248***<br>(-3.05) | -0.454***<br>(-2.97) | -2.136***<br>(-3.66) | -3.616***<br>(-4.72) |
| tstat (2)-(4)         |                      |                      |                      |                      |

Table 1.5: **Returns to Local & Non-Local Senti(SP500 Members)**

This table presents panel regressions of contemporaneous and future abnormal returns on *Senti\_Impact*. Panels A and B present the result for Non S&P500 and S&P500 stocks in our sample respectively. SP 500 membership is determined on a monthly basis. The difference in the estimated coefficients of *Senti\_Impact* and *Senti\_Impact20* across the two groups are reported in the rows marked (1)-(3) and (2)-(4) respectively. *ar* is defined as the abnormal returns from a Fama-French 4 Factor model with 150 day rolling estimation. *CAR2*, *CAR10*, and *CAR20* are defined as the sum of abnormal returns over the (0,1), (0,9), (0,19) event windows. *Sentiment\_Impact* is the day zero value of the social network sentiment variable. *Senti\_Impact20* is the standardized summation of *Sentiment\_Impact* over the (-1,-20) event window. All regressions are firm fixed effects panel data regressions. Ten daily lags of *ar* are added as controls. Driscoll and Kraay (1998) robust t-stats are reported below coefficients

| <b>Panel A: Non SP500 stocks</b> |                      |                      |                      |                      |
|----------------------------------|----------------------|----------------------|----------------------|----------------------|
| VARIABLES                        | (1)<br>ar            | (2)<br>CAR2          | (3)<br>CAR10         | (4)<br>CAR20         |
| Senti_impact                     | 0.147***<br>(3.04)   |                      |                      |                      |
| Senti_Impact20                   | -0.079***<br>(-3.56) | -0.152***<br>(-3.80) | -0.646***<br>(-4.27) | -1.031***<br>(-4.69) |
| Observations                     | 291,897              | 291,897              | 291,897              | 291,897              |
| <b>Panel B: SP500 Stocks</b>     |                      |                      |                      |                      |
| VARIABLES                        | (1)<br>ar            | (2)<br>CAR2          | (3)<br>CAR10         | (4)<br>CAR20         |
| Senti_impact                     | 0.041***<br>(2.73)   |                      |                      |                      |
| Senti_Impact20                   | -0.014*<br>(-1.70)   | -0.017<br>(-1.01)    | -0.093<br>(-1.60)    | -0.197***<br>(-2.75) |
| Observations                     | 92,227               | 92,227               | 92,227               | 92,227               |
| Number of groups                 | 334                  | 334                  | 334                  | 334                  |
| (1)-(3)                          | 0.106**<br>(2.10)    |                      |                      |                      |
| tstat (1)-(3)                    |                      |                      |                      |                      |
| (2)-(4)                          | -0.065***<br>(-2.75) | -0.135***<br>(-3.11) | -0.553***<br>(-3.42) | -0.834***<br>(-3.61) |
| tstat (2)-(4)                    |                      |                      |                      |                      |

**Table 1.6: Information Flow Through a Network**

This table present autoregressions of the level of network sentiment. *Sentiment\_Impact* is the day zero value of the social network sentiment variable. The variables *Week1Senti\_Impact*, *Week2Senti\_Impact*, *Week3Senti\_Impact* and *Week4Senti\_Impact* are defined as the summation of *Senti\_Impact* over the (-1,-5),(-6,-10),(-11,-15) and (-15,-20) event windows respectively. *CARLag5* and *CARLag10* are defined as the summation of abnormal returns over the (-1,-5),(-6,-10) horizon respectively. T-statistics are calculated using Driscoll-Kraay (1998) standard errors that control for cross-sectional and time-series correlation and are reported in parentheses \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

| VARIABLES         | (1)<br>Senti_impact | (2)<br>Senti_impact | (3)<br>Senti_impact |
|-------------------|---------------------|---------------------|---------------------|
| Week1Senti_Impact | 0.084***<br>(5.98)  | 0.083***<br>(5.88)  | 0.082***<br>(5.86)  |
| Week2Senti_Impact | 0.054***<br>(3.11)  | 0.051***<br>(3.20)  | 0.051***<br>(3.21)  |
| Week3Senti_Impact |                     | 0.014<br>(1.10)     | 0.014<br>(1.11)     |
| Week4Senti_Impact |                     | 0.009<br>(0.62)     | 0.009<br>(0.63)     |
| CARLag5           |                     |                     | 0.001***<br>(2.70)  |
| CARLag10          |                     |                     | 0.000<br>(1.40)     |
| Observations      | 384,128             | 384,124             | 384,124             |
| Number of groups  | 1,416               | 1,416               | 1,416               |

**Table 1.7: Factor Regression for Impact and Sentiment Impact**

Portfolios are created based on the previous weeks level of the variables *Impact* and *Senti\_Impact*. Stocks are double sorted by size and previous level of network communication. Size is determined using NYSE break points taken from the Ken French library. All portfolios are equally weighted. The table presents only the intercept term from a Fama-French three-factor regression. T-statistics are calculated using Driscoll-Kraay (1998) standard errors that control for cross-sectional and time-series correlation and are reported in parentheses \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

| Formation Period = 1 week |                      |                   |                   |                    |
|---------------------------|----------------------|-------------------|-------------------|--------------------|
| Hold Period = 1 day       |                      |                   |                   |                    |
| 3 Factor Alphas           |                      |                   |                   |                    |
|                           | Notalk               | Medtalk           | Hightalk          | L-S                |
| Senti_Impact              | 0.0007***<br>(0.000) | 0.0001<br>(0.001) | 0.0011<br>(0.002) | -0.0004<br>(0.002) |
| Impact                    | 0.0007***<br>(0.000) | 0.0005<br>(0.001) | 0.0013<br>(0.001) | -0.0006<br>(0.001) |
| Count                     | 0.0007***<br>(0.000) | 0.0010<br>(0.001) | 0.0006<br>(0.001) | 0.0001<br>(0.001)  |

**Table 1.8: Attention vs No Attention**

This table presents panel regressions of abnormal returns on a measure of the sentiment of network communication as well as a dummy variable to control for lack of firm level attention. *None* is a dummy variable that is equal to one if there is no network communication about a firm in the previous twenty trading days. *ar* is defined as the abnormal returns from a Fama-French 4 Factor model with 150 day rolling estimation. *CAR2*, *CAR5*, *CAR10*, and *CAR20* are defined as the sum of abnormal returns over the (0,1), (0,4), (0,9), (0,19) event windows. *Impact* is the day zero value of the social network sentiment variable. *Impact20* is the standardized summation of *Impact* over the (-1,-20) event window. All regressions are firm fixed effects panel data regressions. Ten daily lags of *ar* are added as controls. T-statistics are calculated using Driscoll-Kraay (1998) standard errors that control for cross-sectional and time-series correlation and are reported in parentheses \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

| VARIABLES         | (1)<br>ar            | (2)<br>CAR2          | (3)<br>CAR5          | (4)<br>CAR10         | (5)<br>CAR20         |
|-------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| Impact            | 0.149***<br>(3.10)   |                      |                      |                      |                      |
| Impact20          | -0.072***<br>(-3.95) | -0.072***<br>(-3.83) | -0.177***<br>(-4.30) | -0.328***<br>(-5.43) | -0.574***<br>(-6.07) |
| None              | 0.067***<br>(3.43)   | 0.134***<br>(3.76)   | 0.309***<br>(3.92)   | 0.543***<br>(4.04)   | 1.022***<br>(5.34)   |
| Observations      | 384,124              | 384,124              | 384,124              | 384,124              | 384,124              |
| Number of groups  | 1,416                | 1,416                | 1,416                | 1,416                | 1,416                |
| Lagged AR Control | Yes                  | Yes                  | Yes                  | Yes                  | Yes                  |

## Chapter 2

# The Impact of Divergence of Opinions about Earnings within a Social Network

### **Abstract**

We collect a unique dataset of Twitter posts and use these posts to contrast investors' private opinions about stocks with opinions available in public press releases. This technique allows us to directly measure the divergence of private investor's opinions from publicly available information thus, providing a solution for the problem of unavailable private opinions cited by Garfinkel (2010). We find that post-earnings announcement returns are significantly more negative when divergence of opinion is present in the pre-announcement period. These negative post announcement returns are concentrated into the stocks whose divergence of opinion is resolved into agreement after the announcement.

# 1 Introduction

Models of trade and price movement often derive their results from heterogeneous agents. If agents in a model are not identical then they have incentives to negotiate in order to determine a mutually agreeable equilibrium. The source of heterogeneity between agents can come from differences in endowments, differences in risk aversion, or differences in information processing. In a unique way, this paper analyzes the impact that divergence of opinion over a public signal has on stock prices and volumes around earnings announcements. Instead of using proxy variables such as income volatility (Berkman, Dimitrov, Jain, Koch and Tice (2009)), analyst forecast dispersion (Deither, Malloy and Scherbina (2002), Berkman et al. (2009)) or abnormal volume (Garfinkel and Sokobin, (2006)) we measure differences of opinion by contrasting the tone of social media posts against the tone of publicly available news releases about the announcing firm. We use our unique dataset of social network investor opinions to better understand how investors' private opinions affect returns around earnings announcements. Using this method we hope to provide information, and a possible solution, to the problem pointed out in Garfinkel (2010) that measures of divergence of opinion found in the literature are proxies for the unobservable private opinions of investors and not those private opinions themselves.

An early theory on the implications of divergence of opinion on stock returns is Miller (1977) who predicts that in the presence of short sale constraints, optimists about the future prospects of the stock set current price. As divergence of opinion increases, the truncated mean of the optimist-only distribution of future valuations also increases; as a result, higher divergence of opinion leads to higher current prices and lower future returns. Using the standard deviation of analyst's earnings forecasts as a measure of divergence of opinion Deither et al. (2002) find that stocks with volatile forecast distributions earn lower future returns, a result they claim is consistent with Miller (1977). Using our direct measure of

divergence of opinion between social network investors and public news we construct a unique test of Miller's (1997) hypothesis and at the same time we investigate the robustness of recent research about divergence of opinion and returns around earnings announcements.

Miller (1997) is a static model in which investors take initial positions and reevaluate these positions as information is revealed. As such, it is of limited usefulness as a model that can be used to generate predictions for divergence of opinion and volume at the time of an earnings release. Well known models such as Harris and Raviv (1993) or Kandel and Pearson (1995) that predict volume at the time of an information event are also of limited usefulness as these models are primarily interested in differential interpretations of information and make assumptions to highlight this effect. For example, Harris and Raviv (1993) impose common priors, an assumption violated if our empirical data finds divergence of opinion before the earnings release. The closest theoretical concept to our measure of divergence of opinion is the *consensus* opinion construct in Holthausen and Verrecchia (1990); investors will either be in general agreement with public news articles or they will not. Using the Holthausen and Verrecchia (1990) theoretical framework, we consider the public news to be a measure of the information investors conjecture from prices. Given this assumption, we consider disagreement between Twitter investors and public news articles a signal of low precision in the information investors receive from prices. In this framework disagreement represents a decline in the commonality of investor's information sets and Twitter posts can provide a proxy for divergence of opinion that is more flexible and useful than proxies based on the static model in Miller (1977). In Holthausen and Verrecchia (1990) a reduction in consensus implies an increase in divergence of opinion and this increase is theorized to be associated with increases in volume after the earnings release.<sup>1</sup>

---

<sup>1</sup>This volume-divergence link is relied on by Garfinkel and Sokobin (2006) who use unexpected volume as a proxy for divergence of opinion around earnings announcements.

Our measure of investors' opinions is constructed using influential posters on Stocktwits.com, an investment-focused subsidiary of Twitter, the popular social media website. We use text readers to identify the tone of both Twitter posts and public news releases for 1,398 companies releasing 5,433 quarterly earnings announcements in the October, 2008 to December, 2009 period. Our textual analysis techniques classify the overall tone and the strength of the opinion for both the Twitter posts and the contemporaneous news releases. We match the tone of the Twitter posts to the tone of the news releases before and after the earnings announcement date. In this way, we can classify whether investors tended to agree or disagree with publicly available news before the earnings announcement. Further, we can document whether investors' opinions converge or diverge from public news announcements after the earnings announcement occurs. This structure of agreement or disagreement allows us to examine predictions about both the level of pre-announcement divergence of opinion but also to test theoretical assumptions about any changes in divergence of opinion subsequent to the release of the earnings information.

Our measure of divergence of opinion differs from existing measures based on the dispersion of financial variables. At its core, our measure is a discrete metric contrasting private opinions in the Twitter social network with the information in publicly available news articles. We also weight the discrete measure of agreement or disagreement by the *Impact* (or following) of particular individual investors. Divergence of opinion in this paper is therefore representative of differential interpretation or uncertainty about the interpretation of public information. Our measure is largely orthogonal to recently published work using measures of dispersion or unexpected volume, yet produces similar quantitative results which are robust to controlling for the measures of dispersion used in other work. This finding suggests that both financial variable proxies for dispersion and our measure of different interpretations of common information can play independent roles in explaining the empirical results on divergence of opinion and stock returns.

Recent work on divergence of opinion and earnings announcements includes Berkman et al. (2009) who find that proxies for divergence of opinion are associated with returns around the earnings announcement. Specifically, high divergence of opinion is correlated with lower short-window excess returns at earnings announcements. Garfinkel and Sokobin (2006) propose a market-corrected unexpected turnover measure as a proxy for divergence of opinion. They find that unexpected turnover at the earnings announcement is positively related to future returns. While these results are not strictly contradictory because they do not use the same event window, they each seem to lead the reader to a different conclusion. In Berkman et al. (2009) the interpretation of their results hinges on the assumption that the earnings announcement resolves uncertainty. Whereas in Garfinkel and Sokobin (2006) they do not assume that uncertainty is resolved, instead their unexpected volume measure is intended to proxy for the level of uncertainty at the time of the announcement.

We find significantly lower future returns when differences of opinion exist prior to the earnings announcement. At the time of the earnings announcement, 2-day announcement returns (-1,0) are comparable across the disagreement (0.61%,  $t$ -statistic = 2.55) and agreement (0.51%,  $t$ -statistic = 4.34) portfolios. Immediately after the announcement period the returns diverge across these two groups. We find that the 60 day post-earnings announcement drift after announcements associated with agreement are statistically insignificant (-0.57%,  $t$ -statistic = -0.89) while the same drift after announcements associated with investor disagreement is significantly negative (-4.00%,  $t$ -statistic = -3.04). The difference between the returns to these two groups is -3.43% ( $t$ -statistic = 1.96). The sign of the actual earnings surprise does not affect the divergence of opinion result; similar results are found for both high standardized unexpected earnings (SUE) stocks and low SUE stocks. Furthermore, our social media measure of divergence of opinion is uncorrelated with unexpected volume. In regressions that control for unexpected volume, as proposed by Garfinkel and Sokobin

(2006), we find that our divergence of opinion measure derived from Twitter posts is still negative and significant and that unexpected volume has no significant effect.

As Berkman et al. (2009) admit, it is difficult to measure changes in divergence of opinion because many of the financial variable proxies are noisy or do not change in a timely manner. Another contribution of our paper is that our method can uniquely classify the *change* in divergence of opinion as a result of the earnings announcement and contrast earnings announcement returns when the announcement resolves prior disagreement against those cases when it does not. This is important as it is often assumed that the earnings announcement reduces asymmetry in the market.<sup>1</sup> We examine ten-day post announcement returns when divergence of opinion measures change. When the correlation between Twitter opinions and the tone of media reports moves from agreement prior to the earnings announcement to disagreement afterwards, abnormal returns are positive. Conversely, when the correlation between Twitter opinions and the tone of media reports move from disagreement before to agreement after the earnings announcements, returns are negative.

We also report the abnormal volume effects of these changes in divergence of opinion. We sort post-earnings announcement volume into four groups depending on whether social network investors and the public press agree or disagree both before and after the earnings announcement. Disagreements that are unresolved have significantly higher post-earnings abnormal volume than prior disagreements that are resolved into agreement. We interpret this result as reflecting the Holthausen and Verrecchia (1990) hypothesis that one effect of earnings announcements can be a change in the level of investor's consensus and changes in consensus are reflected in volume in the expected way.

---

<sup>1</sup>Lower asymmetry as a result of the announcement is not always assumed, Kim and Verrecchia (1994) is one notable exception.

We conclude the paper by using text analysis software to determine if there are any linguistic features of public press releases that are associated with a pre-earnings announcement divergence of opinions. We find that press releases that lack words reflecting optimism and certainty make it more likely that investors will disagree with the press release information. This finding is apparently not a reflection of press releases that do not provide enough repetition as repetitive language tends to induce investors to disagree with the sentiment of press release; perhaps excessive insistence in the press release is regarded as strident and thus, unconvincing to investors. Not only are there potential negative valuation effects associated with divergence of opinions that should matter to investor relations departments but our findings also have implications for corporate disclosure policy that simply wants to convince investors that the expressed opinions are sound.

Section 2 outlines some basic hypotheses regarding divergence of opinions around earnings announcements. Section 3 describes the data, Section 4 the results. Section 5 analyzes the linguistic content of press releases. Section 6 concludes.

## 2 Theory and Hypothesis Development

If investors have heterogeneous priors, they can come to different conclusions about the future value of an asset even when exposed to identical public news. Miller (1977) is an early attempt to loosen the assumption of homogeneity and allow for disagreement among investors. Miller's model makes the key assumption that investors are subject to short sale constraints. In practice, the short sale assumption means that if investors think a stock is overvalued they do not short the stock; they either sell the shares that they own of the stock or stay out of the market. There is strong empirical evidence for the validity of this assumption. For example, Lamont and Stein (2004) show that although hedge funds are known to short sell aggressively, short interest remains a small portion of total shares

outstanding at any one time. Furthermore, most individual investors and mutual funds never take short positions; see Almazan, Brown, Carlson and Chapman (2004).<sup>2</sup>

Miller (1977) shows that in a world with short sale constraints and divergence of opinion among investors, an asset's price will reflect the valuations of the most optimistic investors and will exclude the opinions of pessimistic investors. Furthermore, if a small group of investors is able to absorb the entire supply of a given security then this optimistic group of investors will be able to set the price of the security above the mean valuation of the market. Thus, divergence of opinion between market participants will increase the market clearing price of a security. Since divergence of opinion causes over pricing there should be a negative relationship between divergence of opinion and future expected returns. If an asset is overpriced due to disagreement, if the earnings announcement resolves this disagreement, any overpricing should disappear and the asset's value will fall to reflect the mean evaluation of all investors rather than the just the valuation of the subset of optimistic investors. On the contrary, if investors agree about the valuation of an asset prior to the announcement, it will be fairly priced at the mean valuation of all investors.

In our setting Miller's (1977) theory implies that disagreement before an earnings announcement should cause negative expected abnormal returns after the earnings announcement because the release of the official earnings number should eliminate disagreement among investors about the company's current earnings and future prospects. Our unique data, which measures the sentiment of investors' opinions, allows us to directly test whether divergence of opinion by actual investors is related to the decrease in expected returns predicted by Miller (1977). This relationship between investor sentiment and returns yields Hypotheses 1a and 1b.

---

<sup>2</sup>Consistent with Sorescu, Boehme and Danielson (2006) our results are stronger when it is more likely that short-sale constraints are binding. Stocks with lower market capitalizations or lower institutional ownership exhibit stronger results.

*Hypothesis 1a: Divergence of opinion between network communication and public information before an earnings announcement will cause future negative abnormal returns after the earnings announcement. These negative abnormal returns will persist until the asset's value reflects the mean of all investors' valuations.*

*Hypothesis 1b: If there is agreement between network communication and media sentiment then there will be no impact on abnormal returns after an earnings announcement.*

Empirically, Hypothesis 1a implies that after controlling for all other factors, there should be a negative relationship between investor disagreement and future abnormal returns. Hypothesis 1b implies that investor agreement should have no statistical relationship with future abnormal returns.

## **2.1 Volume and Divergence of Opinion**

As Hong and Stein (2006) discuss, Miller (1997) is a static model and thus, has nothing to say about volume and divergence of opinion. One dynamic approach explaining divergence of opinion and volume is presented in Harris and Raviv (1993) who develop a model where trading is generated by differences of opinion about the value of an asset. These differences of opinion are generated because the speculators in the model have different interpretations of common public information and each group of traders uses their own likelihood function to update their beliefs about an asset after a new signal is revealed. Once each group establishes their beliefs they believe in them absolutely, causing a divergence of opinions. Trading is generated in the model when the two groups disagree on the importance of a specific piece of information. One group will be responsive to the information while the other will not, thus generating trades. One of the central predictions of this model is that disagreement should be positively related to speculative volume. However, to model divergence of opinion around earnings announcements Harris and Raviv (1993) assume that investors have common priors. This assumption is clearly contradicted if we document divergence of opinion prior to the earnings announcement.

A useful alternative perspective is provided by Holthausen and Verrecchia (1990). They contend that the release of earnings information can change both the informedness of investors and the degree of consensus between investors. They contend that both greater informedness and reduced consensus cause greater trading volume. This theoretical link between consensus and trading volume is exploited by Garfinkel and Sokobin (2006) who contend that abnormal volume at the time of the announcement can measure the degree of dispersion of beliefs among investors.

As we construct a continuous measure of divergence of opinion between network communication and public information, we can directly measure the degree of consensus or disagreement in the market both before and after the earnings announcement. Holthausen and Verrecchia’s (1990) theoretical links between consensus and volume provides the testable implications for our measure of divergence of opinion and trading volume.

*Hypothesis 2a: Divergence of opinion between network communication and the public information reflects reduced consensus and should be related to greater trading volume.*

We can also measure changes in the divergence of opinion around the earnings announcement. These changes should reflect changes in Holthausen and Verrecchia’s (1990) consensus measure and have measurable volume effects.

*Hypothesis 2b: Changes in divergence of opinion around the earnings announcement should be related to volume. If the earnings announcement reduces consensus, as measured by our divergence of opinion proxy, volume should increase. If the earnings announcement increases consensus, as reflected by agreement between network communications and public information, then volume should be less.*

### 3 Data Collection and Variable Construction

This section describes the data collection and storage process used for Twitter posts and news articles. This communication data will later be processed for use in testing the predictions of Section 2.

### 3.1 Soft Information Collection

We use the website Stocktwits.com to select the sample of Twitter users. This site provides a list of contributors that are marked as “Recommended” posters. Although the website does not use a quantitative rule for selecting the recommended contributor, in general, the recommended contributors have a large following, have a long track record, post meaningful or interesting comments, and are influential within the social network. We select our sample of Twitter users from the set of recommended users as of October 1, 2008. We require that the Twitter users have at least 100 followers and 300 posts at the beginning of data collection process. This set of criteria produces a sample of 74 Twitter contributors.

All of the posts for the sample of 74 contributors from October 1, 2008 to December 31, 2009 are downloaded from Twitter.com using web scraping software specially designed for this particular website. The software compiles and stores each message in a database and assigns a unique name based on its author, date and time. This process yields 143,640 individual posts that reference 1,469 different tickers. We also add all of the S&P 500 members from 2009 to this dataset. We add these firms to control for the findings of Chan (2003) and Fang and Peress (2009) that “No News” firms outperform firms with news coverage on average. This increases the initial sample to 1,758 companies because 289 tickers which are in the S&P 500 were not tweeted about during the year.

Next, we eliminate all ETFs, companies that are smaller than 100 million in market capitalization, and companies with trading price lower than three dollars in 2009. We are left with 1,398 companies in the sample. Finally, we collect information about the structure of the social network of each of the 74 individual contributors in the sample. Specifically, at the end of the sample period, we record the number of followers, the number of people followed, and the number of posts for each author in the sample.

To get a better understanding of the nature of a “tweet” it is helpful to look at an example from Stocktwits. Figure 2.1 is a sample of the data items on the Stocktwits website. It is a real time aggregated list of messages from all members of the Stocktwits network. Each

message has the author, date, and time listed at the very top of the post, as well as a picture that the author provides for their online profile. Each post can be no longer than 140 characters and it is usually a short declarative statement about a company or the economy. A viewer of the site sees a continual stream of financial topics that are the most interesting to the Twitter community. When a blogger wants to post about a particular company on the Stocktwits website, they tag the company's ticker symbol with a "\$". For example, if you wanted to talk positively about Google and Microsoft you would say, "\$GOOG and \$MSFT, you should buy!". This is a practice called "Hashtagging" which is common place in the Twitter community. The "\$" hashtag allows us to extract the company references in each post with a high level of accuracy. The classification of posts by company ticker symbol is done using a regular expression algorithm that searches the text for a "\$" followed by 1-4 letters that constitute the ticker symbol. In the case where a blogger talks about multiple companies in one post, like in our example above, all references are counted as a unique comment. The final output of the text parsing procedure is a list which associates each post to a unique ticker and author.

Once all posts are parsed for company references we use the author specific data of followers, following and number of posts to construct a variable which measures the specific influence that each contributor has on the network. Only posts during the trading week are included in the analysis because posts on the weekend or on non-trading days cannot be contemporaneously compared to market data. News articles are taken from the Dow Jones Factiva news database. Since we are interested in breaking news and company press releases we only include articles from PR News Wire, Dow Jones News Wire, and Reuters News.

To execute the search process we develop a custom application to query the Factiva website. First, we collect the full company name for each ticker from Compustat. The full company name is input into our program and used as the search term within the Factiva database. Once Factiva runs a search, it provides a list of company names that are the most relevant to your search terms. The program then selects the company name that is the most relevant to the search name that we input. The relevant name is then cross referenced

against the company name. If Factiva does not provide any search results for a particular company name we assume that the company has zero news coverage over the period. All articles for firms with news coverage are downloaded and stored in a database. We include all revisions of articles in the sample. All descriptive information about the article such as, date and time published, news source, and article ID are also recorded. This process yields 74,110 company specific press releases over the sample period.

