The Convergence and Divergence of Investors' Opinions around Earnings News: Evidence from a Social Network

ROBERT GIANNINI

PAUL IRVINE

TAO SHU*

January 2015

*Giannini is at the BlueCrest Capital Management. Irvine is the C. R. Williams professor of finance at Texas Christian University. Shu is from the Terry College of Business at the University of Georgia and is visiting at HKUST Business School. We thank Linda Bamber, John Campbell, Bilal Erturk, Zoran Ivkovic, Andy Puckett, Anna Scherbina, and seminar participants at the 2014 Frontiers of Finance conference, 2014 Finance Down Under conference, 2013 Financial Research Association Meeting, 2013 Asian Finance Association Meeting, University of Georgia, TCU, Tulane University, University of Tennessee, UTSA, Washington State University, University of Memphis, and Clemson University for comments received. We thank Stocktwits.com for providing their data. An earlier version of this paper circulated under the title "The impact of divergence of opinions about earnings within a social network". Send correspondence to p.irvine@tcu.edu.

The Convergence and Divergence of Investors' Opinions around Earnings News: Evidence from a Social Network

January 2015

Abstract

We collect a unique dataset of Twitter posts to examine the change in investor disagreement around earnings announcements. We find that investors' opinions can either converge (reduced disagreement) or diverge (increased disagreement) around earnings announcements. The convergence and divergence of opinion has significant effects on trading volume and return. While the convergence of opinion is associated with lower earnings announcement returns, the divergence of opinion is associated with higher earnings announcement returns. Both the convergence of opinion and the divergence of opinion are associated with greater volume reaction to earnings news, consistent with recent theory.

1. Introduction

Investor disagreement can have a large influence on both asset prices and trading volume. With respect to asset prices, Miller (1977) predicts that when investors disagree on a firm's value, short sale constraints allow optimists to set the price of the firm's stock. As a result, investor disagreement leads to higher current prices and lower future returns. In Varian's (1985) model, however, investor disagreement is an additional risk factor and therefore associated with higher future returns.

Investor disagreement can also provide a rational solution to the puzzle of high trading volume, as the amount of trading observed in the stock markets is much too high to be driven by liquidity or hedging demand. In Kim and Verrecchia's (1991a, 1991b) models, investor disagreement prior to the disclosure of public information makes investors revise their beliefs differently and in turn generates trading volume. On the other hand, the models in Karpoff (1986), Kim and Verrecchia (1994), and Kandel and Pearson (1995) also relate trading volume to investor disagreement, but proposes a completely different mechanism. In their models, volume is generated by investors' different interpretation of the earnings news itself ("belief jumbling").

An earnings announcement is a powerful setting to test theories about investor disagreement because reported earnings can lead investors to revise their beliefs. Regarding asset prices, previous studies find evidence that disagreement is associated with lower returns around earnings announcements (Berkman, Dimitrov, Jain, Koch and Tice 2009) and higher post-earnings announcement returns (Garfinkel and Sokobin 2006). The former finding is consistent with Miller's model, while the latter is consistent with Varian's (1985) prediction. Regarding trading volume, Atiase and Bamber (1994) document a positive relation between investor disagreement and the trading volume around subsequent earnings announcements. Their findings support Kim and Verrecchia's theoretical predictions. Bamber, Barron and Stober (1997) provide empirical evidence of "belief jumbling" based on change in analyst forecasts around earnings announcements.

A limitation in research on the effects of disagreement is that existing studies use indirect measures of investor disagreement such as analyst forecast dispersion, historical trading volume, volatility of accounting performance, volatility of stock returns, or firm age. These proxies do not directly examine investors' opinions, with analyst forecast dispersion being the only exception. However, as discussed in Atiase and Bamber (1994) and Bamber, Barron, and Stevens (2011), analyst forecast dispersion captures only the opinions of analysts, who are a small subset of well-informed market participants. To our knowledge, the only study that directly uses investors' opinions to measure disagreement is Carlin, Longstaff, and Matoba (2014) who use survey data to construct a measure of disagreement in the mortgage-backed securities market. Carlin et al (2014) find that, for mortgage-backed securities, increased disagreement is associated with higher expected returns, higher return volatility, and larger trading volume.

Another limitation to the existing studies is that majority of the currently used proxies for investor disagreement do not change in a timely manner, especially those based on firm characteristics such as historical trading volume, volatility, or firm age, making it difficult to conduct comprehensive tests of the existing theory. To test the predictions of various theories, one needs to measure the change in investor disagreement around earnings announcements. In Miller's (1977) framework the release of news resolves disagreement and eliminates overpricing. The volume reaction in Kim and Verrecchia's (1991a, 1991b) models is accompanied by convergence of opinion around earnings announcements, but the volume reaction caused by "belief jumbling" is accompanied by a divergence of opinion.

We differ from most of the previous studies by using a unique dataset of social network messages to construct a measure of investor disagreement that directly reflects investors' opinions. Additionally, our measure of disagreement is at the daily frequency, enabling us to accurately identify disagreement before and after the earnings announcement date. Therefore, we are able to examine both the convergence and divergence in investors' opinions caused by earnings news, and conduct comprehensive analysis of the existing theories.

We measure investors' opinions using twitter messages about stocks from all the posters on Stocktwits.com, a website powered by Twitter. Stocktwits.com is a popular social media website where members exchange information on financial investments. We then use sentiment classification to identify the tone of Stocktwits posts around 19,751 quarterly earnings announcements by 2,983 firms from July 2009 to June 2011. Our text analysis techniques classify the overall tone and the strength of the opinion for the Stocktwits posts. We then match the sentiment of the Stocktwits posts to the sentiment of the news releases in the two-week windows before and after the earnings announcement date. In this way, we can classify whether investors tend to agree or disagree with publicly available news before the earnings announcement. Further, we investigate whether investors' opinions converge or diverge from public news after the earnings announcement. Our disagreement measure is based on the contrast between stock tweets and news articles. There is a large literature documenting that news articles impact investor opinions and stock returns. Therefore, disagreement between stock tweets and public news articles reflects a dispersion of opinions among investors, indicating low consensus or low correlation between investor's private valuations and public information.[1]

We first examine earnings announcement returns across two groups of investor disagreement in the pre-announcement period. The average two-day announcement abnormal return is 0.33% for the agreement group, but only -0.05% for the disagreement group. The difference between disagreement and agreement group is 0.37% (t-stat 3.40), both economically and statistically significant. These results are also robust in regressions that control for earnings surprise and other factors. This finding is consistent with Berkman et al. (2009) and supports Miller's (1977) theory.

[1] The seminal models of Kim and Verrecchia (1991a, 1991b) promote a distinct role for the dispersion of investor's private valuations. However, the difficulty in constructing timely measures of investor's dispersion has left this source of uncertainty largely ignored. In our attempt to fill this gap, we recognize that our sample contains only a subset of investors. Our choice of contrasting investor sentiment with the public news allows us to involve the opinions of a broader set of investors. A natural alternative is dispersion of opinions between twitter users which captures the opinions in only stock tweets, ignoring the public news. This measure is limited to the opinions of Twitter users in our sample. Using only stock tweets to measure dispersion produces similar results to those in the paper. However, the restriction that there must be multiple Twitter users in the event window to measure within-Twitter dispersion restricts the sample.