The final sample contains 1,138 companies, this includes all of the S&P 500 constituents from 2009 plus any company that was mentioned in a tweet during the sample period and had a price over three dollars and a market capitalization over 100 million.<sup>3</sup> Daily stock returns are taken from CRSP. Other firm level descriptive data comes from the Compustat annually updated database, and analyst information comes from IBES.

## 3.2 Sentiment and Content Classification

In this section, we describe the methodology used to classify the sentiment of Twitter posts and press releases and the content of press releases. Computational linguistics (CL) is a field dealing with statistical or rule based modeling of natural language. CL can be used for categorical classification and well as extraction of tone or sentiment.

Sentiment classification is one of the hardest tasks for CL to handle because levels of sentiments can not always be determined through key words alone. For example, in the phrase, “You would be stupid to buy \$XYZ,” if we look at each word in isolation it is hard to determine that this is a negative post. The word stupid has a negative tone, but the word buy is obviously a positive word. However, the combined tone of the statement is unquestionably negative. Some classification methods such as naive Bayesian classification, the most popular technique for document classification, often have trouble with sentiment classification because it assumes the conditional independence of words in a document. For our purposes, Maximum Entropy (ME) classification is better because it does not make

---

<sup>3</sup>We wish to compare our divergence of opinion measure with the dispersion of analysts’ forecasts measure. The requirement that a stock be covered by analysts results in the final sample truncation to 1,138 firms.

any assumptions about the conditional independence of words; ME tends to perform better when conditional independence does not hold. In our example, ME would be able to use the information that is contained in the grouping of words “stupid to buy” to determine that this post is in fact negative. Pang, Lee and Vaithyanathan (2002) show that ME classification is the most robust for sentiment classification. Thus, we use ME classification throughout this paper.

The general idea that drives ME classification is that when nothing is known about a distribution, the distribution should be as uniform as possible, (having maximum entropy). In the case of sentiment classification, the ME classifier estimates the conditional distribution of the sentiment classes given a particular document, where all documents are represented by the frequency of features contained in them. A feature is a single word or a grouping of words such as “short sell” or “I hate this stock”. We use training data to estimate the expected values of features on a class by class basis. Finally, the classifier will determine the probability that a document falls into a particular class based on the features of the document and the expected values of features conditional on categories.

Consider the example of trying to classify a document as positive, negative or neutral, where we are only told that 50% of documents that contain the word “buy” are considered positive. Intuition tells us that if the document has the word “buy” in it then there is a 50% chance that it is a positive post and a 25% chance of being either negative or neutral. If our document did not have the word “buy” in it then we would just assume an equal distribution of a 33% chance that the document falls into each category. Thus, when we know nothing about our document, we make the distribution uniform. This is the essence of ME classification. In practice, this process is constrained by many features, and the calculations for conditional probabilities become complex, but the logic is still the same as our simple example.

To formally describe the ME procedure, we define the following set of terms. Let  $f_1, \dots, f_m$  be a set of predefined features that can appear in a post. From our previous

example, the word “buy” would be a feature, and the bi-gram ”short sell” could be a feature. Let  $n_i(d)$  be the number of times that the feature  $f_i$  occurs in a document  $d$ . Thus, each document is represented by a document vector that takes the form:  $d = (n_1(d), n_2(d), \dots, n_m(d), )$ . Lastly, let  $c$  be a document category of either positive, negative or neutral. Given this set of variables, the estimate of  $P(c|d)$  is as follows:

$$P_{ME}(c|d) = \frac{1}{Z(d)} \left( \sum_i \lambda_{i,c} F_{i,c}(d, c) \right) \quad (2.1)$$

Where  $Z(d)$  is a normalization function, and  $F_{i,c}$  is a feature class function for the feature  $i$  in the class  $c$ .

$$F_{I,c}(d, c) = 1, n_i(d) > 0, \text{ and } c_i = c_0, 0 \text{ otherwise} \quad (2.2)$$

For example, this feature class function could only return a value of one if the post contained the bi-gram “short sell” and the post is hypothesized to be of negative sentiment. Finally,  $\lambda_{i,c}$  is a weighting parameter that determines the relative value of each of the features  $f_i$  contained in a document. If the value of  $\lambda_{i,c}$  is very large then the feature  $f_i$  is considered to be very strong for the document class  $c$ .<sup>4</sup> Finally, the conditional probabilities  $P_{ME}(c|d)$  are estimated by maximizing the entropy across the three different groups while satisfying the constraint that the expected values of the feature class functions  $F_{i,c}$  are equal to their training data expected values.

To implement this methodology to the Twitter dataset we start by hand classifying an out of sample set of 1,000 posts. This set of 1,000 hand classified posts is defined as the training set. From this training set, the expected values of the  $F_{i,c}$  are calculated. Then the relationships that are documented in the training data are used to calculate  $P_{ME}(c|d)$ , in our case, the probability that a message is positive, negative or neutral. The ME classifier is run on each post, and it produces a conditional probability for each class (negative, neutral, positive). Each post is assigned the sentiment of the class with the highest conditional

---

<sup>4</sup>For example, the word ”buy” could be strongly associated with positive sentiment relative to other words or phrases. Thus, the word ”buy” would have a large lambda.

probability. We assign the values  $(-1,0,1)$  to the sentiments (negative, neutral, positive). We test the accuracy of this procedure by running the ME classifier on a set of 100 posts that are hand classified. The ME classifier worked well in this out of sample test, and it was able to correctly classify 67% of all posts in the test sample. This accuracy rate is similar to the accuracy level that is achieved in other sentiment classification studies (see e.g. Pang, Lee and Vaithyanathan 2002).

This process is repeated on the sample of press releases. First, an out of sample training set of 300 press releases is hand classified for both sentiment as well as content. *Sentiment* is defined in the same way as the Twitter posts  $(-1,0,1)$  and content is defined as a 1 if the article is about earnings and 0 if the article is not about earnings. From this set of pre-classified training data we construct a ME classifier that determines both the sentiment and content of each press release in our sample.

### 3.3 Explanatory Variable Construction

We use the Sentiment measure derived from our ME classification to construct a variable that measures the social network impact of each post as a function of the number of followers an author has. We call this variable *Impact* and it is defined as follows:

$$Impact = ((1 + Followers) * Sentiment), \quad (2.3)$$

where the *Sentiment* measure is defined over the set  $(-1,0,1)$  depending on whether the sentiment of the post is negative, neutral or positive respectively. *Followers* is the number of followers that an author has on Twitter. Posts are summed up over the day to determine the aggregate level of *Impact*. The Demarzo, Vayanos and Zweibel (2003) theory of social network communication states that an agent's importance is a function of the size of his network. *Impact* accounts for this important feature of the DeMarzo et al. (2003) model by measuring the number of followers that are reached each time an agent posts a comment on Twitter. Finally, the social network impact of each post is interacted with its sentiment so

that we measure both the magnitude and direction of each comment. Thus, large positive values of *Impact* denote a broad dissemination of positive sentiment, and large negative values denote a broad dissemination of negative sentiment.

Next, we create a set of variables that utilize the sentiment of press releases to measure the tone of formal news. Once again, *Sentiment* is defined as positive, neutral or negative (+1, 0, -1) for each press release in the sample. We aggregate the sentiment of all news articles that pertain to earnings to a daily frequency to create the variable *News*. The interpretation of *News* is very intuitive because large positive values denote strong positive sentiment, and large negative values denote strong negative sentiment. Finally, we cumulate daily values of *Impact* and *News* over the (-1,-10) event window prior to the earnings announcement:

$$Impact_j = \sum_{t=-1}^{t=-10} Impact_{jt} \quad (2.4)$$

$$News_j = \sum_{t=-1}^{t=-10} News_{jt} \quad (2.5)$$

We are interested in the relationship between the sentiment of popular media coverage and the opinions of the investing public. Thus, to explicitly define the divergence of opinion we create a set of dummy variables that define the nine unique possibilities of the interaction between news and Twitter posts. As we have already defined the *Sentiment* measure over the set (-1, 0, 1) depending on whether the sentiment of the post is negative, neutral or positive, the *Sentiment* component of Twitter posts (*Impact<sub>j</sub>*) and news (*News<sub>j</sub>*) is used to contrast Twitter *Sentiment* against news *Sentiment*. We can think of the interaction between the *Sentiment* of news and Twitter as a 3 × 3 matrix with nine unique cell values. The shorthand *TP* (for Twitter positive) or *N0* (for news neutral) allows us to compactly outline our agreement or disagreement classification below. Using the variables *Impact<sub>j</sub>* and *News<sub>j</sub>* to aggregate the tone of both Twitter posts and news-based articles over the (-1,-10) pre-announcement window we create the variables *DIVOP* and *AGREE* from the set of nine interaction variables which measure divergence or agreement of opinion

between investors and the public news. The specific form of these variables is as follows:

$$\begin{aligned}
DIVOP &= \begin{cases} 1 & \text{if } TP\_NN \vee TP\_N0 \vee TN\_NP \vee TN\_N0 \vee T0\_NN \vee T0\_NP = 1 \\ 0 & \text{Otherwise} \end{cases} \\
AGREE &= \begin{cases} 1 & \text{if } TP\_NP \vee T0\_N0 \vee TN\_NN = 1 \\ 0 & \text{Otherwise} \end{cases}
\end{aligned}$$

In unreported tables we examine the comparative frequency of *DIVOP* and *AGREE* for all the days in the sample and for the subsample of earnings announcement days.<sup>5</sup> The frequency of these divergence of opinion measures is approximately the same between the full sample and the earnings subsample. On many individual days there are no media releases or Twitter posts. This finding is consistent with results in other studies done on firm-specific media coverage such as Fang and Peress (2010) and Tetlock Saar-Tsechansky and Macskassy (2008).

### 3.4 Dispersion Variables

Recent literature has proposed a number of variables that claim to measure divergence of opinion. We calculate these variables to determine how our measure of divergence relates to the standard financial variables used in the literature to measure dispersion of opinion. One common measure of divergence of opinion is the standard deviation of analysts' near-term earnings forecasts. Numerous papers, such as, Ajinkya, Atiase and Gift (1991), Abarbanell, Lanen and Verrecchia (1995), and Doukas, Kim and Pantzalis (2006), use divergence of analysts' forecast as a proxy for divergence of investor sentiment. The presumption that these studies make is that analysts express their unbiased opinions in their forecasts, and investors' opinions follow the opinions of analysts. Thus, if there is divergence in analysts' forecasts, this signals a divergence in analysts' opinions and thus, the opinions of investors.

---

<sup>5</sup>Under this definition note that *DIVOP* = 0 is equivalent to *AGREE* = 1 and defining both variables is a redundant way to define the  $3 \times 3$  matrix of outcomes. We do this as referring separately to agreement as *AGREE* = 1 produces clearer exposition when we discuss the effects of disagreement and agreement.

Deither et al.(2002) find that this measure of dispersion is negatively related to returns. Berkman et al. (2009) find that short-window returns around earnings announcements are lower for high forecast dispersion stocks. To compare our measure of divergence of opinion, we construct a measure of analysts' forecast dispersion using the detail tape from IBES. The forecast dispersion variable (*DISP*) is calculated using all analysts' estimates from the previous month as follows:

$$DISP_{i,t} = \frac{[\sum_{k=1}^K \frac{(Forecast_k - \overline{Forecast})^2}{K-1}]^{1/2}_{month}}{|\overline{Forecast}|}, \quad (2.6)$$

where  $Forecast_k$  is the  $k^{th}$  analyst's forecast of quarterly earnings per share and  $|\overline{Forecast}|$  is the absolute value of the mean analysts' forecast.<sup>6</sup>

Other proxies for divergence of opinion in Berkman et al. (2009) include the inverse of firm age (*AGE*), turnover (*TURN*), the average daily turnover during the pre-event period, stock return volatility (*RETVOL*), a measure of volatility relative to market volatility over the same pre-event period and earnings volatility (*INCVOL*) calculated using 20 quarters prior to the earnings announcement quarter. Berkman et al. (2009) find that these other variables based on financial market data have similar properties to *DISP* around earnings announcements.

Garfinkel and Sokobin (2006), in results somewhat in contrast to those in Berkman et al. (2009) find that the measure of divergence of opinion in their paper, standardized unexpected volume (*SUV*) is positively related to returns in the post-earnings announcement period. Although the results in these two papers do not strictly contradict each other (as they measure returns over different periods), conceptually it is surprising that divergence of opinion is found to be related to negative returns at the time of the announcement and positive returns in the post-announcement period. We estimate three turnover variables based on Garfinkel and Sokobin's (2006) measures including their controls for market-wide

---

<sup>6</sup>One potential pitfall of scaling analysts' forecasts by the mean forecast is that mean forecast near zero will cause very large values of *DISP*. Therefore, we test the robustness of our results using a measure of dispersion that is scaled by stock price. Our results are robust to this alternative specification

volume shocks, market adjusted turnover ( $MATO$ ), change in turnover ( $\Delta TO$ ) as well as  $SUV$ .

Garfinkel (2010) contends that using the raw level of volume as a proxy for disagreement does not adequately control for systematic factors that set our expectations of volume for any given asset. Therefore, we follow the spirit of Garfinkel and Sokobin (2006) and Garfinkel (2010) and estimate unexplained volume in a fashion similar to estimating abnormal returns from a market model. The variable that we create is standardized unexplained volume, or  $SUV$ , which is a standardized value of the prediction error from a regression of trading volume on absolute values of positive and negative returns. The specific form of the variable for firm  $i$  on day  $t$  is constructed by first estimating expected volume as:

$$E[Volume_{it}] = \hat{\alpha}_i + \hat{\beta}_1 |R_{it}|^+ + \hat{\beta}_2 |R_{it}|^-. \quad (2.7)$$

This model treats positive and negative return days as independent events; as the  $+$  and  $-$  superscripts in the regression equation illustrate. This asymmetry is introduced to capture the empirical regularity that volume reacts differently to the absolute returns on positive or negative return days Karpoff (1987). Unexpected volume is then calculated as the deviation from expected volume:

$$UV_{i,t} = Volume_{it} - E[Volume_{it}] \quad (2.8)$$

and finally standardizing unexpected volume by dividing by the standard deviation of volume, which is calculated using sixty trading days prior to the earnings announcement:

$$SUV_{i,t} = \frac{UV_{it}}{\sigma_{it}}. \quad (2.9)$$

This model specification is estimated on a rolling basis using a 60 trading day estimation period.

### 3.5 Control Variables

To measure the level of surprise that occurs around the earnings announcement we create a standardized unexpected earnings variable. We base our expectations model on analysts' forecasts. For each company in our sample, all earnings announcements dates as well as the end of quarter stock prices are collected from the Compustat database. We define our expectations of earnings as the median analyst estimate from the previous 90 days before the earnings announcement.<sup>7</sup> The median estimates are collected from IBES. We calculate standardized unexpected earnings as:

$$SUE = \frac{Actual - Expected}{P_t} \quad (2.10)$$

where  $P_t$  is the stock price at the end of the quarter.

The level of unexpected volume  $SUV$  may be affected by market volume at the time of the earnings announcement. As suggested in Garfinkel and Sokobin (2006) in our empirical tests of the effect of  $SUV$  on earnings announcement returns we control for unexpected market volume using both market adjusted turnover and change in market adjusted turnover. Market adjusted turnover is defined as:

$$MATO = \left\{ \sum_{t=-1}^{t=0} \left[ \left( \frac{Vol_{i,t}}{Shares_{i,t}} \right)_{firm} - \left( \frac{Vol_t}{Shares_t} \right)_{market} \right] \right\} / 2 \quad (2.11)$$

where  $Vol_{i,t}$  is the announcing firm's volume on day  $t$  and  $Shares_{i,t}$  is firm  $i$ 's shares outstanding on day  $t$  and day  $t$  is the earnings announcement day. Change in market adjusted turnover ( $\Delta TO$ ) is defined as daily  $MATO$  minus the median 180 moving average of  $MATO$ .

---

<sup>7</sup>We test  $SUE$  calculated from a random walk model as well as  $SUE$  calculated after exclusion of special items. Our results are robust to both of these alternative unexpected earnings measures.

### 3.6 Dependent Variables

In order to test the relationship between disagreement and returns we use the following method to construct a measure of abnormal returns that controls for well known risk factors. The market excess return  $R_m - r_f$  is the value-weighted return on all NYSE, AMEX, and NASDAQ stocks (from CRSP) minus the one-month Treasury bill rate (from Ibbotson Associates). We construct abnormal returns from the Fama-French four factor model with rolling beta estimation. For each stock in the dataset, we first estimate expected returns:

$$r_{it} = \alpha + \beta_{1i} * (R_m - r_f) + \beta_{2i}HML + \beta_{3i}SMB + \beta_{4i}MOM + e_{it} \quad (2.12)$$

using 150 days of lagged data (event days -195 to -46) to estimate the parameters  $\alpha$  and  $\beta_{1-4}$ . The event period abnormal return  $AR$  is the prediction error from this regression and is estimated as follows:

$$AR_{it} = R_{it} - (\hat{\alpha}_{jt} + \hat{\beta}_{1jt}R_{mt} + \hat{\beta}_{2jt}HML_t + \hat{\beta}_{3jt}SMB_t + \beta_{4i}MOM_t) \quad (2.13)$$

Where  $R_{mt}$ ,  $HML_t$ ,  $SMB_t$ , and  $MOM_t$  are the returns to the market, the book value, the size factor and momentum portfolios respectively on day  $t$ , and  $R_{jt}$  is the return to the  $j^{th}$  firm on day  $t$ . Finally, the abnormal returns are cumulated over the various event windows to create our measures of cumulative abnormal returns ( $CAR$ ).

Panel A of Table 2.1 presents facts about our Twitter data set that allow the reader to infer the how our results generalize to broader populations. The sample stocks are varied in size, with a median market cap of \$1.375 million. Our sample stocks are somewhat larger than the sample in Berkman et al. (2009) probably due to the fact that to enter into our sample a stock must capture the attention of both Twitter followers and the business press.<sup>8</sup> Reflecting their size, our sample stock's  $AGE$  and turnover ( $TURN$ ) are also larger

---

<sup>8</sup>We impose the constraint that analyst coverage must be sufficient to calculate dispersion. Berkman et al. (2009) find that imposing this constraint removes approximately sixty percent of their sample. Berkman et al (2009) do not report the sample statistics for their constrained sample, but we conjecture that the market capitalization numbers in their constrained sample would be closer to those in our sample.

than the sample in Berkman et al. (2009), but the mean income (*INCVOL*) and return (*RETVOL*) volatilities are comparable. Table 2.1 also separately presents the same summary information for low divergence of opinion stocks (*DIVOP* = 0) in Panel B and high divergence of opinion stocks (*DIVOP* = 1) in Panel C. The main difference between these two samples appears to be in firm size, and consequently the levels of analyst coverage. High divergence of opinion stocks are, on average, larger with greater levels of analyst coverage.<sup>9</sup> Note that several of the financial market based proxies for differences of opinion, such as *SUV*, *INCVOL* and *RETVOL* are smaller for the high divergence of opinion stocks than they are for the low divergence of opinion stocks. This fact suggests that our measure of divergence of opinion derived from social network posts is capturing different information than the standard proxies.

To identify these differences, Table 2.2 presents the correlations between the dispersion of opinion measures in Berkman et al. (2009), Garfinkel and Sokobin's (2006) standardized unexpected volume and our divergence of opinion measure *DIVOP*. What is fascinating is that although these three sets of measures are all purported to measure divergence of opinion, they are remarkably uncorrelated. *DIVOP* has a low correlation with standardized unexpected volume (*SUV*) of 0.05 and no correlation at all with analyst forecast dispersion (*DISP*). Further, the correlations between *DIVOP* and the other Berkman et al. (2009) measures *AGE*, *RETVOL*, *TURN* and *INCVOL* are quite low as well ranging from -0.09 to 0.17. Neither are the correlations between the Berkman et al. (2009) dispersion measures and Garfinkel and Sokobin's (2006) standardized unexpected volume large, ranging between -0.13 to 0.02. Clearly, *DIVOP* as a measure of divergence of opinion between investors and the the public information in news is capturing something different than the standard proxies for divergence of opinion. We turn next to determining whether *DIVOP* as a measure of divergence of opinion captures similar effects to the earnings announcement results in Berkman et al. (2009) and Garfinkel and Sokobin (2006).

---

<sup>9</sup>In unreported results we control for market capitalization and analyst coverage in all the subsequent results. We find that including these controls does not affect our conclusions.

## 4 Results

### 4.1 Divergence of Opinion and Returns

Miller (1977) gives rise to Hypothesis 1a which predicts that investor disagreement about earnings will cause the firm to be overvalued because its price will be set by the most optimistic investors. However, once a firm releases their earnings disagreement should be eliminated and the future stock price will reflect the lower mean valuation of all investors. Empirically, this implies that the variable *DIVOP* should be negatively related to future abnormal returns after an earnings announcement. Furthermore, Hypothesis 1b predicts that if investors agree about a firm's future prospects then the firm's price will reflect the mean expectations of all investors and thus be fairly priced. This implies that the variable *AGREE* should not have any significant relationship to future abnormal returns.

Table 2.3 estimates the effect of *DIVOP* on abnormal returns around earnings announcements. We regress the cumulative abnormal returns subsequent to an earnings announcement on *DIVOP* while controlling for the earnings surprise using standardized unexpected earnings. We also include the *IMPACT* and *NEWS* variables in the regression to test for incremental affects of the strength of the sentiment and the level of the news sentiment on returns. The coefficient on *DIVOP* is significantly negative at window lengths of 2, 10, 20 and 60 trading days. There is a significant negative abnormal return association with Twitter opinions diverging from the tone of public news sources. In Table 2.3, our control variables are not consistently significant, although *NEWS* and *IMPACT* are negatively related to returns at the 2 day window.<sup>10</sup>

The divergence of opinion results are illustrated graphically in Figure 2.2, which compares the post-earnings announcement drift in the quarter after earnings announcement for stocks classified where Twitter either agrees with the public news or has a divergent opinion. In the Twitter *AGREE* stocks there is almost no drift in returns over the next quarter. While in

---

<sup>10</sup>Any cross-sectional effect of *IMPACT* may be muted by the fact that our research design includes only high-impact posters.

the *DIVOP* stocks there is a material downward drift of over four percent during the post-earnings announcement quarter. This drift is consistent with the attenuation of differences of opinion after the earnings release. Later, we will take advantage of the continuous changes in our measure to contrast the drift between stocks whose Twitter opinions continue to diverge with public news announcements against those opinions that agree. In this analysis, we find support for the conjecture that a change in investor disagreement is responsible for the drift phenomenon.

The results in Table 2.3 show that our measure of divergence of opinion, *DIVOP*, is significantly negatively related to returns over the 2 day window, a result consistent with Berkman et al. (2009). However, at longer windows divergence of opinion remains negatively related to returns, a result that contradicts the finds of Garfinkel and Sokobin (2006). In Tables 2.4 and 2.5 we introduce the dispersion proxy variables from Berkman et al. (2009) and Garfinkel and Sokobin (2006) to test whether the effect of *DIVOP* remains significant in the presence of these financial market proxies for differences in opinion. Table 2.4 examines the robustness of our results in the presence of the Berkman et al. (2009) divergence of opinion proxy variables. In Table 2.4 we regress cumulative abnormal returns twenty trading days after the earnings announcement against the set of variables including *DIVOP* as well as sequentially including the difference of opinion variables from Berkman et al. (2009). We find that none of the proxies Berkman et al. (2009) use as measures of divergence of opinion have a material impact on the *DIVOP* coefficient. Given the correlation evidence in Table 2.2, this is not surprising. In the presence of *DIVOP*, the set of Berkman et al. (2009) difference of opinion explanatory variables do not predict future negative abnormal performance over this time frame.<sup>11</sup> Of the five divergence of opinion variables: *INCVOL*, *RETVOL*, *DISP*, *AGE*, and *TURN* the coefficients of *INCVOL*, *RETVOL* and *DISP* are significantly positive, while the coefficient of *AGE* is negative but insignificant.

---

<sup>11</sup>We note that Berkman et al. (2009) only claim their dispersion proxies are related to short-window (3-day) returns around earnings announcements.

The significantly positive effect of income volatility, return volatility, and dispersion is reminiscent of the Garfinkel and Sokobin (2009) finding of a positive association between divergence of opinion proxies and post-earnings announcement returns. As both Garfinkel and Sokobin (2006) and Garfinkel (2010) suggest proxies for unexpected volume better capture divergence of opinion, we replicate the tests in Table 2.5 using their standard unexpected volume (*SUV*) measure as a control variable. We also include two other variables from Garfinkel and Sokobin (2006) market-adjusted turnover and change in market-adjusted turnover in some specifications.

These regression results are presented in Table 2.5. However, none of the volume measures, including standardized unexpected volume has a significantly positive effect on returns as predicted in Garfinkel (2010). The *DIVOP* coefficient, derived from investors posted opinions on Twitter continue to have negative and significant effects on abnormal returns.

## 4.2 Divergence of Opinion and Earnings Surprises

Clearly, returns surrounding earnings announcements depend on the content of the announcement itself. We measure the information content using *SUE*, standardized unexpected earnings defined in Equation (10). We thus, control for the affect that unexpected positive or negative news events have on the returns across *DIVOP* groups. We examine this question in Figure 2.4 where we split the sample into thirds and separately examine the differences between *DIVOP* and *AGREE* portfolios in the highest and lowest *SUE* subsamples.