Additionally, our ability to dynamically measure investor disagreement over short windows allows for a sharper test of Miller's (1977) theory. We classify firms into four groups based on opinions before and after the earnings announcement: i) Investors disagree before announcement but agree after announcement (*DA*); ii) Investors disagree both before and after announcement (*DD*); iii) Investors agree before announcement but disagree after announcement (*AD*); iv) Investors agree both before and after announcement (*AA*). Miller (1977) assumes that earnings news reduces disagreement and therefore we expect negative announcement returns for the *DA* group. Hence, we predict *positive* announcement returns for the *AD* group where earnings news *generates* disagreement. Indeed, we observe a -0.16% announcement returns for the *DA* group but a 0.65% announcement return for the *AD* group, with a difference of 0.81% (t-stat 4.48), twice as large as the overall disagreement effect. In contrast, when the level of investor disagreement does not change (the *DD* group and the *AA* group), there is almost no difference in announcement returns (0.07%). These results provide new evidence that supports Miller's theoretical predictions only under the condition that new information resolves disagreement.

Since short sale constraints are a necessary condition for Miller's theory, we also conduct subgroup analyses using institutional ownership as a proxy for short sales constraints (Nagel 2005). We find that, consistent with Miller's theory, the effect of investor disagreement is much stronger for stocks with tighter short sales constraints (lower institutional ownership).

Next, we examine the impact of disagreement on post-earnings announcement returns. Using our measure of investor disagreement, we find a *negative* relation between investor disagreement and returns in various post-earnings announcement windows. For example, cumulative abnormal returns in the post-earnings announcement window of [2,20] is -0.27% for the agreement group but -0.72% for the disagreement group. The return difference is 0.45% (t-stat 2.90), both economically and statistically significant. We observe similar patterns using 10-day or 60-day returns in the post-earnings announcement window. These results are also robust in regressions that control for earnings surprise and other factors. Therefore, our findings show consistently negative relations between disagreement and returns both on and after earnings announcement. These findings support a generalization of Miller (1977) to the post-announcement period.

Our measure differs from the existing measures in that it directly reflects investors' opinions. To validate that investor opinions from Stocktwits can effectively capture information that is not subsumed by existing measures of disagreement, we control for all the existing measures in the regression analyses. The findings using our measures are robust after controlling for the existing measures, and our measure produces stronger results than the existing measures.

We next investigate the relation between investor disagreement and volume reaction to earnings announcements. We find that abnormal trading volume on and after earnings announcements is significantly higher in the presence of investor disagreement prior to the announcement, a finding consistent with Atiase and Bamber (1994). For example, cumulative abnormal turnover (share volume divided by shares outstanding) in the two-day window around earnings announcement [0,1] is 1.70% for the agreement group but 2.78% for the disagreement group. The difference is 1.08% (t-stat 14.11), both economically and statistically significant. This pattern persists in the 10- or 20-day window after earnings announcements, and robust in the multivariate regression setting. These results lend support to the theoretical models proposed by Kim and Verrecchia (1991a, 1991b) that prior investor disagreement generates volume reaction to earnings announcements.

Our data allows us to further divide earnings announcements into four groups according to changes in investor disagreement. We first find that indeed, convergence of opinion generates abnormal trading volume on and after earnings announcements. Specifically, using the *AA* (Agree→Agree) group as a benchmark, we observe positive abnormal volume for the *DA* group (Disagree→Agree). For example, the two-day announcement abnormal trading volume is 1.01% higher in the *DA* group than in the *AA* group. This result further supports the Kim and Verrecchia's (1991a, 1991b) models, where earnings news reduces investor disagreement.

Notably, we find that divergence of opinion also generates abnormal trading volume on and after the earnings announcement. For example, the two-day announcement abnormal trading volume is 1.07% higher in the *AD* group (Agree→Disagree) than in the *AA* group. This finding suggests that "belief jumbling" is also at work in generating abnormal volume upon earnings announcements. Our findings persist into the twenty-day period after the earnings announcement and they are robust in multivariate regressions that control for earnings surprises and other firm-level characteristics. These results clearly demonstrate that trading volume around earnings announcements can be generated by both divergence and convergence of opinion.

Motivated by this finding, we investigate if any commonly available financial or accounting variables can be used to identify whether abnormal volume at the time of the earnings announcement results from divergence or convergence of opinion. We find that idiosyncratic volatility and institutional investor breadth of ownership load positively on divergence of opinion and negatively on convergence of opinion, a result that indicates these variables are useful at identifying the source of abnormal volume. We also confirm that the low-return/high-volume proxy used in Bamber, Barron, and Stober (1999) helps to identify divergence of opinion in response to the earnings announcement.

Our study deepens the understanding of how investor disagreement impacts asset prices and trading volume. Using a unique and direct measure of investor disagreement, we test various theoretical models and provide new evidence on how convergence and divergence of opinion affect returns and trading volume. Regarding asset prices, our findings provide strong evidence that realization of Miller's (1977) theory depends on whether information reduces or exacerbates investor disagreement. Regarding trading volume, we demonstrate that earnings announcements generate trading volume through both divergence of opinion and convergence of opinion. Our findings on trading volume are consistent with the recent study by Banerjee and Kremer (2010) who derive a model that includes both investors' prior disagreement before earnings announcement and investors' different interpretation of earnings news.

Their model suggests that both the convergence and divergence of investors' opinions can generate separate volume reaction to earnings announcements.

Our study also has implications for measuring investor disagreement. Our measure has low correlations with the existing measures, and our findings suggest that both the existing market-based measures of disagreement and our measure may play independent roles in explaining the empirical results on disagreement, stock returns, and trading volume. This conclusion is consistent with Abarbanell, Lanen, and Verrecchia (1995) who contend that analyst forecast dispersion only partially captures investor uncertainty.

The rest of the paper is organized as follows. Section 2 outlines some basic hypotheses regarding investor disagreement around earnings announcements. Section 3 describes the data, Section 4 presents the empirical results. Section 5 analyzes the determinants of divergence and convergence of opinion. Section 6 concludes.

2. Theory and Hypothesis Development

2.1 Stock Return and Investor Disagreement

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. If investors think a stock is overvalued but cannot short the stock, they will either sell the shares they own or stay out of the market. For example, Lamont and Stein (2004) show that 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 (Almazan, Brown, Carlson, and Chapman, 2004).

Miller (1977) shows that in a world with short sale constraints and investor disagreement, an asset's price will reflect the valuations of the most optimistic investors and will exclude the opinions of pessimistic investors. Since investor disagreement causes overpricing, there should be a negative relationship between investor disagreement and future expected returns. If an asset is overpriced due to disagreement *and* the earnings announcement resolves this disagreement, any overpricing should disappear as the price falls to reflect the mean valuation of all investors.

In our setting, Miller's (1977) theory implies that investor disagreement before an earnings announcement causes negative expected abnormal returns after the announcement if the release of the actual earnings number eliminates disagreement among investors. Our unique data, which measures the sentiment of investors' opinions, allows us to directly test this prediction as in Hypotheses 1.

Hypothesis 1: Investor disagreement before an earnings announcement will cause future negative abnormal return after the earnings announcement.

While Miller's (1977) theory assumes information disclosure reduces investor disagreement (convergence of opinion), in reality earnings announcements can also generate investor disagreement (divergence of opinion). In the latter case the binding short sale constraints will lead to a positive price changes around the announcement. Our unique data on the change in investor disagreement allow us to test Hypotheses 2.

Hypothesis 2: Earnings announcement return will be negative when the announcement reduces investor disagreement but positive when the announcement generates investor disagreement.

2.2 Trading Volume and Investor Disagreement

Kim and Verrecchia (1991a) show that when there is a difference in the precision of investors' pre-announcement private information, investors will assign different weights to the announcement and revise their beliefs differently. This differential belief revision generates trading volume. Consistent with this prediction, a number of previous studies document that investor disagreement prior to earnings announcements is associated with greater volume reaction to earnings news (Atiase and Bamber, 1994; Bamber et al., 1997). We therefore use our unique measure to test Hypothesis 3.