The graphs in Figure 2.4 show that, as expected, the stock’s initial abnormal return reaction is strongly positive for the most positive *SUE* tercile and strongly negative for the most negative *SUE* tercile. Immediately after the initial stock price reaction the impact of divergence of opinion can be seen in both *SUE* subsamples. For both high and low *SUE* stocks, the stocks with Twitter opinions that diverge from the public news (*DIVOP* = 1) drift consistently down in price, while *DIVOP* = 0 stocks do not drift appreciably after their initial reaction. Notably we observe a significant *negative* drift for divergence of opinion

stocks with positive  $SUE$ 's, rather than the positive drift often associated with high  $SUE$  stocks due to momentum or the slow diffusion of information. The sign and magnitude of the earnings surprises do not affect our findings on the efficacy of  $DIVOP$  as a measure of divergence of opinion.

### 4.3 Changing Opinions

In the Harris and Raviv (1993) model changes of opinion come naturally and, in fact, are necessary to facilitate trading. This model predicts that trading will only occur when the two groups of speculators switch sides or change opinions about the future value of a firm. In the Miller (1977) model, the information release convinces optimistic investors that their opinions are (on average) in error and the stock price drifts down to the true valuation. Thus, if we see divergence of opinion before an earnings announcement then when the official earnings number is released we should see a switch from disagreement to agreement among investors. Berkman et al. (2009) test Miller's contention by examining the association between changes in their set of measures. As the author's note there are difficulties in testing Miller's theory directly by using their set of proxy variables. Indeed,  $\Delta AGE$  is completely predictable *ex ante* and the reported change in income would at best be a minor increment to the 20 quarter series used to calculate  $INCVOL$ . Instead,  $\Delta RETVOL$ ,  $\Delta DISP$  and  $\Delta TURN$  are used to measure the change in divergence of opinion.

As Berkman et al. (2009) note, the use of these variable as measures of the change in differences of opinion do pose other difficulties. One problem is that the information in the earnings announcement itself could affect the levels of stock volatility, turnover and analyst's forecasts. Berkman et al. (2009) estimate excess stock returns around the earnings announcement on their proxies for changes in differences of opinion. They find mixed results between changes in differences opinion and returns around the earnings announcement:  $\Delta RETVOL$  and  $\Delta DISP$  are significantly negatively related to stock returns, while  $\Delta TURN$  is positively related to stock returns. Aside from the reverse causality problem

that could be caused by earnings announcement information, a further ambiguity in the interpretation of these results arises from the use of the short window abnormal return centered on the event day. A strict interpretation of Miller (1977) is that day -1 abnormal returns should be positively related to increases in differences of opinion, as the optimists still dominate the pricing process before the release of the earnings information. Conversely, after the earnings release, differences of opinion should decrease. Yet increases in differences of opinion are certainly possible, as modeled by Holthausen and Verrecchia (1990), Harris and Raviv (1993) and Kandel and Pearson (1995) among others.

We examine abnormal returns in both the pre-event (-11, -1) and post-event (+1, +60) periods to examine the movement of returns both before and after the earnings announcement rather than grouping them in a single short window announcement return. Figure 2.5 plots the abnormal stock returns from event day -11 to event day +60 surrounding the earnings announcement. After the earnings announcement stocks with prior divergence illustrate the lower stock returns that we previously illustrated in Figure 2.2. Prior to the announcement both groups have runups in stock price, but the runup for *DIVOP* stocks is larger than that of *AGREE* stocks. Note that the returns immediately surrounding the announcement for *DIVOP* stocks exhibit distinct characteristics immediately surrounding the event; increasing immediately before the event and decreasing immediately after the event. These results are broadly consistent with Miller (1977) as returns increase for divergent stocks prior to the announcement and decrease notably after the announcement.

We contend that changes in *AGREE* and *DIVOP* around the earnings announcement can directly test the hypothesis that resolution of disagreement should cause the optimists to correct their positions and result in lower stock returns. We test this contention by contrasting *DIVOP* and *AGREE* before and after the earnings announcement. Figure 2.6 presents a simple matrix that determines whether differences in opinions have converged or diverged around the earnings announcement. In the -11 to -1 pre-event period Twitter investors will either agree (*AGREE*) or diverge (*DIVOP*) from the public news. Similarly, we calculate the same variables in the +1 to +11 post-event period. Thus there are four

possible combinations of pre-event and post-event opinions. Twitter investors can agree or diverge from public news in both periods. In these cases we make no strong opinions about the returns surrounding earnings announcements. If Twitter investors change their opinions then they can go from diverging to agreement, as in Miller (1977), or from agreement to disagreement. The former case contains our main prediction: that stock returns should fall should the opinions of investors change as a result of the earnings announcement from divergence pre-announcement to agreement afterwards. In this case, Miller's (1977) theory would predict that stock prices should generally fall with the decrease in differences of opinion. In the case where agreement turns to disagreement we should see a stock price increase.

Figure 2.7 presents the post-event abnormal returns for the two groups of stocks where Twitter investor opinions change around the earnings announcement. The results are striking, the post-event returns of the two groups diverge dramatically. Stocks that go from prior agreement to post announcement disagreement have cumulative abnormal returns of 2.19 percent ( $t$ -statistic = 3.52) from event day +1 to +11. Stocks that go from prior disagreement to post-announcement agreement decline by -1.63 percent ( $t$ -statistic = -3.06) in the same post-event period. The returns to the group where pre-event disagreement switches to agreement are consistent with optimistic traders reevaluating their valuations down. While the positive post-event agreement to disagreement returns suggest the opposite effect is occurring.

## 4.4 Divergence of Opinion and Volume

Divergence of opinion is a common theoretical explanation for volume (see, for example Kim and Verrecchia (1994), Kandel and Pearson (1995)). Using the simple matrix of changes in divergence of opinion in Figure 2.6, it is straightforward to test how changes in divergence of opinion affects volume in our sample. The theoretical motivation for testing the changes in divergence of opinion comes from Holthausen and Verrecchia's (1990) concept of investor *consensus* around earnings announcements. In their paper, earnings announcement can

effect traders by increasing their level of informedness or changing the degree of consensus between investors. They predict that greater informedness or reduced consensus increase trading volume. Our concept of divergence of opinion between investors on the social network Twitter and public news releases is clearly a measure of consensus. As outlined in Hypotheses 2a and 2b, we first expect divergence of opinion to be positively related to volume, but this is just a test of our measure and nothing new in the well recognized connection between divergence and volume. What is new is that we can determine the volume effects of changes in consensus. Hypothesis 2b predicts that earnings announcements that lead to increased consensus (disagreement to agreement) should be associated with lower volume than earnings announcements that fail to resolve prior disagreement, or even change the investor's opinions from agreement to disagreement.

To test these hypotheses, we calculate cumulative abnormal volume by constructing an unexpected volume measure for every stock every day in the -30 to 60 event period. The methodology of the volume event study is very similar to a standard return event study. To calculate abnormal trading volume we follow the methodology of Campbell and Wasley (1996). We define  $V_{it} = \frac{volume_{it}}{S_{it}}$ , where  $S$  is the total number of shares outstanding. Next we take the log of  $V_{it}$  and estimate a market model of volume over the (-45,-245) window and define abnormal volume  $AV$  as  $AV_{it} = V_{it} - \alpha_i + \beta_i V_{mt}$ , where  $V_{mt}$  is aggregate market volume for all NYSE, AMEX and NASDAQ stocks. An estimate of abnormal volume is made for everyday in our event window, (-30,60) . Finally, we take the sum of abnormal volume everyday in our sample to plot cumulative abnormal volume.

Once abnormal volume is calculated, we divide the sample into the four different cases outlined in Figure 2.6. Investors may either (i) move from prior agreement to post-announcement disagreement, (ii) maintain their agreement over both periods, (iii) move prior disagreement to agreement, or (iv) continue to disagree with the public news even after the earnings announcement information. The cumulative abnormal volume patterns for these four categories is plotted in Figure 2.8. In the 30 trading days prior to the earnings announcement there are no dramatic differences in volume across the four categories nor any large deviations in

cumulative volume. Two categories have below average cumulative volume prior to the earnings announcement, those associated with prior agreement, but the two prior disagreement categories have above average volume. Abnormal volume increases for all four categories at the time of the earnings announcement, but the dramatically different patterns of abnormal volume appear to vary in a consistent way in the post-earnings announcement period.

In the post-earnings announcement period, the *AA* agreement group has the lowest level of abnormal volume, finishing with negative abnormal volume on day +60. The *DD* disagreement group has the highest abnormal volume throughout ending the quarter with almost 1000% cumulative volume increase. Divergence of opinion tends to be associated with higher volume in all periods around the earnings announcement. What is most striking is that the patterns of abnormal volume changes after the earnings announcement coincident with changes in the divergence of opinion between social network investors' opinions and public news releases. As investors' opinions move from agreement to disagreement (*AD* group) around the earnings announcement period, abnormal volume spikes up and continues to increase until the cumulative volume of this group passes the level of the disagreement to agreement group.

The pattern is clear: social network divergence of opinion is correlated with trading volume. Many earnings announcements fail to increase consensus, but in fact can increase investors divergence of opinion, and these changes in divergence are associated with increased volume. Similarly, earnings announcements often produce no disagreement and may even increase agreement by changing investor's opinions from disagreement with the public information to agreement (*DA* group). It is this case that corresponds best to earnings announcements reducing consensus and in this group abnormal volume is indeed attenuated, as it is flat for most of the post-earnings announcement period.

We also note that all earnings announcements tend to be associated with increases in trading volume at the time of the announcement; this makes sense as earnings news is often significant for valuation purposes. We conduct *t*-tests for differences in the level of

abnormal volume at the time of the earnings announcement (event days -1 to 0) sorting on whether the announcement increases or decreases consensus. In cases where investors are in disagreement after the earnings announcement, and thus the Holthausen and Verrecchia (1990) measure of consensus decreases, volume is higher in the -1 to 0 period ( $t$ -statistic = 4.74). The point is that the Garfinkel and Sokobin (2006) finding that  $SUV$  at the time of the earnings announcement is positively related to future returns is based on Varian's (1985) conjecture that asset prices will be lower when investor's opinions are more disperse. Implicit in their finding is the assumption that the earnings announcement resolves uncertainty and increases consensus. Our tests have found that abnormal volume at the time of earnings announcements is higher when there is disagreement after the earnings announcement. It is not that the time frame used for standardized unexpected volume in Garfinkel and Sokobin (2006) and Garfinkel (2010) is an invalid test, rather the point is that  $SUV$  seems to be a relatively low power test, as earnings announcement period volume is not necessarily directly related to resolution of disagreement. For many tests, the more powerful alternative could be to look at differences between groups in the post-earnings announcement period where the relation between divergence and abnormal volume is more dramatic and thus, potentially more powerful.

## 5 What causes divergence of opinion?

In this section we investigate the linguistic and tonal features of pre-earnings announcement press releases. We use the textual analysis software DICTION to summarize the average tonal features of articles that are associated with agreement or disagreement. The DICTION software, or other word frequency textual analysis software such as General Inquirer (GI), have been used in many previous studies to measure tonal or features of financial texts, see Engelberg (2008), Tetlock (2007), Tetlock and Saar-Tsechansky (2008), Demers and Vega (2010), Davis, Piger, and Sedor (2008).<sup>12</sup>

---

<sup>12</sup>See the DICTION website for a full list of scholarly articles that have used DICTION <http://www.dictionsoftware.com>

The DICTION software package centers around thirty one predefined dictionaries which measure various linguistic features of a text, these dictionaries containing more than 10,000 unique words that are used to define the content of a document. The DICTION analysis that we use is based on normative values that are calibrated using a set of firm-specific public relations articles. The calibration process gives us the average linguistic properties for firm-specific public relations documents. This subset of financial news is the best fit to our set of firm specific press releases prior to earnings announcements. Loughran and McDonald (2011) use a set of 10-K reports to show that the Harvard Dictionary of linguistic features is not an appropriate benchmark for the analysis of financial press releases. They show that 73% of the negative words in the Harvard Dictionary are not typically interpreted as negative in a financial press release. We control for this issue because our linguistic priors are explicitly determined by a set of firm-specific public relations articles which measure the central tendencies of this type of financial news and then compare it to the features of our documents.

Thirty one dictionaries define the following list of linguistic characteristics: tenacity, leveling, collectives, numerical terms, ambivalence, self-reference, praise, satisfaction, inspiration, blame, hardship, denial, aggression, accomplishment, communication, motion, cognitive terms, passivity, familiarity, spatial awareness, temporal awareness, present concern, human interest, concreteness, past concern, centrality, cooperation, rapport, diversity, exclusion, and liberation. These characteristics are then used to calculate five master variables that summarize the overall tone of a document and provide an overall robust understanding of its contents. The master variables are constructed by combining additive and subtractive linguistic characteristics so that the master variables measure the net tonal features of a document. The first master variable is Certainty which is defined as  $Certainty = (Tenacity + Leveling\ terms + Collective\ terms + Insistence) - (Numerical\ terms + Ambivalence + Self\ Reference + Variety)$ . Certainty is a measure of overall resoluteness and inflexibility, as well as, a tendency to speak from a position of unquestioned authority. The next master variable is Optimism, which is defined as,  $Optimism =$

$(Praise + Satisfaction + Inspiration) - (Blame + Hardship + Denial)$ . Optimism measures how a text endorses a group or person by highlighting their positive features and attributes. The third master variable is Realism which measures a document's tendency to use language that describes tangible and immediate matters that affect peoples' lives every day. Realism is defined as,  $Realism = (Familiarity + Spatial\ awareness + Temporal\ awareness + Present\ concern + Human\ interest + Concreteness) - (Past\ concern + Complexity)$ . The fourth master variable is Activity, which is defines as,  $Activity = (Aggression + Accomplishment + Communication + Motion) - (Cognitive\ terms + Passivity + Embellishment)$ . Activity is used to measure language featuring change and the implementation of new ideas. The final master variable is Commonality, which is defined as,  $Commonality = (Centrality + Cooperation + Rapport) - (Diversity + Exclusion + Liberation)$ . Commonality measure the use of language that highlights agreed upon values.

In addition to the five master variables, DICTION also creates four ratio variables that measure characteristics of a document that cannot be easily captured by word frequencies from a dictionary. These ratio variables are: Complexity (characters/word), Variety (different words/total words), Insistence (heavily used words  $\times$  total occurrences/10) and Embellishment (adjectives/verbs).

We ask the question: What are the linguistic properties of news releases that are associated with divergence of opinions or agreement? To answer this we separate all of the news articles in our sample into two groups based on whether there was agreement (*AGREE*) or divergence of opinion (*DIVOP*) about the article in the pre-earnings announcement period. This sorting procedure allows us to isolate the linguistic features of a press release that contribute to disagreement among the public. Once articles are sorted into two groups we run a separate analysis of the linguistic features of each group; Table 2.9 reports the output of this analysis. Our analysis compares the features of our news articles to the average linguistic properties, the average values of all of firm-specific public relations documents. If our document contains a large or small number of words from a certain characteristic then its score will be above or below the average. A range of the normal frequencies is available for

this type of document and the range gives us a sense of where our documents differ from the average characteristics of press releases.

The first interesting feature that we see from Table 2.9 is that the master variables for the articles that are associated with agreement are almost all within the normal range for financial press releases. The only unique feature of this set of press releases is that they tend to exhibit a larger variety of words than a normal document. This shows us that investors that read news tend to agree with the sentiment of news if the tonal features in the news are perceived as “normal”.

Next, if we turn to the set of articles that investors disagree with, we see that they have low levels of Optimism and Certainty. This shows us that the tone of these press releases tend to be less resolute and show less authority than normal. At the same time, the press releases tend to have fewer words that highlight the positive virtues of a company. Finally, from the list of calculated variables, we see that press releases that are associated with investor disagreement are more insistent in their language. The text of the press releases tend to use the same words repeatedly in an attempt to persuade the reader. Thus, press releases that induce investor disagreement tend to heavily reuse words, and the tone of the text lacks praise for the company and makes statements that are not authoritative or resolute. These results show us that the tone and content of a corporate press release is associated with perception the investing public has about a company. Tonal and language constructs in press releases affects the level of disagreement between investors and the company. Reducing disagreement is clearly an important goal of corporate investor relations departments because we show that disagreement between Twitter opinions and press releases can have valuation consequences.

## 6 Conclusion

In this paper, we use a unique set of data collected from the social networking site Twitter to measure the divergence of investors’ opinions. We examine the impact of divergence of

opinions on the price and volume reactions to earnings announcements. This dataset allows us to directly measure investor sentiment about earnings through the use of NLP classification techniques. To our knowledge, this is the first study that has been able to explicitly measure divergence of opinion among investors rather than relying on financial market data to construct proxy variables. By measuring the divergence of opinion between press releases and social network communication we find support for the theoretical predictions of Miller (1977) and Holthausen and Verrecchia (1990). We show that divergence of opinion prior to the earnings announcement is generally uncorrelated with financial market proxies for divergence. This prior-period divergence of opinion leads to significant negative abnormal returns in the post-announcement period while agreement prior to the announcement produces no significant abnormal returns. The significant negative abnormal returns associated with divergence of opinion from social network investors persists even when using financial market dispersion measures or standardized unexpected earnings as control variables.

We extend our tests to measure the changes in divergence of opinion from the period prior to the earnings announcement to the post-earnings announcement period. We find that earnings announcements where opinions change from agreement prior to the announcement to divergence after produce positive abnormal returns that markedly differ from the pattern of negative abnormal returns for the sample where divergence of opinion changes from disagreement before to agreement after the earnings announcements. This test illustrates that while earnings announcement can sometimes increase consensus among investors, they can also produce the opposite effect, a divergence of opinion after the announcement is made. These changes in consensus naturally lead to tests of abnormal volume as a function of whether investor consensus increases or decreases after the earnings announcement (Holthausen and Verrecchia (1990)). We find that abnormal volume increases when divergence of opinion is high after the earnings announcement, even when investors were in agreement with publicly available news prior to the announcement. Likewise, when the earnings announcement produces agreement among investors, abnormal volume is muted in the post-earnings announcement period.

Finally, we examine the relation between the linguistic characteristic of press releases and divergence of opinions. We find that the tone of firms' public relations can affect investors' perception of the company. Specifically, a lack of optimism and certainty coupled with repetitive language tends to induce investors to disagree with the sentiment of press releases. We have shown that divergence of opinions about pre-earnings news can substantially impact firms' future valuations. Thus, firms should be continually cognizant of the real time perceptions of the investing public because changes in investors' perceptions can have a material affect on a firm's valuation.

1. Abarbanell, J., W. Lanen and R. Verrecchia, 1995. Analyst' forecasts as proxies for investor beliefs in empirical research. *Journal of Accounting and Economics*, 20, 31-60.
2. Ajinkya, B.; R. Atiase; and M. Gift, 1991. Volume of trading and the dispersion in financial analysts' earnings forecasts, *The Accounting Review*, 66, 389-401.
3. Almazan, A., K. C. Brown, M. Carlson, and D. Chapman, 2004. Why constrain your mutual fund manager? *Journal of Financial Economics*, 73(2):289-321.
4. Berkman, H., V. Dimitrov, P. Jain, P. Koch and S. Tice, 2009. Sell on the news: Differences of opinion, short-sales constraints and returns around earnings announcements. *Journal of Financial Economics*, 92, 376-389.
5. Bushee, B., J. Core and W. Guay, and Sophia J. W. Hamm, 2010, The Role of the Business Press as an Information Intermediary, *Journal of Accounting Research* 48, 1-19.
6. Campbell, C. and C. Wasley, 1996, Measuring abnormal daily trading volume for samples of NYSE/ASE and NASDAQ securities using parametric and nonparametric test statistics. *Review of Quantitative Finance and Accounting* 6, 309-326.
7. Chan, W., 2003. Stock Price Reaction to news and no-news: drift and reversals after headlines, *Journal of Financial Economics*, 70, 223-260.
8. Davis, A., J. Piger, and L. Sedor, 2010, Beyond the Numbers: An analysis of optimistic and pessimistic language in earnings press releases, Working Paper, University of Oregon and University of Washington.
9. Deither, K., C. Malloy and A. Scherbina, 2002. Differences of Opinion and the cross-section of stock returns, *Journal of Finance*, 57(5), 2113-2141.
10. Demers, E., and C. Vega, 2010, Soft information in earnings announcements: News or noise?, Working paper, INSEAD.
11. DeMarzo, P., D. Vayanos and J., Zwiebel, 2003. Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics*, 118(3), 909-968.
12. Doukas, J.; C. Kim; and C. Pantzalis, 2006. Divergence of opinion and equity returns, *Journal of Financial and Quantitative Analysis*, 41, 573-606.
13. Driscoll J. and A. Kraay, 1998. Consistent covariance estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80, 549-560.
14. Engelberg, J., 2008. Costly Information Processing: Evidence from Earnings Announcements. AFA 2009 San Francisco Meetings Paper. Available at SSRN: <http://ssrn.com/abstract>
15. Fang, L. and J. Peress, 2009. Media coverage and the cross-section of stock returns. *Journal of Finance*, 64(5), 2023-2052.

16. Garfinkel, J, and J. Sokobin, 2006. Volume, opinion divergence and returns: A study of post earnings announcement drift, *Journal of Accounting Research*, 44, 85-112.
17. Garfinkel, J., 2010. Measuring Investors' Opinion Divergence, *Journal of Accounting Research*, 47(5), 1317-1348.
18. Harris, M., and A. Raviv, 1993. Differences of opinion make a horse race. *Review of Financial Studies*, 6, 473-506.
19. Holthausen R. and R. Verrecchia, 1990. The effect of informedness and consensus on price and volume behavior. *The Accounting Review*, 65(1), 191-208.
20. Hong, H., and J. Stein, 2006. Disagreement and the stock market. Working Paper, Princeton University.
21. Karpoff, J., 1987. The relationship between price changes and trading volume: A survey. *Journal of Finance and Quantitative Analysis*, 22, 109-26.
22. Kandel E., and N. Pearson, 1995. Differential interpretations of information and trade in speculative markets. *Journal of Political Economy*, 103(4), 831-872.
23. Kim, O. and R. Verrecchia, 1994. Market liquidity and volume around earnings announcements. *Journal of Accounting and Economics*, 17, 41-67.
24. Lamont, O. and J. Stein, 2004. Aggregate short interest and market valuation, *American Economic Review*, 94(2), 29-32.
25. Loughran, T., and B. McDonald, 2011, When is a Liability Not a Liability?, *Journal of Finance* 66, 35-65.
26. Miller, E., 1977. Risk, Uncertainty, and Divergence of Opinion *Journal of Finance* 32, 1151-1168.
27. Pang, B., L. Lee and S. Vaithyanathan, 2002. Thumbs up?: sentiment classification using machine learning techniques. EMNLP '02: Proceedings of the ACL-02 conference on Empirical methods in natural language processing.
28. Sorescu, S., R. Boehme and B. Danielson, 2006. Short sale constraints, dispersion of opinion and overvaluation. *Journal of Financial and Quantitative Analysis*, 41(2), 455-487.
29. Tetlock, P., 2007. Giving Content to Investor Sentiment: The Role of Media in the Stock Market *Journal of Finance*, 62(3), 1139-1168.
30. Tetlock, P., Saar-Tsechansky, M. and S. Macskassy, 2008. More than words: Quantifying language to measure firms' fundamentals. *Journal of Finance*, 63 (3), 1437-1467.
31. Varian, H., 1985. Divergence of opinion in complete markets: A note, *Journal of Finance* 40, 309-317.

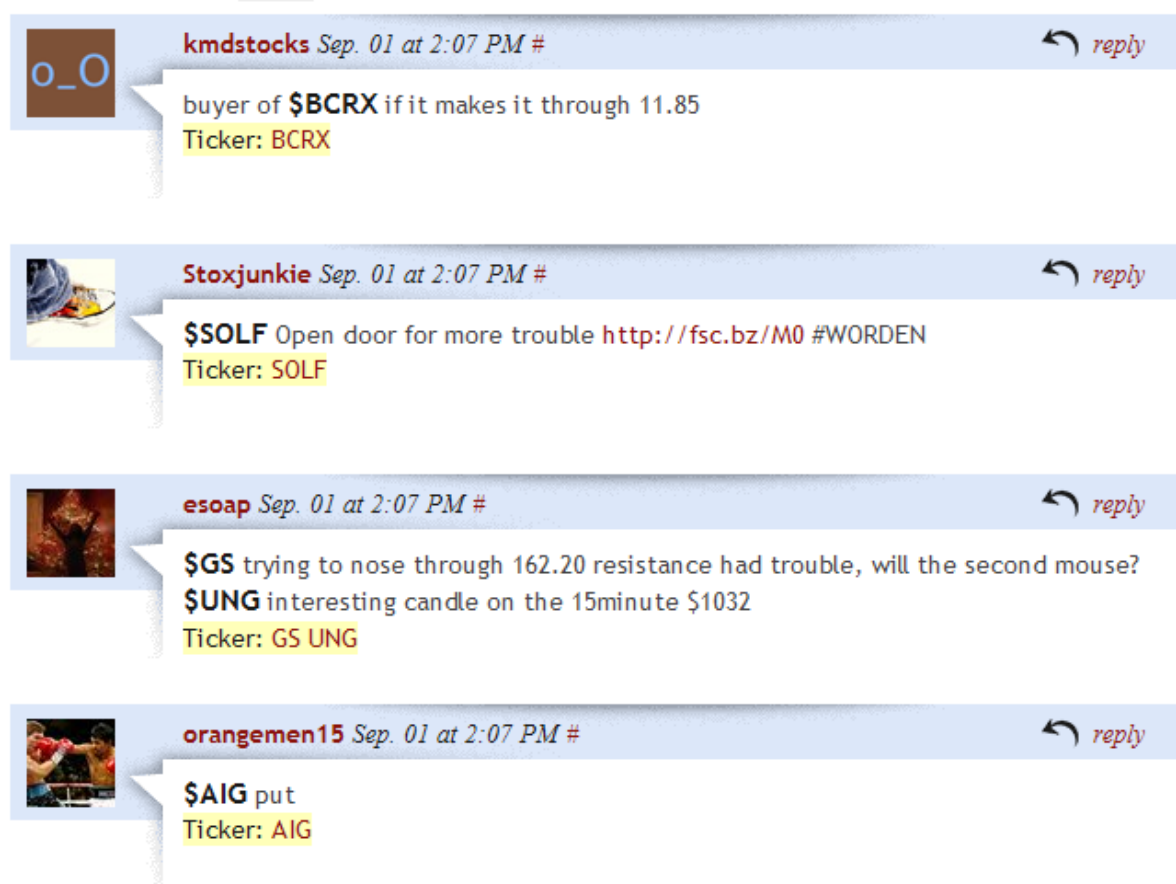


Figure 2.1: www.Stocktwits.com

Figure 2.2: Divergence of Opinions Event Study

This figure depicts post earning announcement cumulative abnormal returns. Abnormal returns are calculated from the Fama-French 3 factor model, and are cumulated over the (1,60) event window. Firms are separated into two groups based on the value of the variable *DIVOP*, which is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the (-1,-10) event window.




Figure 2.3: Divergence of Opinions Event Study top and bottom SUE terciles

This figure depicts post earning announcement cumulative abnormal returns. Abnormal returns are calculated from the Fama-French 3 factor model, and are cumulated over the (1,60) event window. For this study, we first sort firms into terciles based on *SUE*. Then, within the *SUE* terciles, we sort firms into two groups based on the value of the variable *DIVOP*, which is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the (-1,-10) event window. We then plot the results for the high and low *SUE* groups. This process controls for cross sectional differences in unexpected earnings.




Figure 2.4: Divergence of Opinions Event Study top and bottom SUE terciles

This figure depicts post earning announcement cumulative abnormal returns. Abnormal returns are calculated from the Fama-French 3 factor model, and are cumulated over the  $(-11,60)$  event window. For this study, we first sort firms into terciles based on *SUE*. Then, within the *SUE* terciles, we sort firms into two groups based on the value of the variable *DIVOP*, which is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the  $(-1,-10)$  event window. We then plot the results for the high and low *SUE* groups. This process controls for cross sectional differences in unexpected earnings.




Figure 2.5: Divergence of Opinions Event Study

This figure depicts post earning announcement cumulative abnormal returns. Abnormal returns are calculated from the Fama-French 3 factor model, and are cumulated over the  $(-11,60)$  event window. For this study, we first sort firms into two groups based on the value of the variable *DIVOP*, which is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the  $(-1,-10)$  event window. We then plot the results for the groups.




Figure 2.6: **Predicted signs of CARs conditional on opinions**

This figure shows the four possibilities of changes in opinions around earnings. AA is agreement before and after an announcement, and AD is agreement before and disagreement after the announcement. DD is disagreement before and after an announcement, and DA is disagreement before and agreement after. The bottom part of the table gives the predicted sign of abnormal returns after a particular set of opinions.

| Matrix of Changes of Opinions                    |                                         |                    |                   |
|--------------------------------------------------|-----------------------------------------|--------------------|-------------------|
|                                                  |                                         | 1st period         |                   |
|                                                  |                                         | Agree              | Divergence        |
| 2nd Period                                       | Agree                                   | AA                 | DA                |
|                                                  | Divergence                              | AD                 | DD                |
| Predicted Signs of CARs From Changes of Opinions |                                         |                    |                   |
|                                                  |                                         | Pre Event (-11,-1) | Post Event (1,11) |
| DD                                               | <i>Divop</i> $\rightarrow$ <i>Divop</i> | + or 0             | + or 0            |
| DA                                               | <i>Divop</i> $\rightarrow$ <i>Agree</i> | + or 0             | -                 |
| AA                                               | <i>Agree</i> $\rightarrow$ <i>Agree</i> | N/A                | N/A               |
| AD                                               | <i>Agree</i> $\rightarrow$ <i>Divop</i> | N/A                | +                 |

Figure 2.7: Changes of opinions around earnings

This figure plots the Fama French 3 Factor CARs for two groups of changes in opinions around earnings over the (1,11) event window. The first group is AD, which is agreement before and disagreement after the announcement. The second group is DA, which is disagreement before and agreement after an announcement.




Figure 2.8: Changes of opinions around earnings: Volume Test

This figure plots the Fama French 3 Factor CARs for two groups of changes in opinions around earnings over the (1,11) event window. The first group is AD, which is agreement before and disagreement after the announcement. The second group is DA, which is disagreement before and agreement after an announcement. To calculate abnormal trading volume we follow the methodology of Campbell and Wasley (1996).




Table 2.1: Summary Statistics

This table presents summary statistics of the 1,398 sample firms from the sample period of October 1st, 2008 to December 31st, 2009. Market capitalization is measured at the end of 2007. Book-to-market ratio is book equity of fiscal year ending in calendar year 2007 divided by market capitalization at the end of 2007. Analyst coverage is the number of analysts covering a firm as of September 2007. We include summary statistics of the non-standardized values of the communication variables (number of posts, number of articles) used throughout the paper. *DIVOP* is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the (-1,-11) event window. Daily market adjusted turnover, *MATO*, is the firm volume divided by shares outstanding, minus a similarly calculated measure over all NYSE/AMEX firms.  $\Delta TO$  is the change in market adjusted turnover, which is calculated as daily *MATO* minus the median 180 moving average of *MATO*. Standardized unexpected volume, *SUV*, is the scaled prediction errors from a market model-style model regression of volume on absolute valued returns. *INCVOL* is the standard deviation of seasonally differenced ratio of quarterly operating income divided by average total assets measured over the 20 quarters prior to the earnings announcement quarter. *RETVOL* is the standard deviation of of a firm's daily excess returns relative to a value weighted market index, over a 45 day period ending 10 days before the earnings announcement. *DISP*, dispersion of analysts' forecasts. *AGE* is  $\ln(1/\text{firm age})$ , where firm age is the length of time a firm has been covered by CRSP. *TURN* is average daily turnover measures over a 45 day period ending 10 days before the earnings announcement day. We report the same summary statistics after sorting firms into two groups based on their level of *DIVOP* over the entire sample period.

| Full Sample: Panel A |        |         |        |        |       |       |        |
|----------------------|--------|---------|--------|--------|-------|-------|--------|
|                      | Mean   | Std Dev | P10    | P25    | P50   | P75   | P90    |
| Market Cap. (\$M)    | 8,376  | 28,463  | 262    | 498    | 1,375 | 4,810 | 17,343 |
| Book/Market Ratio    | 0.40   | 0.39    | 0.10   | 0.20   | 0.33  | 0.55  | 0.79   |
| Analyst Coverage     | 8.66   | 6.42    | 1.00   | 4.00   | 7.00  | 12.00 | 18.00  |
| Avg. Daily Ret. (%)  | 0.21   | 0.21    | 0.01   | 0.09   | 0.17  | 0.30  | 0.45   |
| No. Posts            | 0.13   | 0.45    | 0.00   | 0.01   | 0.03  | 0.10  | 0.25   |
| No. Articles         | 0.24   | 0.24    | 0.00   | 0.03   | 0.14  | 0.40  | 0.63   |
| DIVOP                | 0.11   | 0.39    | 0.00   | 0.00   | 0.00  | 0.09  | 0.28   |
| IMPACT               | 143.68 | 800.83  | -3.28  | 0.00   | 15.76 | 77.40 | 362.39 |
| MATO                 | -0.40  | 0.74    | -1.10  | -0.91  | -0.59 | -0.13 | 0.55   |
| $\Delta TO$          | 0.01   | 0.09    | -0.03  | -0.01  | 0.01  | 0.02  | 0.03   |
| SUV                  | 2.88   | 32.16   | -19.64 | -11.28 | -2.04 | 9.47  | 26.78  |
| INCVOL               | 2.81   | 5.66    | 0.36   | 0.67   | 1.24  | 2.48  | 5.51   |
| RETVOL               | 3.42   | 1.42    | 1.86   | 2.42   | 3.18  | 4.20  | 5.20   |
| AGE                  | 24.43  | 17.34   | 7.84   | 10.86  | 16.86 | 37.86 | 57.86  |
| TURN                 | 0.92   | 0.65    | 0.14   | 0.48   | 0.81  | 1.25  | 1.81   |
| DISP                 | 0.37   | 0.88    | 0.02   | 0.04   | 0.09  | 0.27  | 0.91   |

Low Divergence of Opinions: Panel B

|                     | Mean  | Std Dev | P10    | P25    | P50   | P75   | P90   |
|---------------------|-------|---------|--------|--------|-------|-------|-------|
| Market Cap. (\$M)   | 4,273 | 24,325  | 241    | 416    | 807   | 2,135 | 6,093 |
| Book/Market Ratio   | 0.42  | 0.37    | 0.12   | 0.21   | 0.35  | 0.58  | 0.79  |
| Analyst Coverage    | 6.49  | 4.92    | 1.00   | 3.00   | 6.00  | 9.00  | 13.00 |
| Avg. Daily Ret. (%) | 0.20  | 0.23    | -0.03  | 0.08   | 0.17  | 0.31  | 0.46  |
| No. Posts           | 0.03  | 0.07    | 0.00   | 0.01   | 0.01  | 0.03  | 0.07  |
| No. Articles        | 0.05  | 0.07    | 0.00   | 0.00   | 0.03  | 0.07  | 0.11  |
| DIVOP               | 0.00  | 0.01    | 0.00   | 0.00   | 0.00  | 0.00  | 0.00  |
| IMPACT              | 38.75 | 169.61  | -2.78  | 0.00   | 6.30  | 24.35 | 84.52 |
| MATO                | -0.67 | 0.56    | -1.18  | -1.01  | -0.82 | -0.49 | 0.00  |
| $\Delta$ TO         | 0.02  | 0.12    | -0.02  | 0.00   | 0.01  | 0.02  | 0.03  |
| SUV                 | 5.64  | 38.35   | -18.95 | -10.94 | -0.55 | 11.88 | 30.34 |
| INCVOL              | 3.31  | 6.94    | 0.35   | 0.70   | 1.31  | 2.72  | 6.95  |
| RETVOL              | 3.56  | 1.35    | 1.96   | 2.60   | 3.44  | 4.30  | 5.32  |
| AGE                 | 20.66 | 15.46   | 6.85   | 9.00   | 14.87 | 26.00 | 47.86 |
| TURN                | 0.69  | 0.54    | 0.00   | 0.33   | 0.59  | 0.93  | 1.38  |
| DISP                | 0.38  | 0.86    | 0.02   | 0.04   | 0.10  | 0.32  | 1.02  |

High Divergence of Opinions: Panel C

|                     | Mean   | Std Dev | P10    | P25    | P50   | P75    | P90    |
|---------------------|--------|---------|--------|--------|-------|--------|--------|
| Market Cap. (\$M)   | 13,538 | 33,424  | 429    | 1,015  | 2,900 | 10,325 | 32,545 |
| Book/Market Ratio   | 0.37   | 0.40    | 0.09   | 0.18   | 0.30  | 0.51   | 0.77   |
| Analyst Coverage    | 11.27  | 6.96    | 3.00   | 6.00   | 10.00 | 16.00  | 21.00  |
| Avg. Daily Ret. (%) | 0.21   | 0.20    | 0.04   | 0.10   | 0.17  | 0.29   | 0.45   |
| No. Posts           | 0.23   | 0.61    | 0.01   | 0.03   | 0.08  | 0.20   | 0.43   |
| No. Articles        | 0.42   | 0.21    | 0.16   | 0.24   | 0.38  | 0.59   | 0.73   |
| DIVOP               | 0.22   | 0.53    | 0.00   | 0.00   | 0.09  | 0.23   | 0.43   |
| IMPACT              | 245.06 | 1092.82 | -3.65  | 7.86   | 38.83 | 174.82 | 668.23 |
| MATO                | -0.17  | 0.81    | -0.95  | -0.73  | -0.40 | 0.21   | 0.83   |
| $\Delta$ TO         | 0.00   | 0.06    | -0.04  | -0.01  | 0.00  | 0.01   | 0.03   |
| SUV                 | 0.45   | 25.27   | -19.86 | -11.62 | -3.25 | 7.76   | 22.94  |
| INCVOL              | 2.35   | 4.07    | 0.36   | 0.62   | 1.19  | 2.31   | 4.95   |
| RETVOL              | 3.29   | 1.47    | 1.80   | 2.29   | 3.01  | 4.04   | 5.06   |
| AGE                 | 27.96  | 18.26   | 8.86   | 12.86  | 19.86 | 45.86  | 58.86  |
| TURN                | 1.14   | 0.68    | 0.38   | 0.69   | 1.03  | 1.56   | 2.07   |
| DISP                | 0.36   | 0.89    | 0.02   | 0.03   | 0.08  | 0.24   | 0.88   |

Table 2.2: Correlation Table

This table presents the simple correlation coefficients between all of the main explanatory variables. *DIVOP* is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the (-1,-11) event window. *SUE* measures standardized unexpected earnings as the difference between the actual and expected (based on IBES median estimates) earnings divided by price. Daily market adjusted turnover, *MATO*, is the firm volume divided by shares outstanding, minus a similarly calculated measure over all NYSE/AMEX firms.  $\Delta TO$  is the change in market adjusted turnover, which is calculated as daily *MATO* minus the median 180 moving average of *MATO*. Standardized unexpected volume, *SUV*, is the scaled prediction errors from a market model-style regression of volume on absolute valued returns. *INCVOL* is the standard deviation of seasonally differenced ratio of quarterly operating income divided by average total assets measured over the 20 quarters prior to the earnings announcement quarter. *RETVOL* is the standard deviation of a firm's daily excess returns relative to a value weighted market index, over a 45 day period ending 10 days before the earnings announcement. *DISP*, dispersion of analysts' forecasts. *AGE* is  $\ln(1/\text{firm age})$ , where firm age is the length of time a firm has been covered by CRSP. *TURN* is average daily turnover measures over a 45 day period ending 10 days before the earnings announcement day.

|        | DIVOP   | SUV     | DTO     | MATO    | TURN    | AGE     | RETVOL  | INCVOL | DISP    | SUE |
|--------|---------|---------|---------|---------|---------|---------|---------|--------|---------|-----|
| DIVOP  | 1       |         |         |         |         |         |         |        |         |     |
| SUV    | 0.0493  | 1       |         |         |         |         |         |        |         |     |
| DTO    | 0.0648  | 0.5615  | 1       |         |         |         |         |        |         |     |
| MATO   | 0.1646  | 0.4552  | 0.7045  | 1       |         |         |         |        |         |     |
| TURN   | 0.1686  | 0.0157  | 0.3302  | 0.7777  | 1       |         |         |        |         |     |
| AGE    | -0.1232 | -0.0865 | -0.1658 | -0.2318 | -0.2048 | 1       |         |        |         |     |
| RETVOL | -0.0925 | -0.1294 | -0.0739 | -0.0013 | 0.1516  | 0.3311  | 1       |        |         |     |
| INCVOL | -0.019  | -0.0574 | -0.1023 | -0.0767 | -0.0056 | 0.2524  | 0.2207  | 1      |         |     |
| DISP   | -0.0009 | -0.0104 | -0.0157 | -0.0017 | 0.0182  | 0.0317  | 0.0313  | 0.178  | 1       |     |
| SUE    | -0.0012 | -0.0339 | 0.0098  | 0.0184  | 0.0157  | -0.0328 | -0.0919 | 0.0046 | -0.0057 | 1   |

Table 2.3: Divergence of Opinion and Returns

This table tests the relationship between divergence of opinion and abnormal returns. Abnormal returns are calculated from the Fama-French 3 Factor model. *CAR2*, *CAR10* and *CAR20* are cumulative abnormal returns over the (0,1), (0,9) and (0,19) event windows respectively. We include *DIVOP* as the key explanatory variable, which is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the (-1,-11) event window. We add the key control variables *IMPACT*, which measures the level of investor sentiment, *NEWS*, which measures the sentiment of firm press releases and *SUE*, which measures standardized unexpected earnings as the difference between the actual and expected (based on IBES median estimates) earnings divided by price. We also include, but do not report, the additional controls of firm market capitalization and ten days of lagged abnormal returns. All regressions are firm fixed effects panel data regressions on our 1,398 sample firms. Driscoll and Kraay (1998) robust t-stats, which control for time series and cross-sectional correlation, are reported below coefficients in parentheses \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

| VARIABLES    | (1)<br>CAR2          | (2)<br>CAR10         | (3)<br>CAR20         | (4)<br>CAR60         |
|--------------|----------------------|----------------------|----------------------|----------------------|
| DIVOP        | -0.876***<br>(-3.05) | -1.293***<br>(-3.06) | -1.568***<br>(-2.82) | -2.110***<br>(-2.76) |
| SUE          | 0.824<br>(0.74)      | 1.819<br>(1.06)      | -0.792<br>(-0.31)    | -8.855***<br>(-2.86) |
| IMPACT       | 0.252**<br>(2.19)    | 0.034<br>(0.14)      | 0.208<br>(0.63)      | -0.159<br>(-0.44)    |
| NEWS         | -0.276**<br>(-2.13)  | -0.057<br>(-0.25)    | -0.435<br>(-1.49)    | 0.065<br>(0.18)      |
| Observations | 5,422                | 5,422                | 5,422                | 5,422                |
| R-sq         | 0.006                | 0.009                | 0.013                | 0.030                |

Table 2.4: Divergence of Opinion and Returns (Berkman et al. Controls)

This table tests the relationship between divergence of opinion and abnormal returns. Abnormal returns are calculated from the Fama-French 3 Factor model. *CAR20* is the cumulation of abnormal returns over the (0,19) event window. We include *DIVOP* as the key explanatory variable, which is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the (-1,-11) event window. We add the key control variables *IMPACT*, which measures the level of investor sentiment, *NEWS*, which measures the sentiment of firm press releases and *SUE*, which measures standardized unexpected earnings as the difference between the actual and expected (based on IBES median estimates) earnings divided by price.. We also include, but do not report, the additional controls of firm market capitalization and ten days of lagged abnormal returns. We include five proxies for divergence of opinion following Berkman et al. (2009) *INCVOL*, *RETVOL*, *DISP*, *AGE* and *TURN*. *INCVOL* is the standard deviation of seasonally differenced ratio of quarterly operating income divided by average total assets measured over the 20 quarters prior to the earnings announcement quarter. *RETVOL* is the standard deviation of a firm's daily excess returns relative to a value weighted market index, over a 45 day period ending 10 days before the earnings announcement. *DISP*, dispersion of analysts' forecasts. *AGE* is  $\ln(1/\text{firm age})$ , where firm age is the length of time a firm has been covered by CRSP. *TURN* is average daily turnover measures over a 45 day period ending 10 days before the earnings announcement day. All regressions are firm fixed effects panel data regressions on our 1,398 sample firms. Driscoll and Kraay (1998) robust t-stats, which control for time series and cross-sectional correlation, are reported below coefficients in parentheses \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

| VARIABLES    | (1)<br>CAR20        | (2)<br>CAR20         | (3)<br>CAR20         | (4)<br>CAR20         | (5)<br>CAR20         |
|--------------|---------------------|----------------------|----------------------|----------------------|----------------------|
| DIVOP        | -1.302**<br>(-2.44) | -1.465***<br>(-2.66) | -1.589***<br>(-2.87) | -1.522***<br>(-2.81) | -1.650***<br>(-3.44) |
| SUE          | -1.304<br>(-0.51)   | -0.674<br>(-0.26)    | -0.799<br>(-0.31)    | -0.741<br>(-0.29)    | 1.083<br>(0.64)      |
| IMPACT       | 0.403<br>(1.20)     | 0.193<br>(0.59)      | 0.207<br>(0.63)      | 0.261<br>(0.81)      | 0.446<br>(1.32)      |
| NEWS         | -0.521*<br>(-1.77)  | -0.476*<br>(-1.65)   | -0.434<br>(-1.49)    | -0.480<br>(-1.64)    | -0.574*<br>(-1.66)   |
| INCVOL       | 4.826**<br>(2.52)   |                      |                      |                      |                      |
| RETVOL       |                     | 0.987**<br>(2.08)    |                      |                      |                      |
| DISP         |                     |                      | 0.081*<br>(1.88)     |                      |                      |
| AGE          |                     |                      |                      | -5.572<br>(-1.01)    |                      |
| TURN         |                     |                      |                      |                      | 0.604<br>(0.51)      |
| Observations | 5,132               | 5,419                | 5,422                | 5,328                | 4,285                |
| R-sq         | 0.018               | 0.015                | 0.014                | 0.013                | 0.018                |

Table 2.5: **Divergence of Opinion and Returns (Garfinkel et al. Controls)**

This table tests the relationship between divergence of opinion and abnormal returns. Abnormal returns are calculated from the Fama-French 3 Factor model. *CAR20* is the cumulation of abnormal returns over the (0,19) event window. We include *DIVOP* as the key explanatory variable, which is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the (-1,-11) event window. We add the key control variables *IMPACT*, which measures the level of investor sentiment, *NEWS*, which measures the sentiment of firm press releases and *SUE*, which measures standardized unexpected earnings as the difference between the actual and expected (based on IBES median estimates) earnings divided by price. We also include, but do not report, the additional controls of firm market capitalization and ten days of lagged abnormal returns. Daily market adjusted turnover, *MATO*, is the firm volume divided by shares outstanding, minus a similarly calculated measure over all NYSE/AMEX firms.  $\Delta TO$  is the change in market adjusted turnover, which is calculated as daily *MATO* minus the median 180 moving average of *MATO*. Standardized unexpected volume, *SUV*, is the scaled prediction errors from a market model-style model regression of volume on absolute valued returns. All regressions are firm fixed effects panel data regressions on our 1,398 sample firms. Driscoll and Kraay (1998) robust t-stats, which control for time series and cross-sectional correlation, are reported below coefficients in parentheses \*\*\* p<0.01, \*\* p<0.05, \* p<0.1

| VARIABLES    | (1)<br>CAR20        | (2)<br>CAR20         | (3)<br>CAR20         |
|--------------|---------------------|----------------------|----------------------|
| DIVOP        | -1.056**<br>(-2.39) | -1.587***<br>(-3.45) | -1.384***<br>(-2.97) |
| SUE          | 3.161<br>(0.57)     | 2.801<br>(0.65)      | 2.763<br>(0.50)      |
| IMPACT       | -0.114<br>(-0.65)   | -0.205<br>(-1.36)    | -0.166<br>(-1.07)    |
| NEWS         | -0.156<br>(-0.67)   | -0.137<br>(-0.50)    | -0.165<br>(-0.64)    |
| MATO         |                     | -0.624<br>(-1.26)    | -0.422<br>(-0.73)    |
| $\Delta TO$  |                     | -0.293<br>(-0.78)    | -0.558<br>(-1.44)    |
| SUV          | -0.172<br>(-0.69)   |                      | 0.292<br>(1.26)      |
| Observations | 4,208               | 4,230                | 4,111                |
| R-sq         | 0.013               | 0.012                | 0.013                |

Table 2.6: **Divergence of Opinions Event Study**

This table presents an earnings announcement day event study which sorts on the variable *DIVOP*, which is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the (-1,-11) event window. Pre and post earnings announcement cumulative returns over various event window from 30 days before an earnings announcement to 60 days after the announcement. The top panel present the returns to firms were there is divergence of opinions and the bottom panel present the returns to firms without divergence before the earnings announcement.

| Divergence of Opinions |     |        |           |
|------------------------|-----|--------|-----------|
| Window                 | N   | CAR    | T-stat    |
| (-30,-2)               | 730 | 0.54%  | 0.591     |
| (-1,0)                 | 730 | 0.61%  | 2.545**   |
| (+1,+5)                | 730 | -0.67% | -1.767*   |
| (+1,+10)               | 730 | -0.87% | -1.624    |
| (+1,+15)               | 730 | -1.01% | -1.539    |
| (+1,+20)               | 730 | -1.56% | -2.048**  |
| (+1,+30)               | 730 | -1.89% | -2.029**  |
| (+1,+50)               | 730 | -3.66% | -3.040*** |
| (+1,+60)               | 730 | -4.00% | -3.036*** |

| No Divergence of Opinions |      |        |          |
|---------------------------|------|--------|----------|
| Window                    | N    | CAR    | T-stat   |
| (-30,-2)                  | 4139 | -0.85% | -1.905*  |
| (-1,0)                    | 4139 | 0.51%  | 4.343*** |
| (+1,+5)                   | 4139 | 0.22%  | 1.177    |
| (+1,+10)                  | 4139 | -0.14% | -0.531   |
| (+1,+15)                  | 4139 | -0.16% | -0.516   |
| (+1,+20)                  | 4139 | -0.10% | -0.275   |
| (+1,+30)                  | 4139 | 0.02%  | 0.038    |
| (+1,+50)                  | 4139 | -0.39% | -0.664   |
| (+1,+60)                  | 4139 | -0.57% | -0.89    |