Hypothesis 3: Investor disagreement prior to earnings announcements is associated with a greater trading volume reaction.

A holistic alternative is provided by Banerjee and Kremer (2010) who contend that trading volume reflects revisions to the level of disagreement. In their model, volume consists of both convergence trades and divergence trades. When agents agree about the interpretation of a common signal such as an earnings release, volume in this case depends on the degree of prior disagreement - more prior disagreement results in more convergence trading. On the other hand, differential interpretations of an earnings release can also result in ongoing disagreement and significant divergence trading after the release. While neither convergence trades nor divergence trades are unique to the Banerjee and Kremer (2010) model, their theoretical synthesis of these two types of trading motivates our empirical tests of volume at the time of earnings announcements.

The flexibility of our measure allows us to examine the effects of both divergence of opinion and convergence of opinion on trading volume.

Hypothesis 4a: Divergence of opinion around the earnings announcement is associated with a greater volume reaction to earnings announcement.

Hypothesis 4b: Convergence of opinion around the earnings announcement is associated with a greater volume reaction to earnings announcement.

3. Data Collection and Variable Construction

3.1 Collection of Stock Twits

Twitter is a micro blogging application where users are able to post short thoughts of no more than 140 characters, called tweets. While Twitter was started as a social network, its worldwide popularity and the broad user base have gained it fast growing impact on many aspects of people's lives.

Twitter is also related to the financial markets. Paul Hawtin, founder of Twitter hedge fund Derwent Capital, claims "Today, social media creates a vast amount of information and it has been proven that the sentiment derived from it can predict stock market movements." In April 2013, the U.S. Securities and Exchange Commission approved using Twitter to communicate company announcements. On April 24, 2013, the Dow Jones industrial average plunged by more than 140 points immediately after a hacker sent out a false tweet from Associated Press's account.[3] We collect Twitter posts from Stocktwits.com, an open micro-blogging site which is powered by Twitter with a focus 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. In 2010, Stocktwits was named Time.com's top 50 best websites as well as Fast Company's top 10 innovative companies in finance. Eggers (2014) discusses how Stocktwits, among other social networking platforms, is actively engaged in marketing their data directly to financial media outlets, and to data analytics firms who in turn market a refined feed to algorithmic traders and hedge funds.

We obtain twitter messages from all members of Stocktwits.com from July 10, 2009 to June 10, 2011. Figure 1 presents a sample of the original twitter messages on the Stocktwits website. Each message has the author, date, and time, as well as a picture that the author provides for their online profile. A post 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 Stocktwits community. When a blogger wants to post about a particular company on the Stocktwits website, they tag the company's ticker symbol with a "$".[4] For example, if you wanted to talk positively about Google and Microsoft you would say, "$GOOG and $MSFT, you should buy!" This is a practice called

---

[3] Reflecting the trend of increasing use of social media by firms and investors, Chen, Hwang, and Liu (2013) contend that the Twitter activity from company executives can help predict abnormal returns, while Chen, De, Hu, and Hwang (2014) suggest that negative sentiment from the social media site Seeking Alpha can predict future negative price performance.

[4] Although anyone can register and post opinions, popular members of Stocktwits.com are typically investment professionals such as newsletter writers, who use Stocktwits as a means of advertising their opinions to the investing public.

"Hashtagging" which is common place in the Twitter community. The "$" hashtag provides a mechanism to extract the company references in each post with a high level of accuracy.

For each post in the data, we have the content of the post, the associated ticker symbol(s), the date and the time of the post, and the blogger's ID and the number of followers. When a blogger talks about multiple companies in one post, we count all references as a unique post.[5] This initial sample contains 1,048,575 posts covering 7,757 security symbols. Since some of the symbols represent non-stock assets such as gold, foreign currencies, or indices, we further identify stock tweets by matching to stock tickers in CRSP. This procedure yields in total 782,904 stock tweets by 9,472 user IDs covering 5,927 stock tickers, with each post associated to a unique ticker and author. Finally, we match stock tickers to PERMNOs which is the unique firm identifier in our analysis, and the final sample contains 778,764 posts covering 5,806 unique PERMNOs. Section A1 of the Appendix describes the details of the matching procedures.

To examine the factors that impact Twitter activity, we estimate a firm-level cross-sectional regression of the total number of Tweets in our sample period on a number of firm characteristics. The results reveal that Stocktwits posting activity is generally correlated with other measures of investor attention. Stocktwits post activity is positively related to the frequency of news articles released during the sample period, analyst coverage, market-to-book ratio, and average daily return, though negatively correlated with market capitalization. We standardize the independent variables to assess the economic significance and find that the most important variable is the frequency of news articles, where one standard deviation increase in the frequency of news articles is associated with 0.574 standard deviation increase in the frequency of Tweets. The importance of news articles in explaining the frequency of Stocktwits posts suggests that many of the posts are commentaries on stock news from other sources.

---

[5] Among the original posts, 88% cover only one security symbol, 7% covering two symbols, 5% covering more than two symbols.

This finding supports our research design choice of contrasting Twitter sentiment with news sentiment, as many of the posts will either confirm or disagree with the news articles.[6]

3.2 Collection of News Stories

We collect news articles 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. We search the news stories for firms covered by the Stocktwits sample based on their stock tickers and company names. The initial sample contains 640,283 news stories during our sample period, each associated with a unique stock ticker. We match tickers of the news stories to PERMNOs to create our final news sample of 615,637 news stories covering 5,096 unique firms (PERMNOs). Section A2 of the Appendix describes the details of our news search procedures.

3.3 Construction of the Earnings Announcement Samples

Our initial sample of quarterly earnings announcements contains 38,773 announcements from the CRSP-COMPUSTAT merged dataset with non-missing announcement dates between July 24, 2009 and June 11, 2011. We also require the CRSP-COMPUSTAT firms to be in the Stocktwits sample. To exclude penny stocks, for the announcements in year *y*, we drop the firms with prices below \$2 or market capitalization below \$100 million at the end of the year *y*-1. We also require the share code of the announcing firm to be 10 or 11 (ordinary common shares) on the announcement date. Finally, we require the firms to have enough information to calculate cumulative abnormal announcement returns (described in Section 4.2). Our final sample contains 19,751 earnings announcements by 2,983 firms during the sample period.

3.4 Sentiment and Content Classification

We use the maximum entropy (ME) approach to classifying the information in Twitter posts. The ME approach derives meaning from the language in posts by applying a maximum likelihood algorithm to

[6] This regression is unreported for brevity, but available from the authors on request.

qualitative data. Since the information in a Twitter post can be subtle, 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. 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 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 classify.[7]

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 post, a 25% chance of being negative, and a 25% chance of being neutral. If our document did not have the word "buy" in it then we would just assume an equal

[7] 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.

distribution of a 33% chance that the document falls into each category. Thus, if we knew nothing about our document, we begin with a uniform distribution with equal likelihoods for each sentiment category. 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 $F = (f_1, \ldots f_m)$ 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. A list of the most common features is presented in Panel A of Table 1. 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: $\bar{d} = (n_1(d), n_2(d), , n_m(d))$. 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 $P(c=c_0 \mid \bar{d})$ is as follows:

$$P_{ME}(c = c_0 \mid \bar{d}) = \frac{1}{Z(d)} \left(\sum_i \lambda_{i,c} F_{i,c}(d, c) \right) \tag{1}$$

where $Z(d)$ is a normalization function, and $F_{i,c}$ is a feature category function for the feature i and for each category c defined as:

$$F_{i,c}(d,c) = \begin{cases} 1, \text{ if } n_i(d) > 0 \text{ and } c_i = c_0 \\ 0, \text{ otherwise} \end{cases} \tag{2}$$

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 positive 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. Panel B of Table 1 presents examples of representative weighting parameters. Using the weighting parameter allows us to implement Jegadeesh and Wu's (2013) finding that weighting can be an important tool in content analysis.