Table 2.7: Divergence of Opinion Event Study (SUE Ranks)

This table presents an earnings announcement day event study which sorts on the variables *DIVOP* and *SUE*. *DIVOP* is a dummy variable that takes the value of 1 if there is disagreement between firm press releases and investor sentiment over the (-1,-11) event window, and *SUE* measures standardized unexpected earnings as the difference between the actual and expected (based on IBES median estimates) earnings divided by price. Firms are first sorted into three groups based on *SUE* and then sorted into two subsets based on the presence of *DIVOP*. Pre and post earnings announcement cumulative returns over various event window from 30 days before an earnings announcement to 60 days after the announcement. The results for high *SUE* firms are reported on the left side of the table and low *SUE* firms are reported on the right.

| High SUE |      |       |           | Low SUE  |      |        |           |
|----------|------|-------|-----------|----------|------|--------|-----------|
| Window   | N    | CAR   | T-stat    | Window   | N    | CAR    | T-stat    |
| (-30,-2) | 1407 | 1.09% | 1.503     | (-30,-2) | 1154 | -1.55% | -2.328**  |
| (-1,0)   | 1407 | 2.28% | 11.955*** | (-1,0)   | 1154 | -1.41% | -8.082*** |
| (+1,+5)  | 1407 | 2.83% | 9.393***  | (+1,+5)  | 1154 | -2.47% | -8.953*** |
| (+1,+10) | 1407 | 2.80% | 6.561***  | (+1,+10) | 1154 | -2.97% | -7.622*** |
| (+1,+15) | 1407 | 2.89% | 5.531***  | (+1,+15) | 1154 | -2.77% | -5.789*** |
| (+1,+20) | 1407 | 3.15% | 5.224***  | (+1,+20) | 1154 | -2.41% | -4.371*** |
| (+1,+30) | 1407 | 3.41% | 4.627***  | (+1,+30) | 1154 | -1.99% | -2.947*** |
| (+1,+50) | 1407 | 3.41% | 3.578***  | (+1,+50) | 1154 | -2.18% | -2.504**  |
| (+1,+60) | 1407 | 3.31% | 3.171**   | (+1,+60) | 1154 | -2.12% | -2.214**  |

| High SUE (Div_op) |     |        |          | Low SUE (Div_op) |     |        |           |
|-------------------|-----|--------|----------|------------------|-----|--------|-----------|
| Window            | N   | CAR    | T-stat   | Window           | N   | CAR    | T-stat    |
| (-30,-2)          | 216 | 6.04%  | 3.369*** | (-30,-2)         | 164 | -0.11% | -0.058    |
| (-1,0)            | 216 | 2.83%  | 6.017*** | (-1,0)           | 164 | -0.92% | -1.809*   |
| (+1,+5)           | 216 | 1.36%  | 1.831*   | (+1,+5)          | 164 | -3.17% | -3.951*** |
| (+1,+10)          | 216 | 1.46%  | 1.384    | (+1,+10)         | 164 | -3.37% | -2.973*** |
| (+1,+15)          | 216 | 1.77%  | 1.372    | (+1,+15)         | 164 | -3.38% | -2.433**  |
| (+1,+20)          | 216 | 1.47%  | 0.989    | (+1,+20)         | 164 | -3.84% | -2.394**  |
| (+1,+30)          | 216 | 1.38%  | 0.759    | (+1,+30)         | 164 | -4.19% | -2.133*   |
| (+1,+50)          | 216 | -0.28% | -0.118   | (+1,+50)         | 164 | -4.51% | -1.779*   |
| (+1,+60)          | 216 | -0.54% | -0.211   | (+1,+60)         | 164 | -5.27% | -1.897*   |

Table 2.8: **Changes in Opinions**

This table presents an event study condition on changes in opinions around earnings. If the variable *DIVOP* takes the value of one then there is disagreement, if *DIVOP* is zero then there is agreement. The pre-event period is measured over the (-11,-1) window, and the post event period is measured over the (1,11) window. We define four possible outcomes for changes in opinion. AA, which is defined as agreement before and after an earnings announcement, AD, agreement before and disagreement after DD, disagreement before and after and DA, disagreement before and agreement after the earnings announcement. We report various event windows over the (-11,20) event time period.

| Shifts of Opinion Around Earnings |      |        |           |                  |     |        |          |
|-----------------------------------|------|--------|-----------|------------------|-----|--------|----------|
| Agree-Agree (AA)                  |      |        |           | Divop-Divop (DD) |     |        |          |
| Window                            | N    | CAR    | T-stat    | Window           | N   | CAR    | T-stat   |
| (-11,-1)                          | 3715 | -0.11% | -0.412    | (-11,-1)         | 437 | 1.52%  | 1.962*   |
| (-1,+1)                           | 3715 | 0.47%  | 3.442***  | (-1,+1)          | 437 | 0.26%  | 0.652    |
| (+1,+5)                           | 3714 | -0.36% | -2.034*   | (+1,+5)          | 437 | -0.03% | -0.049   |
| (+1,+11)                          | 3714 | -0.96% | -3.678*** | (+1,+11)         | 437 | 0.14%  | 0.181    |
| (+1,+15)                          | 3714 | -0.85% | -2.771**  | (+1,+15)         | 437 | -0.11% | -0.118   |
| (+1,+20)                          | 3714 | -0.75% | -2.126*   | (+1,+20)         | 437 | -0.83% | -0.793   |
| Agree-Divop (AD)                  |      |        |           | Divop-Agree (DA) |     |        |          |
| Window                            | N    | CAR    | T-stat    | Window           | N   | CAR    | T-stat   |
| (-11,-1)                          | 561  | 0.86%  | 1.384     | (-11,-1)         | 730 | -0.06% | -0.106   |
| (-1,+1)                           | 561  | 1.65%  | 5.078***  | (-1,+1)          | 730 | 0.17%  | 0.612    |
| (+1,+5)                           | 561  | 1.02%  | 2.440*    | (+1,+5)          | 730 | -0.95% | -2.632** |
| (+1,+11)                          | 561  | 2.19%  | 3.524***  | (+1,+11)         | 730 | -1.63% | -3.064** |
| (+1,+15)                          | 561  | 1.80%  | 2.483*    | (+1,+15)         | 730 | -1.61% | -2.591** |
| (+1,+20)                          | 561  | 1.05%  | 1.257     | (+1,+20)         | 730 | -1.88% | -2.613** |

Table 2.9: Linguistic Characteristics of News

This table presents a linguistic analysis produced by DICTION software. The normative values are set by using a set of firm public relations news. The high and low range of normative values set a range that is average for this type of communication. This table present the results for all articles that are associated with either agreement or DIVOP in the pre-earnings-announcement period. The tables present the 5 master variables used by DICTION as well as four calculated variables. See section 5 for a description of these variables. A \* indicates that the features of a document set fall outside of the normal parameters for public relations news.

| Master Variable      | AGREE Score | DIVOP Score | Normative Values |            |
|----------------------|-------------|-------------|------------------|------------|
|                      |             |             | Low Range        | High Range |
| Activity             | 50.38       | 50.70       | 48.16            | 52.43      |
| Optimism             | 48.64       | 47.27*      | 48.21            | 55.58      |
| Certainty            | 50.65       | 45.09*      | 48.44            | 52.71      |
| Realism              | 46.40       | 47.97       | 44.40            | 50.67      |
| Commonality          | 50.39       | 50.01       | 48.40            | 54.08      |
|                      |             |             | Normative Values |            |
| Calculated Variables | AGREE Score | DIVOP Score | Low Range        | High Range |
| Insistence           | 49.43       | 133.55*     | 9.40             | 99.67      |
| Embellishment        | 0.44        | 0.33        | 0.27             | 0.94       |
| Variety              | 0.54*       | 0.5         | 0.45             | 0.53       |
| Complexity           | 4.91        | 5.04        | 4.62             | 5.40       |

## Chapter 3

# Do Local Investors Know More? A Direct Examination of Individual Investors Information Set

### **Abstract**

We examine whether individual investors have value-relevant information about local firms using a unique dataset of 37,065 Twitter posts from 53 notable individual investors covering 1,184 U.S. companies from October 1, 2008 to December 31, 2009. While these investors on average exhibit a surprisingly negative stock picking ability, they are significantly more informed about local companies than about nonlocal companies. The differential stock return predictability between local and nonlocal Twitter posts is 19 basis points per week, both statistically

and economically significant. Additionally, local advantage concentrates in firms without public news coverage and firms with greater information asymmetry, indicating that local advantage is attributable to individual investors' access to private information about local companies.

## 1 Introduction

Do investors have more value-relevant information about local firms? This question has important implications for financial market efficiency, the information diffusion process, and investment practices. In addition, local advantage can shed light on the puzzling “home bias” in which investors exhibit a strong preference for locally headquartered stocks (e.g., Coval and Moskowitz (1999, 2001), Ivkovic and Weisbenner (2005)). The pioneering work of Coval and Moskowitz (2001) provides evidence of local advantage by showing that mutual fund managers earn abnormal returns on their local investments.<sup>1</sup> Researchers have also documented local advantage for various financial market participants including analysts (Malloy (2005), Bae, Stulz, and Tan (2008)), commercial and investment banks (Bulter (2008), Agarwal and Hauswald (2010)), institutional shareholders as monitors (Gaspar and Massa (2007), Ayers, Ramalingegowda, and Yeung (2011)), and acquirers (Almazan, Motta, Titman and Uysal (2010)).

Despite the extensive studies on local advantage, the empirical evidence is mixed for individual investors. Ivkovic and Weisbenner (2005) observe local advantage in a sample of 34,517 households from 1991 to 1996. They find that individual investors earn an abnormal return of 3.2% per annum on their local holdings relative to their nonlocal holdings. However, a recent paper by Seasholes and Zhu (2010) uses a calendar-time portfolio approach and shows that individual investors earn only zero alphas on their local holdings relative to

---

<sup>1</sup> Baik, Kang, and Kim (2010) and Bernile, Kumar, and Sulaeman (2010) further find that local ownerships of general institutions positively predict stock returns.

their nonlocal holdings, and their purchases of local stocks significantly *underperform* their sales of local stocks. Seasholes and Zhu conclude that “individuals do not seem to have value-relevant information about the local stocks they trade.”

In this paper we take a novel approach to investigate whether individuals have value-relevant information about local stocks. While previous studies focus on the performance of local investments by individual investors, we use a large sample of Twitter posts to directly examine investors’ information about local and nonlocal firms. Twitter is an electronic social network where users post short thoughts of no more than 140 characters, called tweets. As of April 2011, Alexa.com, a website that provides statistics on web traffic, rated Twitter as the 9th most popular website on the Internet, and estimated that Twitter accounted for 3.0% to 5.8% of all global internet traffic in 2010. We collect Twitter posts on publicly traded U.S. companies from the website Stocktwits.com, a popular platform for Twitter users to tweet about stocks. We focus on Twitter users in Stocktwits.com’s list of “recommended” users, who generally have a large following, a long track record, post meaningful or interesting comments, and are notable within the social network. Our sample contains 37,065 Twitter posts from 53 Twitter users, covering 1,184 publicly traded U.S. companies from October 1, 2008 to December 31, 2009.

Our sample Twitter users have two unique features. First, they are considered “smart” individual investors; most sample users have extensive trading experience, publish books on trading, sell financial services, manage private accounts, or have held positions in the financial industry (we tabulate their professions and additional investing information in the Appendix). This feature effectively reduces the potential noise in the Twitter posts.<sup>2</sup> More importantly, this feature provides a powerful setting to detect local advantage because these individual investors may have sufficient knowledge and incentives to acquire value-relevant

---

<sup>2</sup>On the contrary, Twitter posts from random users may contain noise. For example, a Twitter user may recommend selling Wal-Mart’s stock after a bad experience at the local Wal-Mart store caused by her ill temper.

information about a local company. Second, our sample Twitter users have available location information, enabling us to classify local versus nonlocal posts. These Twitter users are geographically dispersed, which avoids the issue of geographical clustering.

We start by examining the overall stock-picking ability of our sample Twitter users. We quantify the evaluations in the Twitter posts using a maximum entropy (ME) classification approach and examine whether these evaluations predict stock returns. Our setting offers an interesting test on the informedness of individual investors because it is free of various trading complications such as liquidity and various trading constraints.<sup>3</sup> Surprisingly, we find that the sample individual investors, despite their superior experience and knowledge about the stock market, have a *negative* return predictability. For example, a one standard deviation increase in the two-week evaluation measure is associated with a 14 basis point *decrease* (t-stat -3.05) in abnormal stock returns in the subsequent week. This result is robust to the alternative return windows from two days to one month, alternative measurements of evaluations, and skipping a week before return measurements to control for microstructure effects.

Next, we examine whether the Twitter users have a local advantage. We further classify Twitter posts into local and nonlocal posts according to whether the distance between a user's location and corporate headquarters is within 100 miles.<sup>4</sup> We find a sharp contrast in return predictability between local and nonlocal Twitter posts. While nonlocal posts still exhibit strongly negative return predictability, local posts have insignificantly positive return predictability. For example, a one standard deviation increase in nonlocal evaluations predicts a decrease of 17 basis points (t-stat -3.63) in weekly stock returns, but a one standard deviation increase in local evaluations predicts a increase of 2 basis points (t-stat

---

<sup>3</sup>For example, Kaniel, Saar, Titman (2008) find evidence that individuals make trading profits by providing liquidity.

<sup>4</sup>Ivokovic and Weisbenner (2007) show that local information diffusion among individual investors dissipates dramatically beyond 50 miles. For robustness, we also classify local and nonlocal posts by whether the Twitter user and the company are located in the same state, and find similar results.

0.36) in weekly stock returns. Local advantage, measured as the differential predictability between local and nonlocal evaluations, is a statistically and economically significant 19 basis points (t-stat 2.31) per week. The local advantage persists for return measurement windows from two days to one month, and for alternative constructions of the evaluation measures. The local advantage is also robust to controls of price factors, return autocorrelations, microstructure effects, and geographical factors. Additionally, the differential predictability is not likely to be caused by Twitter followers' stronger responses to local posts.

The superior stock picking ability of local investors is consistent with individual investors possessing value-relevant private information about local firms ("information hypothesis"). However, an alternative explanation is that local investors are more sophisticated ("sophistication hypothesis"). For example, individual investors may be better at analyzing public information of local firms or exhibiting less irrationality regarding local firms. To disentangle private information from sophistication, we examine local advantage across firms with and without public news coverage during the period of Twitter posts.<sup>5</sup> If local advantage is due to local investors' access to private information, then we expect the local advantage to be stronger among firms with no public news coverage. Our results shows that local advantage concentrates in the firms without public news coverage (36 basis points per week) but becomes insignificantly negative for firms with public news coverage (-8 basis points per week). This finding lends strong support to the information hypothesis.

For a more thorough investigation, we also conduct cross-sectional analyses of local advantage across levels of information asymmetry. If local advantage is due to investors' private information, then we expect the local advantage to be more pronounced in firms with more severe information asymmetry. Previous studies provide evidence that information asymmetry is severe in small firms (e.g., Coval and Moskowitz (1999)), non-S&P 500 firms

---

<sup>5</sup>We collect news articles from the three major news wires including Reuters News, Dow Jones News Wire, and PR News Wire.

(e.g., Ivkovic and Weisbenner (2005)), low analyst coverage firms (e.g., Brennan and Subrahmanyam (1995)), and firms with high idiosyncratic volatility (e.g., Krishnaswami and Subramaniam (1999)). Our results based on these proxies consistently show a significantly positive association between local advantage and information asymmetry, indicating that local advantage is caused by individual investors' access to private information about local firms.

We further examine local advantage using a rolling long-short strategy based on local advantage that is similar to the rolling momentum strategy proposed by Jegadeesh and Titman (1993). We first contrast local and nonlocal posts by constructing a differential evaluation measure which is local evaluation minus nonlocal evaluation. We then design a rolling strategy that goes long the stocks with positive differential evaluations in the previous two weeks and goes short the stocks with negative differential evaluations. The average daily abnormal profits of this zero-investment strategy are significantly positive, ranging from 4.6 basis points to 8.9 basis points for the holding windows from two days to one month. The profits of this local-advantage-based strategy increase significantly when we focus on firms with greater information asymmetry (small firms, Non-S&P 500 firms, low analyst coverage firms, or high idiosyncratic volatility firms). For example, when we focus on firms in the bottom tercile of analyst coverage, we observe daily abnormal profits from 22.3 basis points to 27.2 basis points for holding windows from two days to one month.

Our findings shed light on whether individual investors have a local advantage (e.g., Ivkovic and Weisbenner (2005), Seasholes and Zhu (2010)). We are the first to directly examine investors' information about local and nonlocal companies, and our evidence suggests that individual investors can have a statistically and economically significant informational advantage about local firms. The recent work by Seasholes and Zhu (2010) suggests that individuals do not have a local advantage because they earn zero or negative abnormal returns on local investments relative to nonlocal investments. We differ from their approach

in that we focus on a much smaller group of notable individual investors, and we directly study investors' information instead of investment returns. Together with Seasholes and Zhu (2010), our results indicate that local advantage may be limited to a sub-group of individual investors who are sufficiently motivated and capable of acquiring information about local firms. Additionally, individual investors may encounter difficulties in converting their value-relevant information into trading profits, possibly due to their constraints or irrationalities.<sup>6</sup>

We also contribute to a rapidly growing literature on the informedness of individual investors. Some studies provide evidence that individual investors lose significantly from their trading (Barber and Odean (2000)) or that individuals make trading profits but only through liquidity provision (Kaniel, Saar, and Titman (2008)). However, several recent papers suggest that individual investors can be informed prior to earnings announcements (Kaniel, Saar, Titman, Liu (2011)), takeovers (Griffin, Shu, and Topaloglu (2011)), or when they use market orders (Kelley and Tetlock (2011)). Our paper performs a unique examination of the informedness of individual investors since we directly study the information possessed by individual investors and this test design is free of trading complications such as liquidity or price impact. We find that, on the one hand, even the "smart" individual investors who have extensive trading experience and investing knowledge exhibit *negative* stock return predictability, which underlines the lack of stock-picking ability for individual investors. On the other hand, individual investors are able to acquire private information when they live close to a firm's headquarters.

Our study also extends the literature on internet communication about the financial markets. Motivated by the rapid growth in internet communication in the past two decades, financial researchers have started to investigate whether stock specific internet messages

---

<sup>6</sup>Short sales constraints (borrowing constraints) can prevent investors from capitalizing on negative (positive) information. Investor irrationalities may also prevent individual investors from utilizing information. For instance, the disposition effect (Odean (1998), Grinblatt and Han (2005)) can make investors hold on to a loser despite negative information.

contain value-relevant information or just noise. Several studies find that messages posted on internet stock message boards (e.g., Yahoo! Finance) have little to no predicting power for stock returns, suggesting that these messages may simply contain noise (Tumarkin and Whitelaw (2001), Antweiler and Frank (2004), Das, Martinez-Jerez, and Tufano (2005)). Consistent with these studies, we also document a negative stock return predictability for the Twitter posts in our sample. However, we extend the previous studies by documenting that the Twitter posts may contain value-relevant information about local companies. Our results therefore provide interesting evidence that instead of all noise, internet messages also contain value-relevant information about financial markets.

The outline of our paper is as follows. Section II describes the Twitter data and sample construction. Section III presents the evidence of local advantage. Section IV analyzes the source of local advantage, and Section V concludes.

## 2 Data and Sample Selection

### 2.1 Why do People Share Financial Information in Twitter Posts?

Twitter is a micro blogging application where users are able to post short thoughts of no more than 140 characters, called tweets. We collect Twitter posts from Stocktwits.com, an open micro-blogging site which is powered by Twitter and focused on financial markets. Stocktwits.com was founded in 2008 and has since then become a popular website for Twitter users to exchange investment information. Since its inception, Stocktwits.com has been covered by major news media such as The New York Times and CNNMoney.com.<sup>7</sup>

A natural question is why people would share value-relevant information in their Twitter posts. More broadly, why do people post internet messages on the stock markets? Despite

---

<sup>7</sup>In 2010, Stocktwits was named Time.com's top 50 best websites and Fast Company's top 10 innovative companies in finance.

the rapid growth in internet communication about financial markets, there is not much theoretical literature that rationalizes this type of communication. DeMarzo, Vayanos, and Zweibel (2003) propose a model in which investors fail to account for the repetition of opinions (“persuasion bias”). In equilibrium, well-connected agents in a social network can have significant influence on the actions of other members, and therefore, impact the market. This model can potentially explain why Twitter users may have incentives to gain popularity and followers by sharing value-relevant information in their Twitter posts. Consistent with this view, the models in Cao, Coval, and Hirshleifer (2002), Colla and Mele (2010), and Hong, Hong, and Ungureanu (2011) show that information sharing among investors can cause trading and affect the outcomes of the financial markets.<sup>8</sup>

Additionally, gaining popularity in the Twitter world may also bring direct financial benefits. For example, Stocktwits.com now offers a full “marketplace” of premium blogs to users of the site. These premium blogs are based on the tweets and trading ideas of successful investors from the Stocktwits.com community. The themes of the premium streams range from value investing to swing trading, and annual subscriptions can cost in excess of \$800. Therefore the potential financial benefits can also motivate Twitter users to gain followers by sharing value-relevant information.

## 2.2 Collection of Twitter Posts and Construction of Sample

Figure 1 provides an example of the stream of Twitter posts that comprise our sample. Twitter users comment about a specific company by referring to the company’s ticker preceded by a “\$” hashtag. An example of this would be “\$MSFT and \$AAPL are a buy!”. Hashtagging allows us to extract specific company references with a high level of accuracy by looking for the “\$” hashtag followed by one to four capital letters that constitute the ticker symbol. In

---

<sup>8</sup>Empirical evidence also suggests that people listen to ideas from friends to make financial decisions (Duflo and Saez (2002)). Additionally, Hong, Kubik, and Stein (2004, 2005) show that institutional and individual investors’ investment decisions are affected by other institutions in the same area or neighbors.

the case of multiple company references in one post, like the example above, each reference is counted as a unique post.

We focus on users in Stocktwits.com’s “recommended” list of contributors.<sup>9</sup> Although the website does not use a quantitative rule for selecting the recommended contributor, in general, the recommended contributors have a large following, a long track record, post meaningful or interesting comments, and are influential within the social network. We select our sample of Twitter users from the set of recommended users as of October 1, 2008; we also require a Twitter user to have at least 100 followers and 300 posts at the beginning of the data collection process. This set of criteria produces a sample of 74 Twitter users.

Most Twitter users provide a short personal profile that consists of their location, personal website, and a short biography. The following is a sample profile:

“Based in Atlanta, GA, Tony Johnson has over 16 years of trading experience during which he developed his investment strategy. Tony worked in Atlanta for ABC Group, Inc., as an Institutional Account Representative. During his tenure at ABC, Tony gained extensive knowledge of the technical analysis of stocks, market forecasting, and risk management. Tony completed his undergraduate work at XYZ University, receiving a B.A. in Economics in 1992. Outside of trading, Tony enjoys playing the drums, travel, and sports.”<sup>10</sup>

42 users from our initial sample provide an explicit location in their profiles. We further identify the locations of 11 users using the information from their bios or personal websites. The remaining 21 users from our initial sample either do not provide enough information to identify their locations, or live outside of the continental U.S. Therefore, our final sample contains 53 Twitter users with location information. Figure 2 plots the locations of our sample Twitter users, showing that their geographical distribution is quite dispersed.

---

<sup>9</sup>The Twitter posts by recommended contributors are classified as “suggested stream” by Stocktwits.com and are presented at <http://stocktwits.com/streams/suggested>.

<sup>10</sup>Although a real example, we change the user’s name, location, former employer name, college name, and year of graduation.

For readers to better understand the investing backgrounds of the sample Twitter users, we collect professions and additional information for all but three of the users. We present the information in the Appendix. Most of the sample users have a long trading history, have published books on investing, sell financial services, manage private accounts, have held positions in financial industries, post about financial markets on a regular basis, or engage in derivatives trading. This unique feature reduces the noise in our sample Twitter posts and provides a powerful setting to detect local advantage because these individuals should be the most capable at acquiring value-relevant information about a local company.

We require the sample firms to have available CRSP data and at least one Twitter post over our sample period. To control for microstructure effects, we drop penny stocks that are priced below two dollars. Our final sample contains 37,065 Twitter posts covering 1,184 publicly traded companies from October 1, 2008 to December 31, 2009. We obtain accounting data for our sample firms from Compustat and analyst data from IBES. We also obtain the daily Fama-French four factors from Kenneth French’s data library for the construction of abnormal returns.<sup>11</sup> Some of our tests use news articles collected from Factiva, and we will describe the details of the construction of news data when we discuss the corresponding tests.

## 2.3 Classifying Local and Nonlocal Posts

We calculate the straight line geographic distance between the location of each Twitter user and the headquarters of each company in our sample using longitude and latitude coordinates. If a Twitter user provides a city as her location instead of longitude and latitude coordinates, we assign longitude and latitude coordinates based on the zip code that is central to the city.<sup>12</sup> We obtain zip codes of corporate headquarters from Compustat

---

<sup>11</sup> We thank Professor Kenneth French for making the data available.

<sup>12</sup>We match longitude and latitude coordinates to zip codes using the database from <http://www.getzipcodedata.com/#>.

and assign the corresponding longitude and latitude coordinates. We then calculate the distance between Twitter user and company headquarters using the following equation:

$$Distance = 7921.6623 * \arcsin(\sqrt{(\sin((0.017 * lat2 - 0.017 * lat1)/2))^2 +} \quad (3.1)$$

$$\cos(0.017 * lat1) * \cos(0.017 * lat2) * (\sin((0.017 * long2 - 0.017 * long1)/2))^2) \quad (3.2)$$

where *lat1* and *long1* are the latitude and longitude coordinates of a Twitter user and *lat2* and *long2* are the latitude and longitude coordinates of corporate headquarters.<sup>13</sup>

We classify a Twitter post as local (nonlocal) if the distance between the Twitter user and the corporate headquarters is within (more than) 100 miles. Previous studies use various criteria of distance to classify local stocks, from 62 to 250 miles (e.g., Coval and Moskowitz (2001), Ivokovic and Weisbenner (2005), Malloy (2005), Seasholes and Zhu (2010)). We adopt the moderate 100-mile criterion because Ivokovic and Weisbenner (2007) show that local information diffusion among individual investors dissipates dramatically beyond 50 miles. Our approach classifies 3,498 posts as local and 33,567 posts as nonlocal. For robustness we also try classifying a post as local (nonlocal) if the Twitter user is in the same state as the firm's (a different state from the firm's) and find similar results from our tests.

## 2.4 Quantifying the Information in Twitter Posts

We use the maximum entropy (ME) approach to classify the information in Twitter posts. The ME approach derives meaning from natural language in the posts by applying a statistical algorithm to qualitative data. Since the information in a Twitter post can be subtle,

---

<sup>13</sup> This equation is provided by SAS at <http://www2.sas.com/proceedings/sugi31/143-31.pdf>. This approach is based on the great circle distance model which is similar to the distance equations used in the literature (e.g., Ivokovic and Weisbenner (2007)) but provide greater accuracy at small distances. More details about the distance models can be found at [http://en.wikipedia.org/wiki/Great-circle\\_distance](http://en.wikipedia.org/wiki/Great-circle_distance).

using key word frequencies alone can cause misclassification. For example, the statement “You would be crazy to sell \$GOOG right now” contains the word “sell” which unconditionally we would assume has a negative connotation. However, the statement “crazy to sell” is obviously a positive statement for Google. ME classification is considered the most robust technique for information classification because it controls for the conditional dependence of words (Pang, Lee, and Vaithyanathan (2002)). Unlike the less sophisticated procedures which handle each word as an unconditional feature, ME classification uses the information contained in multiple word phrases such as “crazy to sell” to more accurately classify information.

In addition to controlling for the conditional dependence of words, the ME classification also avoids the misidentification issue associated with the alternative approaches that simply rely on key-word frequencies. For example, Loughran and McDonald (2011) show that in the textual analysis of 10-K reports, almost three-fourths (73.8%) of the negative word counts according to the widely used Harvard Dictionary are attributable to words that are typically not negative in a financial context (e.g., tax, cost, capital, board, liability). Other words on the Harvard list (e.g., mine, cancer, crude, tire, or capital) are more likely to identify a specific industry segment than reveal a negative financial event. ME classification does not suffer the noise introduced by key-word selection because the identification is based on a large training sample of Twitter posts that we hand classified.<sup>14</sup>

The general idea of ME classification is that when nothing is known about a distribution, the distribution should be uniform, i.e., have maximum entropy. Consider the example of trying to classify a document as positive, negative or neutral, where we are only told that 50% of documents that contain the word “buy” are considered positive. Intuition tells us that if the document has the word “buy” in it then there is a 50% chance that it is a positive

---

<sup>14</sup>Additionally, many previous studies using the Harvard list only count negative words because they find little incremental information in the Harvard positive word list (e.g., Tetlock (2007), Engelberg (2008)). In contrast, ME classification is based on both positive and negative comments in the messages.

post, a 25% chance of being negative, and a 25% to of being neutral. If our document did not have the word “buy” in it then we would just assume an equal distribution of a 33% chance that the document falls into each category. Thus, when we know nothing about our document, we make the distribution uniform. This is the essence of ME classification. In practice, this process is constrained by many features, and the calculations for conditional probabilities become complex, but the logic is still the same as our simple example.

To formally describe the ME procedure, we define the following set of terms. Let  $f_1, \dots, f_m$  be a set of predefined features that can appear in a post. From our previous example, the word “sell” would be a feature, and the tri-gram “crazy to sell” would also be a feature. Let  $n_i(d)$  be the number of times that the feature  $f_i$  occurs in a post  $d$ . Thus, each post is represented by a post vector that takes the form:  $\bar{d} = (n_1(d), n_2(d), \dots, n_m(d))$ . Lastly, let  $c$  be a post category that takes the value of  $c_0$  (positive, negative, or neutral). Given this set of variables, the estimate of  $P(c|d)$  is as follows:

$$P_{ME}(c|d) = \frac{1}{Z(d)} \left( \sum_i \lambda_{i,c} F_{i,c}(d, c) \right) \quad (3.3)$$

$$F_{i,c}(d, c) = 1, n_i(d) > 0, \text{ and } c_i = c_0, 0 \text{ otherwise} \quad (3.4)$$

Where  $Z(d)$  is a normalization function, and  $F_{i,c}$  is a feature class function for the feature  $i$  in the class  $c$  defined as.

$$P_{ME}(c|d) = \frac{1}{Z(d)} \left( \sum_i \lambda_{i,c} F_{i,c}(d, c) \right) \quad (3.5)$$

Where  $Z(d)$  is a normalization function, and  $F_{i,c}$  is a feature class function for the feature  $i$  in the class  $c$ .

$$F_{I,c}(d, c) = 1, n_i(d) > 0, \text{ and } c_i = c_0, 0 \text{ otherwise} \quad (3.6)$$

For example, this feature category function only returns a value of one if the post contains the tri-gram “crazy to sell” and the post is hypothesized to be of negative sentiment.  $\lambda_{i,c}$  is a weighting parameter that determines the relative strength of each of the features  $f_i$  contained in a document. If the value of  $\lambda_{i,c}$  is very large then the feature  $f_i$  is considered to be very strong for a specific category  $c_0$ . We implement the ME classifier by hand classifying a corpus of 1,000 twitter posts. This out of sample set of categorized data is called training set, and is used to calculate the expected values of  $F_{i,c}$ . Next, we use all the Twitter posts to estimate the conditional probabilities  $P_{ME}(c = c_0 | d)$  by maximizing the entropy across the three different categories while satisfying the constraint that the expected values of the feature category functions  $F_{i,c}$  are equal to their training data expected values. Each post in our dataset is then assigned a value of  $(-1, 0, 1)$  based on the highest conditional probability of a post being positive, negative, or neutral. We test the accuracy of this procedure by running the ME classifier on a set of 100 posts that are hand classified. The ME classifier worked well in this out of sample test, and it was able to correctly classify 67% of all posts in the test sample. This accuracy rate is similar to the accuracy level that is achieved in other sentiment classification studies, such as Pang, Lee and Vaithyanathan (2002).<sup>15</sup>

---

<sup>15</sup> It is difficult to compare the accuracy of ME classification with previous studies in the finance literature because they generally use key-word counts directly in the empirical analyses without examining the proportions of correct and incorrect identifications of sentiments. Loughran and McDonald (2011) report that 73.8% of the negative word counts based on Harvard Dictionary are not associated with negative meanings in a financial context, but their sample is 10-K reports instead of internet messages.

## 3 Do Individual Investors Have Local Advantage?

### 3.1 Summary Statistics

Panel A of Table I summarizes the characteristics of our sample. A typical firm in our sample has a market capitalization of \$5,981 million, a book-to-market ratio of 0.40, and is followed by 8.5 analysts. For a comparison, an average firm in the contemporaneous CRSP universe has a market cap of \$2,982 million, a book-to-market ratio of 0.46, and is followed by 3.5 analysts. Since we require sample firms to have at least one Twitter post during the sample period, these comparisons suggest that the firms covered by Twitter users tend to have larger size, higher analyst coverage, and slightly lower book-to-market ratios.<sup>16</sup> We also report average idiosyncratic volatility and daily return for sample firms. Idiosyncratic volatility for a firm-day is the standard error of residuals from the time-series regressions of a firm's excess returns on the daily market factor (MKT) in the previous 150-day rolling window. A typical sample firm has an idiosyncratic volatility of 0.045 and an average daily return of 2.4 basis points.

Panel B further presents firm characteristics sorted on the total number of Twitter posts, local posts, and nonlocal posts, respectively. The results show that Twitter users, whether local or nonlocal, tend to cover larger firms, growth firms, high analyst coverage firms and winner firms. In addition, the number of Twitter posts does not vary with idiosyncratic volatility.

#### *B. Stock Return Predictability of Sample Twitter Users*

We start by examining the stock return predictability of sample Twitter users. Specifically, we estimate the following daily panel regression:

---

<sup>16</sup>We drop penny stocks priced below \$2, which also makes our sample firms larger than the CRSP universe.

$$CAR[t, t+k]_i = \alpha_0 Eval_{it} + \sum \beta_j AR_{it-j} + \sum \gamma_i D_i + \epsilon_{it} \quad (3.7)$$

where  $CAR[t, t+k]_i$  is cumulative abnormal returns of firm  $i$  from day  $t$  to  $t+k$ . For our tests, we examine abnormal returns in the two- ( $k=1$ ), five- ( $k=4$ ), ten- ( $k=9$ ) and twenty-day ( $k=19$ ) windows. We follow the literature (e.g., Fama, Fisher, Jensen, and Roll (1969)) to calculate daily abnormal return as residuals from the Fama-French four-factor model. Specifically, for each firm  $i$  on day  $t$ , we estimate factor loadings from the Fama-French four-factor regressions in the  $[t-150, t-1]$  window and then calculate daily abnormal returns using the estimated factor loadings. The abnormal return  $AR_{it}$  therefore captures price response to the new information arriving on day  $t$ .

The independent variable in equation (3.7),  $Eval_{it}$ , is the aggregate evaluation of Twitter users for firm  $i$  over the two-week period prior to day  $t$ . Specifically, we first assign the scores of either -1 (negative), 0 (neutral), or 1 (positive) to each Twitter post about firm  $i$  in the two weeks prior to day  $t$  using the maximum entropy classification techniques described in Section II.C, and then sum up the scores. We assign zero to the evaluation measure if a firm is not covered by any Twitter post in the two-week period. To ease the assessment of economic significance, we standardize the evaluation variable in each cross-section. We further include firm fixed effects ( $D_i$ ) to control for firm-specific characteristics, and ten lags of daily returns ( $AR_{it-j}$ ) to control for short-term return reversals and microstructure effects. We calculate t-statistics using Driscoll and Kraay (1998) robust standard errors that control for both cross-sectional and time-series correlations.<sup>17</sup>

The coefficients  $\alpha_0$  in equation (3.7) measures the stock return predictability of our sample Twitter users, and the sign of  $\alpha_0$  is unclear ex-ante. On the one hand, previous studies suggest

---

<sup>17</sup>The Driscoll-Kraay standard errors are in the similar spirit as Newey-West standard errors but corrects both time-series and cross-sectional correlations in the panel regression setting.

that individual investors generally lose significantly from their trading (Barber and Odean (2000)), and that when they make trading profits they do so by providing liquidity (Kaniel, Saar, and Titman (2008)). Additionally, researchers have documented that messages on the internet message boards have little stock return predictability (e.g., Antweiler and Frank (2004)). On the other hand, our sample Twitter users may have positive stock picking ability because they are considered “smart” individual investors who have extensive trading experience and investing knowledge. If there is heterogeneity in stock picking ability among individual investors, then our sample investors are among the most likely candidates to demonstrate a positive stock return predictability.

Panel A of Table II presents the results of the return regressions. Interestingly, we observe that the coefficients on investor evaluations are significantly negative for all of the return windows. For example, in the model of five-day returns, the estimated coefficient on investor evaluation is -0.138 (t-stat -3.05), indicating that a one standard deviation increase in the evaluation measure is associated with a 13.8 basis point *decrease* in subsequent weekly returns. In Panel B, we control for microstructure effects by skipping a week before return measurements and find similar results. Panel C repeats the regressions but with evaluations in the one-week period prior to return measurement instead of two-week period, and the negative return predictability persists for all return windows.

Table II paints a pessimistic picture about the investing ability of individual investors. Specifically, even the messages from those among the best individual investors have a significantly negative stock return predictability. This finding also illustrates that one may lose money in the stock market by simply following the seemingly trustworthy advice from the internet. In the next subsection, we explore whether there is a difference in the stock return predictability between local and nonlocal posts.

### 3.2 Do Individual Investors Have Local Advantage?

In this section, we examine whether local Twitter users are more informed of future stock returns than their nonlocal peers. Specifically, we estimate the following daily panel regression:

$$CAR[t, t + k]_i = \alpha_0 NonLocal\_Eval_{it} + Local\_Eval_{it} + \sum \beta_j AR_{it-j} + \sum \gamma_i D_i + \epsilon_{it} \quad (3.8)$$

This regression is similar to equation (3.7) except that we construct investor evaluation measures separately for local posts ( $Local\_Eval_{it}$ ) and nonlocal posts ( $NonLocal\_Eval_{it}$ ). Specifically, the local and nonlocal evaluation measures are the aggregate evaluations by local and nonlocal Twitter users for firm  $i$  over the two-week period prior to day  $t$ , respectively. The coefficients  $\alpha_1$  and  $\alpha_2$  indicate the stock return predictability of local and nonlocal investors. If local investors are better at predicting returns than nonlocal investors, then we expect that  $(\alpha_1 - \alpha_2) > 0$ .

Panel A of Table III presents the results of these regressions. Interestingly, the negative return predictability reported in Table II remains for nonlocal investors but disappears for local investors. For example, in the model of five-day returns, the estimated coefficient is -0.165 (t-stat -3.63) for nonlocal evaluations but 0.022 (t-stat 0.36) for local evaluations. The local advantage ( $\alpha_1 - \alpha_2$ ) is 0.187 (t-stat 2.31), suggesting that a one standard deviation increase in local evaluation relative to nonlocal evaluation predicts an 18.7 basis point increase in weekly stock returns. This local advantage is both economically and statistically significant. Additionally, local advantage is large and significant for the other return windows of two, ten and twenty days. For robustness, we further repeat the regressions with raw returns instead of abnormal returns in Panel B of Table III and observe similar results. The results

in Table III provide strong evidence of local advantage among the individual investors in our sample.

### 3.3 Robustness Tests

We conduct various robustness tests on the local advantage results in Table III. First, we examine whether the results in Table III are specific to the two-week measurement window of evaluations. Panel A of Table IV repeats the regressions in Table III but with the local and nonlocal evaluations measured in the previous one-week window instead of the two-week window. These results consistently show statistically and economically significant coefficients on  $(\alpha_1 - \alpha_2)$  in all models, suggesting that local advantage persists with the alternative measurement windows of evaluations.

We also examine whether our results on local advantage are sensitive to an alternative classification of local posts. In Panel B of Table IV, we classify local and nonlocal posts according to whether the Twitter users are in the same states as the firms' headquarters. Our finding of local advantage persists with the local classification based on the state criterion. For example, in the model of five-day abnormal returns the estimated local advantage  $(\alpha_1 - \alpha_2)$  is 0.204 (t-stat 2.41), slightly larger than the 0.187 in Panel A of Table III.

When no Twitter post covers a firm during the two-week evaluation period, we do not drop the observation but treat it as a neutral evaluation by assigning zero to the evaluation measure. For a robustness test, we repeat the regression analysis but include only the firms with at least one Twitter post in the evaluation period. We present the results in Panel C of Table IV, which shows that local advantage persists for all windows of abnormal returns.

Seasholes and Zhu (2010) point out the importance of controlling for geographical return factors in the examination of local advantage. For example, if both sample firms and sample investors cluster in certain areas (e.g., New England or the Bay area) and if stocks of firms in these areas happen to perform well during the sample period, then one can observe a

mechanically positive relation between local investment and stock performance. Although this concern is alleviated by the firm fixed effects in our regressions, we nevertheless construct state-adjusted return for a firm-day by subtracting the average daily returns of all firms located in the same state. Table V presents the regression analyses with state-adjusted abnormal returns (Panel A) and state-adjusted raw returns (Panel B). The local advantage is both statistically and economically significant in all models.

We further examine whether the results on local advantage are due to local posts causing greater price impact. Specifically, followers of a Twitter user may buy (sell) after reading the user’s positive (negative) evaluation, causing a positive relation between the evaluation and subsequent stock returns. If followers expect local posts to contain more reliable information than nonlocal posts and therefore respond more strongly to local posts than nonlocal posts, then one will observe that local posts predict returns better than nonlocal posts. Our findings on local advantage are unlikely driven by price impact. While price impact is temporary, our results on local advantage persist for the return window up to one month (twenty trading days). We nevertheless conduct three robustness tests to investigate this explanation.

First, we skip one week between the measurement of evaluations and returns. Investor responses to the posts should concentrate in the week after the posts, so if the results on local advantage are caused by investor responses, then local advantage should be significantly reduced in the skip-a-week setting. Panel A of Table VI presents the results of the skip-a-week regressions, which show that the magnitude of local advantage in all models is very similar to those in Table III. For example, the skip-a-week local advantage in terms of five-day abnormal returns is 0.176 percent (t-stat 2.36), very close to the 0.187 percent in Table III.

Second, we construct weighted evaluation measures that assign larger weights to the Twitter users with more followers. Specifically, we first multiply the evaluation score of each post by the number of followers of the Twitter user, and then sum up the weighted scores

for local and nonlocal posts. If the result on local advantage is caused by a stronger investor response to local posts, then we should observe a stronger local advantage with the weighted evaluation measure. Panel B of Table VI presents the results on the regressions with the weighted evaluation measures. Interestingly, the coefficients on both local and nonlocal evaluations are more positive (or less negative) than those in Table III. The improved return predictability is consistent with price impact of the followers but also consistent with the fact that Twitter users with more followers are more informed. Contradicting the price impact story, the increase in the coefficients are actually larger for nonlocal posts than for local posts, causing a smaller local advantage (0.120, t-stat 3.30) than that in Table III. The fact that the weighted measures do not lead to a larger local advantage suggests that price impact, if at work, is not driving the local advantage.

Finally, we examine the corresponding trading volume in the return windows. If our finding on local advantage is caused by followers' stronger response to local posts than nonlocal posts, then we expect to observe a greater increase in trading volume for local posts than nonlocal posts. Panel A of Table VII repeats the regressions in Table II but with cumulative abnormal turnover. We calculate daily turnover as daily trading volume scaled by total shares outstanding, and then follow the literature (e.g., Tkac (1999), Gebhardt, Lee, and Swaminathan (2001)) to control for firm-specific and market-wide factors that affect volume. Specifically, we first calculate daily excess turnover by subtracting cross-sectional average turnover of the CRSP universe, and then obtain abnormal turnover for a firm-day by subtracting the firm's average daily excess turnover in the previous 180-day rolling window. Panel A of Table VII shows that the coefficients on local and nonlocal posts are both positive but do not differ significantly from each other. We further estimate regressions of abnormal turnover on the absolute values of evaluations because both positive and negative evaluations could trigger abnormal volume. Panel B of Table VII shows similar results to those in Panel

A. Overall, the results in Table VII confirm that our findings of local advantage are not driven by the stronger response of followers to local posts.

## **4 Is Local Advantage Private Information or Investor Sophistication?**

Our finding of a local advantage is consistent with individual investors possessing private information about local firms (“information hypothesis”). However, an alternative explanation is that local investors are more sophisticated than nonlocal investors (“sophistication hypothesis”). For example, local investors may be better at analyzing public information than nonlocal investors or they might exhibit less irrationality than nonlocal investors. In this section, we perform a number of cross-sectional analyses to investigate whether local advantage is accounted for by access to private information or investor sophistication.

### **4.1 The Effect of Public News Coverage on Local Advantage**

If local advantage is caused by investors’ access to private information about local firms, then local advantage will be stronger among firms with no public news coverage. We therefore set out to examine the effect of public news coverage on local advantage. To ensure the reliability of news sources we collect news articles from the three major news wires including Reuters News, Dow Jones News Wire, and PR News Wire using Factiva. For each firm in our sample, we first retrieve a list of Factiva firm names that are cross-referenced to the firm’s Compustat name and manually screen the list to delete the incorrect matches. We then use the matched Factiva firm names to download all news articles during our sample period. We collect 40,609 news articles that cover 458 of our sample firms. The 39% coverage during our fourteen-month sample period is consistent with Chan (2003) who finds news coverage for

about 30% of firms in any given year during 1980-2000.<sup>18</sup> For each day of our sample period, we sort firms into two groups based on whether the firms have public news coverage in the previous two weeks (measurement window for Twitter posts). We then estimate regressions of abnormal returns for the sub-samples of firms with and without news coverage.

Table VIII presents the regressions of abnormal stock returns for the no-news (Panel A) and news (Panel B) sub-samples. We observe a local advantage in the no-news sample that is almost twice as large as the local advantage found in the full sample. For example, for the five-day window of abnormal returns in Panel A, local advantage is 36.4 basis points per week (t-stat 3.95). In contrast, Panel B shows that the corresponding local advantage is -8.2 basis points (t-stat -0.57) for firms with news coverage. Panel C further presents the difference in local advantage between no-news and news samples. The spread of local advantage in the five-day return window is a large 44.7 basis points (t-stat 2.59), both economically and statistically significant. These results suggest that local advantage concentrates in the firms that have no public news coverage. This finding lends strong support to the hypothesis that individual investors have access to private information about locally headquartered firms.

## 4.2 The Effect of Information Asymmetry on Local Advantage

If local advantage is caused by local investors' access to private information, then we would expect a positive association between local advantage and information asymmetry. In this section we examine the effect of information asymmetry on local advantage using a number of commonly used proxies proposed by the previous studies.

---

<sup>18</sup> Fang and Peress (2010) examines a sample of large firms (NYSE stocks plus 500 randomly selected NASDAQ stocks), and find that news coverage by the four nationwide newspapers decreased from 77% in 1993 to 57% in 2002.

## **The effect of firm size on local advantage**

Our first proxy of information asymmetry is firm size, which is widely used in the literature (e.g., Coval and Moskowitz (1999), Hong, Lim, and Stein (2000)). Previous studies suggest that small firms have greater information asymmetry than large firms because investors, facing fixed information costs, may exert more effort to learn about large firms in which they can make larger investments. Therefore, if local advantage is caused by private information then we expect local advantage to be stronger in small firms.

For each day in our sample period, we classify firms into three groups according to their market capitalizations from the previous year, and calculate local advantage based on regressions of abnormal returns as in Panel A of Table III for small (bottom tercile) and large firms (top tercile), respectively. Panel A of Table IX shows that local advantage for small firms is significant and much larger than that of the full sample. For example, in the five-day window of abnormal returns, local advantage for small firms is 0.620 (t-stat 3.54), more than three times as large as that of the full sample (0.187, Panel A of Table III). In contrast, local advantage becomes insignificant for large firms (top tercile) in all models. The differences in local advantage between small and large firms are statistically significant in all models. Therefore, the results of the sub-sample analysis based on firm size support the information hypothesis. Since firm size also captures many other aspects of a firm, we perform more cross-sectional analyses using a broad set of proxies for information asymmetry.

## **S&P 500 composite index as a proxy for information asymmetry**

Previous studies on local advantage use the S&P 500 composite index to classify high and low information asymmetry firms. For example, Ivkovic and Weisbenner (2005) document a much stronger local advantage for firms that are not a component of the S&P 500 index. Seasholes and Zhu (2010) also use the S&P 500 index as a measure of information asymmetry, but find no evidence of local advantage in either index or non-index firms. For each day of

our sample period, we sort sample firms into two groups according to whether they are in the S&P 500 composite index at the beginning of the month and then report, in Panel B of Table IX, the local advantage for index and non-index firms, respectively.

Panel B of Table IX demonstrates a sharp contrast between non-index firms and index firms. Specifically, local advantage among non-index firms is large and significant in all return windows, but local advantage among index firms shrinks and becomes insignificant. We further examine the differences in local advantage between non-index and index firms, which are all positive but insignificant. To summarize, these results present some evidence that, consistent with the information hypothesis, local advantage is stronger in non-index firms than index firms.

### **The effect of analyst coverage on local advantage**

Analyst coverage is another commonly used proxy for information asymmetry (e.g., Brennan and Subrahmanyam (1995), Hong, Lim, and Stein (2000), Irvine (2004)). Specifically, firms followed by larger numbers of analysts tend to have lower information asymmetry. Since analyst coverage and firm size are strongly correlated, we construct size-adjusted analyst coverage as the residual from cross-sectional regressions of analyst coverage on firm size. For each day in our sample period, we sort firms into terciles according to their size-adjusted analyst coverage for the month, and examine local advantage for low coverage firms (bottom tercile of coverage) and high coverage firms (top tercile of coverage), respectively.

In Panel C of Table IX, we observe that local advantage for low coverage firms is significantly positive in all return windows, and more than twice as large as that of the full sample (Panel A of Table III). In contrast, local advantage completely disappears for the high coverage firms. The spread in local advantage between low and high coverage firms are also quite large and statistically significant. Thus, the analyst coverage results are consistent with the information hypothesis.

## The effect of idiosyncratic volatility on local advantage

We further use idiosyncratic stock return volatility as a proxy for information asymmetry. A number of studies suggest that higher idiosyncratic volatility indicates a larger amount of firm-specific information not shared by the market, and therefore, greater information asymmetry (e.g., Bhagat, Marr, and (1985), Blackwell, Marr, and Spivey (1990), Krishnaswami and Subramaniam (1999), Zhang (2006)). Since small firms tend to have more volatile returns, we further construct size-adjusted volatility as the residual from cross-sectional regressions of idiosyncratic volatility on firm size. For each day of our sample period, we sort firms into three groups based on idiosyncratic volatility, and examine local advantage among high volatility firms (top tercile of volatility) and low volatility firms (bottom tercile of volatility), respectively.

Panel D of Table IX presents the results, which show that local advantage among high volatility firms is strong for all return windows examined. For example, the local advantage is 38.3 basis points (t-stat 2.66) in the five-day return window, both economically and statistically significant. On the contrary, local advantage among low volatility firms is insignificantly negative in all models. The differences in local advantage between high and low volatility groups are large and statistically significant. Therefore, the evidence from idiosyncratic volatility is also consistent with the information hypothesis.

To summarize, our results with all the proxies for information asymmetry consistently present a positive association between information asymmetry and local advantage. These results suggest that individual investors' local advantage is due to their access to value-relevant private information about local firms.

### 4.3 Performances of Long-Short Trading Strategies based on Local Advantage

In this subsection, we examine the profitability of zero-investment trading strategies based on local advantage. This examination is not only of interests to practitioners, but also helps verify the validity of our local advantage findings. We first construct a differential evaluation measure to capture the difference between local and nonlocal evaluations. Specifically, for a firm-day, we calculate the differential evaluation measure as the firm’s local evaluations in the previous two-week window minus the firm’s nonlocal evaluations in the previous two-week window. On each day of our sample period, we form a portfolio containing firms with negative differential evaluation measures (“negative portfolio”) and a portfolio containing firms with positive differential evaluation measures (“positive portfolio”). We then hold these portfolios for  $J$  days, where  $J=2, 5, 10$ , or  $20$ . This strategy is similar to the rolling momentum strategy proposed by Jegadeesh and Titman (1993) except that we form portfolios based on differential evaluations rather than momentum.

We report the average daily abnormal profits of these zero-investment strategies in Table X. Specifically, we first calculate for each day the difference in average abnormal returns between the two portfolios (“positive portfolio” – “negative portfolio”) and then report the time-series means. Daily abnormal returns are constructed based on the Fama and French 4 Factor model as defined in Section III. To control for time-series correlations, we report t-statistics using Newey-West robust standard errors with 30 lags.

We observe in Table X that the daily abnormal profits range from 4.6 basis points to 8.9 basis points for the two to twenty day holding windows, and are statistically significant except for the twenty-day window. These results provide strong evidence for the existence of local advantage. Since Section IV.B shows that local advantage is associated with firm characteristics including firm size, index identity, analyst coverage, and idiosyncratic volatil-

ity, we further examine the performance of the local-advantage-based strategy among small firms, Non-S&P 500 firms, low analyst coverage firms, or high idiosyncratic volatility firms (their classifications are defined in Section IV.B). We find that the daily abnormal profits significantly increase in these sub-samples. For example, among the low analyst coverage firms (bottom tercile of coverage), the daily abnormal profits range from 22.3 basis points to 27.2 basis points across different holding windows. These results are also consistent with our previous finding that local advantage is increasing in information asymmetry, suggesting that local advantage is likely caused by investors' access to private information about locally headquartered companies.

## 5 Conclusion

This paper investigates the local advantage of individual investors using a unique dataset of Twitter posts that cover publicly traded U.S. companies. While previous studies on individual investors' local advantage focus on the abnormal returns on investors' local investments, we directly examine individual investors' information about local and nonlocal companies. In addition, our sample is comprised of a group of notable and experienced Twitter users who may have enough knowledge and incentives to extract value-relevant information about local companies.

We first examine the overall stock-picking ability of the sample Twitter users. Interestingly, these individual investors, despite their extensive trading experience and knowledge about the stock market, exhibit significantly negative stock return predictability. We then contrast the stock return predictability between local and nonlocal investors and observe a large and significant local advantage. For example, when we examine weekly returns subsequent to investor evaluations, local advantage is 19 basis points per week, both economically and statistically significant. Further analyses show that local advantage concentrates in firms

without public news coverage, and firms with severe information asymmetry. These results indicate that local advantage is due to individual investors' access to private information about local firms.

We contribute to the debate on whether local advantage can exist for individual investors, one of the major participants in the stock market. While the recent work by Seasholes and Zhu (2010) finds little evidence that individual investors earn abnormal returns on their local investments, we directly examine investors' information set and document a significant local advantage for a group of relatively experienced and sophisticated individual investors. Together with Seasholes and Zhu, our results suggest possible heterogeneity in local advantage among individual investors, and the possibility that individual investors may fail to convert their value-relevant information about local firms into trading profits.

Our findings also have interesting implications for the rapidly growing internet communication about financial markets. Many people perceive that internet messages on the stock markets simply contain noise or reflect investor sentiment that is unrelated to firm fundamentals. We find that, indeed, even the seemingly trustworthy Twitter posts from notable contributors have a large negative return predictability. However, we also observe that local posts significantly outperform nonlocal posts and such advantage seems to result from contributors possessing private information about local firms. This finding suggests that internet communication about the financial markets contains value-relevant information as opposed to all noise.

## 6 References

Agarwal, Sumit and Robert Hauswald, 2010, Distance and private information in lending, *Review of Financial Studies* 23, 2757-2788.

Almazan, Andres, Adolfo De Motta, Sheridan Titman and Vahap Uysal, 2010, Financial structure, acquisition opportunities, and firm locations, *Journal of Finance* 65, 529-563.

Antweiler, Werner, and Murray Z. Frank, 2004, Is all that talk just noise? The information content of internet stock message boards, *Journal of Finance* 59, 1259-1294.

Ayers, Benjamin, Santhosh Ramalingegowda and Eric Yeung, 2011, Hometown advantage: The Effects of monitoring institution location on financial reporting discretion, *Journal of Accounting and Economic*, forthcoming.

Bae, Kee-Hong, Rene M. Stulz, and Hongping Tan, 2008, Do local analysts know more? A cross-country study of the performance of local analysts and foreign analysts, *Journal of Financial Economics* 88, 581-606.

Baik, Kang, Jun-Koo Kang and Jin-Mo Kim, 2010, Local institutional investors, information asymmetries, and equity returns, *Journal of Financial Economics* 97, 81-106.

Barber, Brad M., and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* 55, 773-806.

Bernile, Gennaro, Alok Kumar, and Johan Sulaeman, 2010, Home away from home: economic relevance and local investors, Working Paper, University of Miami.

Bhagat, Sanjai, Wayne M. Marr, and Rodney Thomson, 1985, The Rule 415 experiment: Equity markets, *Journal of Finance* 85, 1385-1401.

Blackwell, David W., Wayne M. Marr, and Michael F. Spivey, 1990, Shelf registration and the reduced due diligence argument: Implications of the underwriter certification and the implicit insurance hypotheses, *Journal of Financial & Quantitative Analysis* 25, 245-259.

- Brennan, Michael J., and Avanidhar Subrahmanyam, 1995, Investment analysis and price formation in securities markets, *Journal of Financial Economics* 38, 361-381.
- Butler, Alexander W., 2008, Distance still matters: Evidence from municipal bond underwriting, *Review of Financial Studies* 21, 763-784.
- Cao, Henry H., Joshua D. Coval, and David Hirshleifer, 2002, Sidelined investors, trade-generated news, and security returns, *Review of Financial Studies* 15, 615-648.
- Chan, Wesley S., 2003, Stock price reaction to news and no-news: Drift and reversal after headlines, *Journal of Financial Economics* 70, 223-260.
- Colla, Paolo, and Antonio Mele, 2010. Information linkages and correlated trading, *Review of Financial Studies* 23, 203-246.
- Coval, Joshua D., and Tobias J. Moskowitz, 1999, Home bias at home: Local equity preference in domestic portfolios, *Journal of Finance* 54, 2045-2073.
- Coval, Joshua D., and Tobias J. Moskowitz, 2001, The geography of investment: informed trading and asset prices, *Journal of Political Economy* 109, 811-41.
- Das, Sanjiv, Asis Martinez-Jerez, and Peter Tufano, 2005, eInformation: A clinical study of investor discussion and sentiment, *Financial Management* 34, 103-137.
- DeMarzo, Peter, Dimitri Vayanos, Jeffrey Zwiebel, 2003, Persuasion bias, social influence, and unidimensional opinions, *Quarterly Journal of Economics* 118, 909-968.
- Driscoll, John and Aart Kraay, 1998, Consistent covariance matrix estimation with spatially dependent panel data, *Review of Economics and Statistics* 80, 549-560
- Duflo, Esther, and Emmanuel Saez, 2002, Participation and investment decisions in a retirement plan: The influence of colleagues' choices, *Journal of Public Economics*, 85,121-148.

- Engelberg, Joseph, 2008, Costly information processing: Evidence from earnings announcements, Working paper, University of North Carolina.
- Gaspar, Jose-Miguel and Massimo Massa, 2007, Local ownership as private information: Evidence on the monitoring-liquidity trade-off, *Journal of Financial Economics* 83, 751-792.
- Fama, Eugene F., Lawrence Fisher, Michael C. Jensen, and Richard Roll, 1969, The adjustment of stock prices to new information. *International Economic Review* 10, 1-21.
- Fang, Lily H., and Joel Peress, 2011, Media coverage and the cross-section of stock returns, *Journal of Finance*, *forthcoming*.
- Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan, 2001, Toward an implied cost of capital, *Journal of Accounting Research* 39, 135-176.
- Griffin, John M., and Lemmon, Michael L., 2002, Book-to-market equity, distress risk, and stock returns, *Journal of Finance* 57, 2317-2336.
- Griffin, John M., Tao Shu and Selim Topaloglu, 2011, Examining the dark side of financial markets: Do institutions trade on information from investment banks connections? Working Paper, University of Texas at Austin.
- Grinblatt, Mark and Bing Han, 2005, Prospect theory, mental accounting, and momentum, *Journal of Financial Economics* 78, 311-339.
- Hong, Dong, Harrison Hong, and Andrei Ungureanu, 2011, An epidemiological approach to opinion and asset price-volume dynamics, Working Paper, Princeton University.
- Hong, Harrison, Jeffery Kubik, and Jeremy Stein, 2004, Social interactions and stock market participation, *Journal of Finance* 59, 137 – 163.

Hong, Harrison, Jeffery Kubik, and Jeremy Stein, 2005, The neighbor's portfolio: Word-of-mouth effects in the holdings and trades of money managers, *Journal of Finance* 60, 2801–2824.

Hong, Harrison, Terence Lim and Jeremy C. Stein., 2000, Bad News Travels Slowly: Size, Analyst Coverage, And The Profitability of Momentum Strategies, *Journal of Finance* 55, 265-295

Irvine, Paul J., 2004, Analysts' forecasts and brokerage-firm trading, *Accounting Review* 79, 125-149.

Ivkovi, Zoran, and Scott Weisbenner, 2005, Local does as local is: Information content of the geography of individual investors' common stock investments, *Journal of Finance* 60, 267-306.

Ivkovi, Zoran, and Scott Weisbenner, 2007, Information diffusion effects in individual investors' common stock purchases: Covet thy neighbors' investment choices, *Review of Financial Studies* 20, 1327-1357.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65-91.

Kaniel, Ron, Gideon Saar, and Sheridan Titman, 2008, Individual investor trading and stock returns, *Journal of Finance* 2008, 273-310.

Kaniel, Ron, Gideon Saar, Sheridan Titman, and Shuming Liu, 2011, Individual investor trading and return patterns around earnings announcements, *Journal of Finance*, *forthcoming*.

Kelley, Eric, and Paul Tetlock, 2011, How wise are crowds? Insights from retail orders and stock returns, Working Paper, Columbia University.

Krishnaswami, Sudha and Venkat Subramaniam, 1999, Information asymmetry, valuation, and the corporate spin-off decision, *Journal of Financial Economics* 53, 73-112.

Loughran, Tim and Bill McDonald, 2011, When is liability not a liability? Textual Analysis, Dictionaries, and 10-Ks, *Journal Finance* 66, 35-65.

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

Odean, Terrance, 1998, Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775-1798.

Pang, Bo, Lillian Lee, and Shivakumar Vaithyanathan, 2002, Thumbs up? Sentiment classification using machine learning techniques, Proceedings, *ACL-02 Conference on Empirical methods in natural language*.

Seasholes, Mark S., and Ning Zhu, 2010, Individual investors and local bias, *Journal of Finance* 65, 1987-2010.

Tetlock, Paul C., 2007, Giving content to investor sentiment: The role of media in the stock market, *Journal of Finance* 62, 1139–1168.

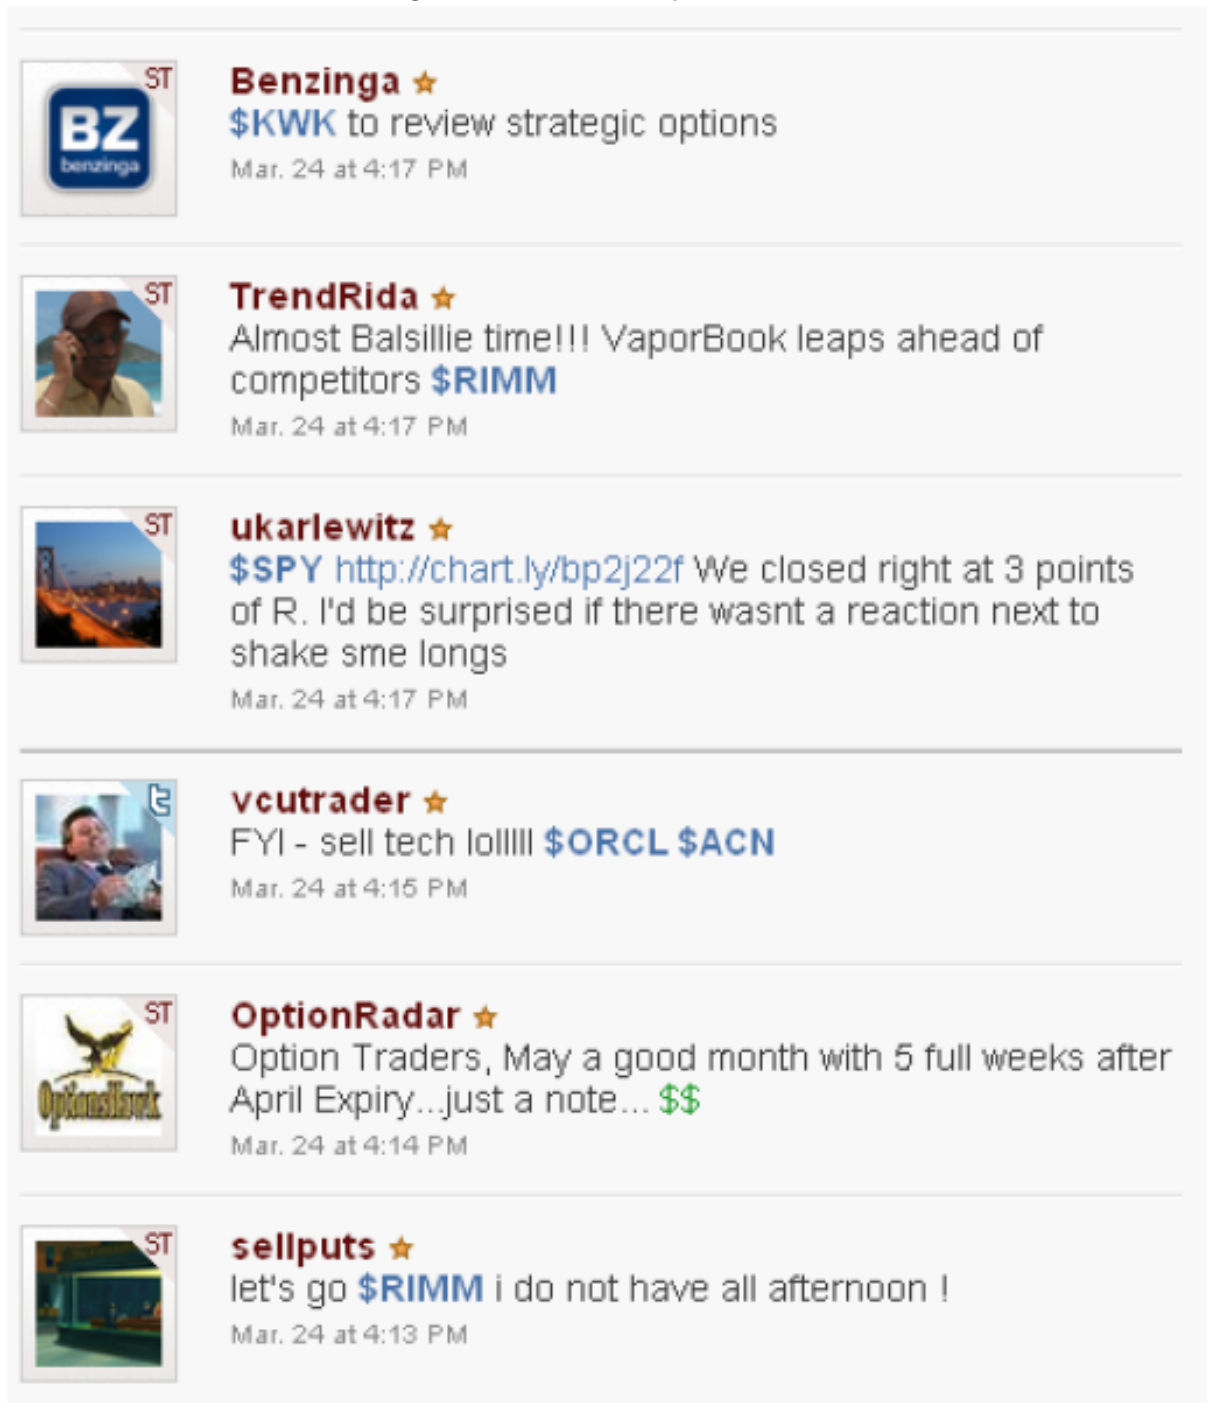
Tkac, Paula A., 1999, A trading volume benchmark: theory and evidence, *Journal of Financial and Quantitative Analysis* 34, 89-114.

Tumarkin, Robert, and Robert F. Whitelaw, 2001, News or noise? Internet postings and stock prices, *Financial Analysts Journal* 57, 41-51.

Figure 3.1: Example of Twitter Stream

This figure shows the interface that a Stocktwits.com user will see. Company tickers can be seen in blue after the \$ hashtags.

Figure 3.2: Summary Statistics



**Figure 3.3: Geographical Distribution of Sample Twitter Users**

This figure plots the geographical distribution of Twitter users in our sample. Each arrow represents a Twitter user.



Table 3.1: **Summary Statistics**

Panel A reports summary statistics for the 1,184 firms in our sample from October 1st, 2008 to December 31st, 2009. Market capitalization is measured at the end of 2007. Book-to-market ratio is book equity divided by market capitalization at the end of 2007 fiscal year. Analyst coverage is the number of analysts covering the firm as of September 2007. Idiosyncratic volatility for a firm-day is the standard error of residuals from time-series regressions of the firms excess returns on the market excess returns (MKT) in the previous 150-day rolling window. Average daily return is the average of daily raw returns of sample firms during our sample period. Panel B reports average firm characteristics for the subsets of firms sorted into two groups according to the numbers of all Twitter posts, local posts, and nonlocal posts, respectively. A Twitter user is local (nonlocal) to a firm if the users location is less than (more than) 100 miles from the firms headquarters.

| <b>Panel A: Characteristics of Sample Firms</b>                             |                  |             |                    |             |                       |             |            |
|-----------------------------------------------------------------------------|------------------|-------------|--------------------|-------------|-----------------------|-------------|------------|
|                                                                             | <b>Mean</b>      | <b>STD</b>  | <b>P10</b>         | <b>P25</b>  | <b>P50</b>            | <b>P75</b>  | <b>P90</b> |
| Market Capitalization (\$M)                                                 | 5,981            | 21,499      | 129                | 324         | 942                   | 3,157       | 12,247     |
| Book/Market Ratio                                                           | 0.40             | 0.38        | 0.11               | 0.20        | 0.33                  | 0.54        | 0.79       |
| Analyst Coverage                                                            | 8.5              | 6.5         | 1.0                | 4.0         | 7.0                   | 12.0        | 18.0       |
| Idiosyncratic Volatility                                                    | 0.035            | 0.015       | 0.019              | 0.025       | 0.033                 | 0.043       | 0.054      |
| Average Daily Ret. (%)                                                      | 0.210            | 0.226       | 0.016              | 0.094       | 0.178                 | 0.307       | 0.463      |
| <b>Panel B: Average Firm Characteristics Sorted on the Numbers of Posts</b> |                  |             |                    |             |                       |             |            |
|                                                                             | <b>All Posts</b> |             | <b>Local Posts</b> |             | <b>Nonlocal Posts</b> |             |            |
|                                                                             | <b>Low</b>       | <b>High</b> | <b>Low</b>         | <b>High</b> | <b>Low</b>            | <b>High</b> |            |
| Market Capitalization (\$M)                                                 | 3,070            | 8,787       | 4,370              | 11,680      | 3,306                 | 8,551       |            |
| Book/Market Ratio                                                           | 0.43             | 0.38        | 0.41               | 0.35        | 0.43                  | 0.38        |            |
| Analyst Coverage                                                            | 6.2              | 10.68       | 8.1                | 10.1        | 6.4                   | 10.5        |            |
| Idiosyncratic Volatility                                                    | 0.034            | 0.036       | 0.035              | 0.036       | 0.034                 | 0.037       |            |
| Average Daily Ret.                                                          | 0.19             | 0.23        | 0.21               | 0.22        | 0.21                  | 0.24        |            |

Table 3.2: **Panel Regressions of Abnormal Stock Returns on Investor Evaluations**

Panel A presents panel regressions of abnormal stock returns on the previous local and nonlocal evaluations. The dependent variables are two-, five-, ten-, and twenty-day cumulative abnormal returns (measured in percentage), respectively. To calculate daily abnormal return for a firm-day, we first estimate a Fama-French 4 Factor regression for the firm in the previous 150-day rolling window, and then use the estimated factor loadings to calculate abnormal returns for the firm-day. The independent variables include investor evaluations in the two-week windows prior to return measurements. To calculate investor evaluations, we first use the maximum entropy classification to measure the evaluation of each post, and then sum the evaluation measures of the posts in the two weeks prior to return measurement. All regressions include firm fixed effects with lagged returns in the previous ten trading days as controls. Panel B is similar to Panel A but skip one week before the return measurement. Panel C is similar to Panel A but with one-week evaluations rather than two-week evaluations. T-statistics (reported in parentheses) are calculated using Driscoll and Kraay (1998) robust standard errors that control for both cross-sectional and time-series correlations. \*\*\*, \*\*, and \* represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

|                                                                          | Dependant Variables  |                      |                      |                      |
|--------------------------------------------------------------------------|----------------------|----------------------|----------------------|----------------------|
|                                                                          | 2-Day CAR            | 5-Day CAR            | 10-Day CAR           | 20-Day CAR           |
| <b>Panel A: Regressions on Two-Week Evaluations</b>                      |                      |                      |                      |                      |
| Two-Week Evaluation                                                      | -0.060***<br>(-3.27) | -0.138***<br>(-3.05) | -0.249***<br>(-3.06) | -0.465***<br>(-4.33) |
| Controls of Lagged Returns                                               | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                                       | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                            | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                                                        | 1,184                | 1,184                | 1,184                | 1,184                |
| <b>Panel B: Regressions on Two-Week Evaluations: Skip-a-Week Returns</b> |                      |                      |                      |                      |
| Two-Week Evaluation                                                      | -0.043**<br>(-2.39)  | -0.111***<br>(-2.75) | -0.223***<br>(-3.67) | -0.447***<br>(-4.80) |
| Controls of Lagged Returns                                               | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                                       | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                            | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                                                        | 1,184                | 1,184                | 1,184                | 1,184                |
| <b>Panel C: Regressions on One-Week Evaluations</b>                      |                      |                      |                      |                      |
| One-Week Evaluation                                                      | -0.057***<br>(-4.35) | -0.115***<br>(-3.27) | -0.227***<br>(-3.03) | -0.397***<br>(-3.89) |
| Controls of Lagged Returns                                               | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                                       | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                            | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                                                        | 1,184                | 1,184                | 1,184                | 1,184                |

Table 3.3: **Panel Regressions of Stock Returns on Local and Nonlocal Evaluations**

Panel A presents panel regressions of stock returns on previous local and nonlocal evaluations. The dependent variables are cumulative two-, five-, ten- and twenty-day abnormal returns (measured in percentage), respectively. To calculate daily abnormal return for a firm-day, we first estimate a Fama-French 4 Factor regression for the firm in the previous 150-day rolling window, and then use the estimated factor loadings to calculate abnormal returns for the firm-day. The independent variables include local evaluation and nonlocal evaluation in the two-week window prior to return measurement. To calculate local and nonlocal evaluations, we first classify Twitter posts into local and nonlocal posts according to whether the Twitter users locations are within 100 miles of the headquarters of the firms mentioned in the posts. We use maximum entropy classification to measure the evaluation of each post, and then sum the evaluation measures of the local and nonlocal posts, respectively, in the two weeks prior to return measurement. For each regression, we further report the difference between the coefficients on local evaluation and nonlocal evaluation. All regressions include firm fixed effects with lagged returns in the previous ten trading days as controls. T-statistics (reported in parentheses) are calculated using Driscoll and Kraay (1998) robust standard errors that control for both cross-sectional and time-series correlations. Panel B is similar to Panel A except that the dependent variables are raw returns instead of abnormal returns. \*\*\*, \*\*, and \* represent statistical significance at the 0.01, 0.05, and 0.10 levels, respectively.

| <b>Panel A: Regressions of Abnormal Returns</b> |                      |                      |                      |                      |
|-------------------------------------------------|----------------------|----------------------|----------------------|----------------------|
|                                                 | Dependant Variables  |                      |                      |                      |
|                                                 | 2-Day CAR            | 5-Day CAR            | 10-Day CAR           | 20-Day CAR           |
| Local Evaluation                                | 0.003<br>(0.12)      | 0.022<br>(0.36)      | 0.056<br>(0.49)      | -0.023<br>(-0.18)    |
| Nonlocal Evaluation                             | -0.068***<br>(-3.52) | -0.165***<br>(-3.63) | -0.307***<br>(-3.74) | -0.500***<br>(-4.04) |
| Local – Nonlocal                                | 0.070**<br>(2.23)    | 0.187**<br>(2.31)    | 0.363**<br>(2.43)    | 0.477**<br>(2.42)    |
| Controls of Lagged Returns                      | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                              | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                   | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                               | 1,184                | 1,184                | 1,184                | 1,184                |
| <b>Panel B: Regressions of Raw Returns</b>      |                      |                      |                      |                      |
|                                                 | Dependant Variables  |                      |                      |                      |
|                                                 | 2-Day Ret.           | 5-Day Ret.           | 10-Day Ret.          | 20-Day Ret.          |
| Local Evaluation                                | 0.026<br>(1.05)      | 0.082<br>(1.19)      | 0.141<br>(1.04)      | 0.096<br>(0.51)      |
| Nonlocal Evaluation                             | -0.060**<br>(-1.85)  | -0.135*<br>(-1.76)   | -0.260*<br>(-1.94)   | -0.500**<br>(-2.23)  |
| Local – Nonlocal                                | 0.086**<br>(2.40)    | 0.217**<br>(2.47)    | 0.401**<br>(2.58)    | 0.596***<br>(2.59)   |
| Controls of Lagged Returns                      | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                              | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                   | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                               | 1,184                | 1,184                | 1,184                | 1,184                |

Table 3.4: **Panel Regressions of Stock Returns: Alternative Construction/ Sample Selection**

Panel A presents the regressions of abnormal returns on local and nonlocal evaluations in the one-week window prior to return measurements. The definition of abnormal returns, local and nonlocal evaluations, and regression settings are similar to the Panel A of Table III. Panel B presents the regressions of abnormal returns on local and nonlocal evaluations in the two-week window prior to return measurements. These regressions are similar to the Panel A of Table III except that we classify Twitter posts into local and nonlocal according whether the users and the company headquarters locate in the same state. Panel C presents the regressions of abnormal returns similar to the Panel A of Table III except that we only include firms that have at least one Twitter post in the two-week period of evaluation measurement. All regressions include firm fixed effects and lagged returns of the previous ten trading days as controls. T-statistics (reported in parentheses) are calculated with the Driscoll and Kraay (1998) robust standard errors that control for both cross-sectional and time-series correlations. \*\*\*, \*\*, and \* represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

|                                                                                | Dependant Variables  |                      |                      |                      |
|--------------------------------------------------------------------------------|----------------------|----------------------|----------------------|----------------------|
|                                                                                | 2-Day CAR            | 5-Day CAR            | 10-Day CAR           | 20-Day CAR           |
| <b>Panel A: Return Regressions: Evaluations Measured in the Past One Week</b>  |                      |                      |                      |                      |
| Local Evaluation                                                               | -0.005<br>(-0.38)    | -0.002<br>(-0.06)    | 0.036<br>(0.36)      | 0.005<br>(0.04)      |
| Nonlocal Evaluation                                                            | -0.056***<br>(-3.58) | -0.116***<br>(-3.12) | -0.250***<br>(-3.75) | -0.407***<br>(-3.89) |
| Local – Nonlocal                                                               | 0.051**<br>(2.08)    | 0.114**<br>(1.99)    | 0.286**<br>(2.29)    | 0.411**<br>(2.47)    |
| Controls of Lagged Returns                                                     | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                                             | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                                  | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                                                              | 1,184                | 1,184                | 1,184                | 1,184                |
| <b>Panel B: Return Regressions: Local Classification based on State</b>        |                      |                      |                      |                      |
| Local Evaluation                                                               | 0.008<br>(0.37)      | 0.034<br>(0.57)      | 0.087<br>(0.79)      | 0.038<br>(0.29)      |
| Nonlocal Evaluation                                                            | -0.070***<br>(-3.50) | -0.170***<br>(-3.54) | -0.322***<br>(-3.61) | -0.530***<br>(-3.93) |
| Local – Nonlocal                                                               | 0.078**<br>(2.42)    | 0.204**<br>(2.41)    | 0.409**<br>(2.56)    | 0.568**<br>(2.56)    |
| Controls of Lagged Returns                                                     | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                                             | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                                  | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                                                              | 1,184                | 1,184                | 1,184                | 1,184                |
| <b>Panel C: Return Regressions: At Least One Post in the Evaluation Period</b> |                      |                      |                      |                      |
| Local Evaluation                                                               | 0.006<br>(0.25)      | 0.033<br>(0.52)      | 0.083<br>(0.70)      | 0.026<br>(0.19)      |
| Nonlocal Evaluation                                                            | -0.058***<br>(-2.66) | -0.135***<br>(-2.92) | -0.244***<br>(-3.08) | -0.377***<br>(-3.16) |
| Local – Nonlocal                                                               | 0.064*<br>(1.85)     | 0.168**<br>(1.97)    | 0.327**<br>(2.01)    | 0.404*<br>(1.77)     |

|                            | Dependant Variables |           |            |            |
|----------------------------|---------------------|-----------|------------|------------|
|                            | 2-Day CAR           | 5-Day CAR | 10-Day CAR | 20-Day CAR |
| Controls of Lagged Returns | Yes                 | Yes       | Yes        | Yes        |
| Firm Fixed Effects         | Yes                 | Yes       | Yes        | Yes        |
| Number of Obs              | 67,270              | 67,270    | 67,270     | 67,270     |
| Number of PERMNOs          | 1,184               | 1,184     | 1,184      | 1,184      |

**Table 3.5: Panel Regressions of Stock Returns: State Adjusted Returns**

Panel A presents the regressions of state-adjusted abnormal returns on local and nonlocal evaluations in the two-week period prior to return measurements. The dependent variables are two-, five-, ten- or twenty-day cumulative state-adjusted abnormal returns. We calculate a firms daily state-adjusted abnormal return as the firms daily abnormal return minus the average daily abnormal returns of all firms in the same state as the firm. The independent variables and regression settings are defined in the heading of Table III. Panel B repeats regressions in Panel A but with state-adjusted raw returns. Daily state adjusted return for a firm is calculated as the daily return of the firm minus the average daily return of all firms in the same state as the firm. All regressions include firm fixed effects with lagged returns of the previous ten trading days of as controls. T-statistics (reported in parentheses) are calculated with Driscoll and Kraay (1998) robust standard errors that control for both cross-sectional and time-series correlations. \*\*\*, \*\*, and \* represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

| <b>Panel A: Regressions of State-Adjusted Abnormal Returns</b> |                      |                      |                      |                      |
|----------------------------------------------------------------|----------------------|----------------------|----------------------|----------------------|
|                                                                | Dependant Variables  |                      |                      |                      |
|                                                                | 2-Day CAR            | 5-Day CAR            | 10-Day CAR           | 20-Day CAR           |
| Local Evaluation                                               | 0.022<br>(1.01)      | 0.072<br>(1.09)      | 0.138<br>(1.17)      | 0.100<br>(0.58)      |
| Nonlocal Evaluation                                            | -0.051***<br>(-2.98) | -0.121***<br>(-3.01) | -0.212***<br>(-2.88) | -0.399**<br>(-2.33)  |
| Local – Nonlocal                                               | 0.073**<br>(2.55)    | 0.193***<br>(2.60)   | 0.351**<br>(2.47)    | 0.488**<br>(2.23)    |
| Controls of Lagged Returns                                     | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                             | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                  | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                                              | 1,184                | 1,184                | 1,184                | 1,184                |
| <b>Panel B: Regressions of State-Adjusted Raw Returns</b>      |                      |                      |                      |                      |
|                                                                | Dependant Variables  |                      |                      |                      |
|                                                                | 2-Day Ret.           | 5-Day Ret.           | 10-Day Ret.          | 20-Day Ret.          |
| Local Evaluation                                               | 0.003<br>(0.16)      | 0.024<br>(0.40)      | 0.056<br>(0.50)      | -0.020<br>(-0.16)    |
| Nonlocal Evaluation                                            | -0.062***<br>(-3.21) | -0.152***<br>(-3.51) | -0.282***<br>(-3.59) | -0.452***<br>(-3.89) |
| Local – Nonlocal                                               | 0.066**<br>(2.12)    | 0.176**<br>(2.23)    | 0.339**<br>(2.34)    | 0.432**<br>(2.26)    |
| Controls of Lagged Returns                                     | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                             | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                  | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                                              | 1,184                | 1,184                | 1,184                | 1,184                |

Table 3.6: **Panel Regressions of Stock Returns: Alternative Measurements of Evaluations or Returns**

Panel A presents regressions of abnormal returns on local and nonlocal evaluations in the two-week period ending one week before the return measurements. The constructions of abnormal returns, investor evaluations, and regression settings are defined in the heading of Table III. Panel B presents regressions of abnormal returns on the weighted evaluation measures in the two-week period before return measurement. Specifically, during the two-week period before return measurement, we multiply the evaluation of each Twitter post by the number of followers of the Twitter user, and then sum up the weighted evaluations. All regressions include firm fixed effects with lagged returns of the previous ten trading days as controls. T-statistics (reported in parentheses) are calculated with the Driscoll and Kraay (1998) robust standard errors that control for both cross-sectional and time-series correlations. \*\*\*, \*\*, and \* represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

| <b>Panel A: Regressions of State-Adjusted Abnormal Returns</b> |                      |                      |                      |                      |
|----------------------------------------------------------------|----------------------|----------------------|----------------------|----------------------|
|                                                                | Dependant Variables  |                      |                      |                      |
|                                                                | 2-Day CAR            | 5-Day CAR            | 10-Day CAR           | 20-Day CAR           |
| Local Evaluation                                               | 0.022<br>(1.01)      | 0.072<br>(1.09)      | 0.138<br>(1.17)      | 0.100<br>(0.58)      |
| Nonlocal Evaluation                                            | -0.051***<br>(-2.98) | -0.121***<br>(-3.01) | -0.212***<br>(-2.88) | -0.399**<br>(-2.33)  |
| Local – Nonlocal                                               | 0.073**<br>(2.55)    | 0.193***<br>(2.60)   | 0.351**<br>(2.47)    | 0.488**<br>(2.23)    |
| Controls of Lagged Returns                                     | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                             | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                  | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                                              | 1,184                | 1,184                | 1,184                | 1,184                |
| <b>Panel B: Regressions of State-Adjusted Raw Returns</b>      |                      |                      |                      |                      |
|                                                                | Dependant Variables  |                      |                      |                      |
|                                                                | 2-Day Ret.           | 5-Day Ret.           | 10-Day Ret.          | 20-Day Ret.          |
| Local Evaluation                                               | 0.003<br>(0.16)      | 0.024<br>(0.40)      | 0.056<br>(0.50)      | -0.020<br>(-0.16)    |
| Nonlocal Evaluation                                            | -0.062***<br>(-3.21) | -0.152***<br>(-3.51) | -0.282***<br>(-3.59) | -0.452***<br>(-3.89) |
| Local – Nonlocal                                               | 0.066**<br>(2.12)    | 0.176**<br>(2.23)    | 0.339**<br>(2.34)    | 0.432**<br>(2.26)    |
| Controls of Lagged Returns                                     | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                             | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                  | 317,583              | 317,583              | 317,583              | 317,583              |
| Number of PERMNOs                                              | 1,184                | 1,184                | 1,184                | 1,184                |

Table 3.7: Panel Regressions of Trading Volumes

Panel A presents regressions of turnovers on local and nonlocal evaluations in the two-week period prior to turnover measurements. The independent variables are two-, five-, ten- or twenty-day cumulative abnormal turnovers. Daily turnover is a firms daily trading volume scaled by total shares outstanding. We obtain daily excess turnover by subtracting cross-sectional average turnover of the CRSP universe and then calculate abnormal turnover for a firm-day by subtracting average daily excess turnover of the firm in the previous 180-day rolling window. Investor evaluations are defined in the heading of Table III. Panel B is similar to Panel A but the independent variables are absolute values of the local and nonlocal evaluation measures. All regressions include firm fixed effects with lagged returns of the previous ten trading days as controls. T-statistics (reported in parentheses) are calculated with the Driscoll and Kraay (1998) robust standard errors that control for both cross-sectional and time-series correlations. \*\*\*, \*\*, and \* represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

|                                                                                            | Dependant Variables |                    |                    |                    |
|--------------------------------------------------------------------------------------------|---------------------|--------------------|--------------------|--------------------|
|                                                                                            | 2-Day Vol.          | 5-Day Vol.         | 10-Day Vol.        | 20-Day Vol.        |
| <b>Panel A: Regressions of Trading Volumes on Local and Nonlocal evaluations</b>           |                     |                    |                    |                    |
| Local Evaluation                                                                           | 0.027***<br>(4.42)  | 0.022***<br>(3.62) | 0.017***<br>(2.67) | 0.013***<br>(2.59) |
| Nonlocal Evaluation                                                                        | 0.033***<br>(2.98)  | 0.029***<br>(2.80) | 0.023**<br>(2.35)  | 0.015*<br>(1.73)   |
| Local – Nonlocal                                                                           | -0.006<br>(-0.46)   | -0.007<br>(-0.55)  | -0.005<br>(-0.48)  | -0.002<br>(-0.22)  |
| Controls of Lagged Returns                                                                 | Yes                 | Yes                | Yes                | Yes                |
| Firm Fixed Effects                                                                         | Yes                 | Yes                | Yes                | Yes                |
| Number of Obs                                                                              | 271,573             | 271,573            | 271,573            | 271,573            |
| Number of PERMNOs                                                                          | 1,064               | 1,064              | 1,064              | 1,064              |
| <b>Panel B: Regressions of Trading Volumes on Un-Signed Local and Nonlocal Evaluations</b> |                     |                    |                    |                    |
| Local Evaluation                                                                           | 0.044***<br>(5.07)  | 0.037***<br>(3.53) | 0.030***<br>(2.79) | 0.025***<br>(2.69) |
| Nonlocal Evaluation                                                                        | 0.034***<br>(3.07)  | 0.029***<br>(2.73) | 0.021**<br>(2.14)  | 0.014*<br>(1.50)   |
| Local  –  Nonlocal                                                                         | 0.010<br>(0.70)     | 0.009<br>(0.58)    | 0.009<br>(0.65)    | 0.011<br>(1.01)    |
| Controls of Lagged Returns                                                                 | Yes                 | Yes                | Yes                | Yes                |
| Firm Fixed Effects                                                                         | Yes                 | Yes                | Yes                | Yes                |
| Number of Obs                                                                              | 271,573             | 271,573            | 271,573            | 271,573            |
| Number of PERMNOs                                                                          | 1,064               | 1,064              | 1,064              | 1,064              |

Table 3.8: **Panel Regressions of Stock Returns: Stocks without Public News vs. Stocks with Public News**

This table reports regressions of two-, five-, ten-, or twenty-day abnormal returns on local and nonlocal evaluations in the two-week period prior to return measurement for stocks with and without public news coverage, respectively. We collect news articles from PR News Wire, Dow Jones News Wire, and Reuters News and classify stocks into two groups based whether they have news coverage in the two-week period of evaluation measurement. We then estimate regressions for the non-news firms in Panel A and for the news firms in Panel B. We further report the difference in local advantage between no-news and news samples in Panel C. All regressions include firm fixed effects with lagged returns of the previous ten trading days as controls. T-statistics (reported in parentheses) are calculated with Driscoll and Kraay (1998) robust standard errors that control for both cross-sectional and time-series correlations. \*\*\*, \*\*, and \* represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

|                                                                                    | Dependent Variables  |                      |                      |                      |
|------------------------------------------------------------------------------------|----------------------|----------------------|----------------------|----------------------|
|                                                                                    | 2-Day CAR            | 5-Day CAR            | 10-Day CAR           | 20-Day CAR           |
| <b>Panel A: Stocks without Public News Coverage in the Period of Twitter Posts</b> |                      |                      |                      |                      |
| Local Evaluation                                                                   | 0.041**<br>(1.98)    | 0.113**<br>(1.88)    | 0.140<br>(1.15)      | 0.002<br>(0.01)      |
| Nonlocal Evaluation                                                                | -0.101***<br>(-4.24) | -0.251***<br>(-4.53) | -0.460***<br>(-4.07) | -0.791***<br>(-3.90) |
| Local - Nonlocal (1)                                                               | 0.143***<br>(4.04)   | 0.364***<br>(3.95)   | 0.600***<br>(3.26)   | 0.793***<br>(2.83)   |
| Controls of Lagged Returns                                                         | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                                                 | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                                      | 243,157              | 243,157              | 243,157              | 243,157              |
| Number of PERMNOs                                                                  | 1,165                | 1,165                | 1,165                | 1,165                |
| <b>Panel B: Stocks with Public News Coverage in the Period of Twitter Posts</b>    |                      |                      |                      |                      |
| Local Evaluation                                                                   | -0.089*<br>(-1.72)   | -0.195<br>(-1.58)    | -0.144<br>(-0.73)    | -0.073<br>(-0.32)    |
| Nonlocal Evaluation                                                                | -0.047**<br>(-2.20)  | -0.113**<br>(-2.18)  | -0.221**<br>(-2.47)  | -0.338***<br>(-2.63) |
| Local - Nonlocal (2)                                                               | -0.042<br>(-0.68)    | -0.082<br>(-0.57)    | 0.077<br>(0.32)      | 0.264<br>(0.87)      |
| Controls of Lagged Returns                                                         | Yes                  | Yes                  | Yes                  | Yes                  |
| Firm Fixed Effects                                                                 | Yes                  | Yes                  | Yes                  | Yes                  |
| Number of Obs                                                                      | 74,426               | 74,426               | 74,426               | 72,426               |
| Number of PERMNOs                                                                  | 457                  | 457                  | 457                  | 457                  |
| <b>Panel C: Difference in Local Advantage: Non-News versus News Stocks</b>         |                      |                      |                      |                      |
| (1) - (2)                                                                          | 0.185***<br>(2.61)   | 0.447***<br>(2.59)   | 0.523*<br>(1.73)     | 0.528<br>(1.28)      |

**Table 3.9: Local Advantage across Proxies of Information Asymmetry**

Panel A reports local advantage for small and large firms. On each day of our sample period, we sort stocks into three groups based their market capitalizations. We then estimate regressions of abnormal returns as in the Panel A of Table III for small firms (lowest tercile of market capitalization) and large firms (highest tercile of market capitalization), respectively. We then report local advantage (Local – Nonlocal in the Panel A of Table III) for small firms, large firms, and their differences. For Panel B, on each day of our sample period, we sort stocks into two groups based whether they are in the S&P 500 Composite Index at the beginning of the month. We then report local advantage for non-S&P 500 firms, S&P 500 firms, and their differences. For Panel C, on each day of our sample period, we sort stocks into three groups based on size-adjusted analyst coverage, where size-adjusted analyst coverage is residual from cross-sectional regression of analyst coverage on size. We then report local advantage for low coverage firms (lowest tercile of coverage), high coverage firms (highest tercile of coverage), and their differences. For Panel D, on each day of our sample period, we sort stocks into three groups based on size-adjusted idiosyncratic volatility. Idiosyncratic volatility for a firm-day is standard deviation of the residuals from the time-series regression of daily stock returns on the market factor (MKT) in the previous 150-day rolling window. Size-adjusted idiosyncratic volatility is residual from cross-sectional regression of idiosyncratic volatility on firm size. We then report local advantage for high volatility firms (highest tercile of volatility), low volatility firms (lowest tercile of volatility), and their differences. \*\*\*, \*\*, and \* represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

|                                                                            | Dependent Variables |                    |                    |                    |
|----------------------------------------------------------------------------|---------------------|--------------------|--------------------|--------------------|
|                                                                            | 2-Day CAR           | 5-Day CAR          | 10-Day CAR         | 20-Day CAR         |
| <b>Panel A: Local Advantage for Small versus Large Firms</b>               |                     |                    |                    |                    |
| Small Firms                                                                | 0.278***<br>(3.41)  | 0.620***<br>(3.54) | 1.123***<br>(3.79) | 1.903***<br>(3.54) |
| Large Firms                                                                | 0.168<br>(0.51)     | 0.080<br>(0.93)    | 0.188<br>(1.17)    | 0.146<br>(0.74)    |
| Small – Large                                                              | 0.261***<br>(2.97)  | 0.540***<br>(2.77) | 0.935***<br>(2.78) | 1.757***<br>(3.07) |
| <b>Panel B: Local Advantage for Non-S&amp;P versus S&amp;P Firms</b>       |                     |                    |                    |                    |
| Non-S&P Firms                                                              | 0.084***<br>(2.66)  | 0.214***<br>(2.64) | 0.393***<br>(2.57) | 0.507***<br>(2.47) |
| S&P Firms                                                                  | 0.024<br>(0.75)     | 0.095<br>(1.10)    | 0.206<br>(1.25)    | 0.147<br>(0.72)    |
| Non-S&P – S&P                                                              | 0.060<br>(1.34)     | 0.119<br>(1.00)    | 0.187<br>(0.83)    | 0.360<br>(1.24)    |
| <b>Panel C: Local Advantage for Low versus High Analyst Coverage Firms</b> |                     |                    |                    |                    |
| Low Coverage Firms                                                         | 0.173***<br>(4.17)  | 0.489***<br>(5.16) | 1.000***<br>(4.97) | 1.463***<br>(3.95) |
| High Coverage Firms                                                        | 0.076<br>(1.46)     | 0.158<br>(1.30)    | 0.261<br>(1.29)    | 0.356<br>(1.52)    |
| Low – High                                                                 | 0.097<br>(1.45)     | 0.331**<br>(2.13)  | 0.740**<br>(2.59)  | 1.107**<br>(2.53)  |

| Panel D: Local Advantage for High versus Low Idiosyncratic Volatility Firms |                    |                    |                    |                   |
|-----------------------------------------------------------------------------|--------------------|--------------------|--------------------|-------------------|
| High Idio. Volatility Firms                                                 | 0.158***<br>(2.73) | 0.383***<br>(2.66) | 0.700***<br>(2.68) | 1.016**<br>(2.56) |
| Low Idio. Volatility Firms                                                  | -0.068<br>(-0.99)  | -0.149<br>(-1.12)  | -0.196<br>(-0.96)  | -0.088<br>(-0.33) |
| High – Low                                                                  | 0.226**<br>(2.52)  | 0.532***<br>(2.71) | 0.896***<br>(2.70) | 1.103**<br>(2.31) |

Table 3.10: **Daily Abnormal Profits (%) of Rolling Long-Short Strategies Based on Local Advantage**

Panel A reports local advantage for small and large firms. On each day of our sample period, we sort stocks into three groups based their market capitalizations. We then estimate regressions of abnormal returns as in the Panel A of Table III for small firms (lowest tercile of market capitalization) and large firms (highest tercile of market capitalization), respectively. We then report local advantage (Local - Nonlocal in the Panel A of Table III) for small firms, large firms, and their differences. For Panel B, on each day of our sample period, we sort stocks into two groups based whether they are in the S&P 500 Composite Index at the beginning of the month. We then report local advantage for non-S&P 500 firms, S&P 500 firms, and their differences. For Panel C, on each day of our sample period, we sort stocks into three groups based on size-adjusted analyst coverage, where size-adjusted analyst coverage is residual from cross-sectional regression of analyst coverage on size. We then report local advantage for low coverage firms (lowest tercile of coverage), high coverage firms (highest tercile of coverage), and their differences. For Panel D, on each day of our sample period, we sort stocks into three groups based on size-adjusted idiosyncratic volatility. Idiosyncratic volatility for a firm-day is standard deviation of the residuals from the time-series regression of daily stock returns on the market factor (MKT) in the previous 150-day rolling window. Size-adjusted idiosyncratic volatility is residual from cross-sectional regression of idiosyncratic volatility on firm size. We then report local advantage for high volatility firms (highest tercile of volatility), low volatility firms (lowest tercile of volatility), and their differences. \*\*\*, \*\*, and \* represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

|                              | Hold 2 Days        | Hold 5 Days        | Hold 10 Days       | Hold 20 Days       |
|------------------------------|--------------------|--------------------|--------------------|--------------------|
| <b>Full Sample</b>           | 0.089**<br>(2.30)  | 0.073**<br>(2.01)  | 0.074*<br>(1.68)   | 0.046<br>(1.43)    |
| <b>Small Firms</b>           | 0.357***<br>(2.97) | 0.253**<br>(2.47)  | 0.283***<br>(2.84) | 0.144*<br>(1.86)   |
| <b>Non-S&amp;P Firms</b>     | 0.195***<br>(4.06) | 0.148***<br>(3.07) | 0.130**<br>(2.25)  | 0.091*<br>(1.79)   |
| <b>Low Coverage Firms</b>    | 0.223***<br>(3.29) | 0.240***<br>(3.27) | 0.275***<br>(2.76) | 0.272***<br>(3.51) |
| <b>High Volatility Firms</b> | 0.192***<br>(2.87) | 0.130*<br>(1.74)   | 0.114<br>(1.50)    | 0.083<br>(1.16)    |