We implement the ME classifier by hand classifying a corpus of 2,000 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 $P_{ME}(c = c_0 \mid \bar{d})$ using maximum likelihood estimation 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 implement the ME approach using the Natural Language Toolkit (Bird, Klein, and Loper, 2009), a widely used text processing package. 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).[8]

We also try classifying the sentiment in Twitter posts using the Naïve Bayesian (NB) approach used by the existing literature (Li, 2010; Huang, Zang, and Zheng, 2014). We repeat the tests in this paper using the NB approach and find similar results. Section 4.5 describes the NB approach and the corresponding results.

3.5 Construction of the Investor Disagreement Measure

We use the *Sentiment* measure derived from our ME classification of Stocktwits posts 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:

$$IMPACT = (1 + Followers) \times Sentiment \tag{3}$$

[8] 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.

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.[9] *Followers* is the number of followers that an author has on Stocktwits. Posts are summed up over the day to determine the aggregate level of *IMPACT*. DeMarzo, Vayanos and Zwiebel's (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 investor posts a comment on Stocktwits. Finally, each post's number of followers interacts 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 news articles to measure the tone of public news. Once again, sentiment is defined as negative, neutral or positive (-1,0,1) for each news article 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 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 two-week window prior to the earnings announcement.

We are interested in the relationship between the sentiment of public news and the opinions of the investing public. Thus, to explicitly define the divergence of opinion we create a dummy variable *DIVOP* that equals 0 if the signs of *IMPACT* and *NEWS* are the same, and 1 if the signs of *IMPACT* and *NEWS* are different. The specific form of *DIVOP* is as follows:

$$DIVOP = \begin{cases} 0 \text{ if } (IMPACT > 0, NEWS > 0) \text{ or } (IMPACT = 0, NEWS = 0) \text{ or } (IMPACT < 0, NEWS < 0) \\ 1 \text{ } \textit{Otherwise} \end{cases} \tag{4}$$

[9] When there is no stock tweet or news article for a firm-day, we treat the Twitter sentiment or news sentiment for the firm-day as neutral (sentiment equals to zero).

Given this definition *DIVOP*=1 indicates disagreement between Stocktwits sentiment and news sentiment, while *DIVOP*=0 indicates sentiment agreement.

3.6 Other Measures of Investor Disagreement

Recent literature has proposed a number of variables to measure investor disagreement. We calculate these variables to determine how our measure of investor disagreement relates to the measures commonly used in the literature. One common measure of investor disagreement is the standard deviation of analysts' near-term earnings forecasts (e.g., Ajinkya, Atiase, and Gift, 1991; Abarbanell, Lanen, and Verrecchia, 1995; Doukas, Kim, and Pantzalis, 2006). We follow the literature and 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 days [-47,-3] prior to the earnings announcement as:

$$DISP = \frac{\left[\sum_{k=1}^{K} \frac{(Forecast_k - \overline{Forecast})^2}{k-1}\right]^{1/2}}{\left|\overline{Forecast}\right|} \tag{5}$$

where $Forecast_k$ is the k^{th} analyst's forecast of quarterly earnings per share and $\left|\overline{Forecast}\right|$ is the absolute value of the mean analyst forecast.[10]

We also construct other proxies for investor disagreement used in previous studies (e.g., Berkman et al. 2009) including the inverse of firm age (*AGE*); turnover (*TURN*), the average daily turnover during the pre-event period; stock return volatility (*RETVOL*), a measure of firm volatility relative to market volatility over the pre-event period; and earnings volatility (*INCVOL*), calculated using 20 quarters prior to the earnings announcement quarter. We described the construction of these proxies in Section A3 of the Appendix.

[10] 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.

We also follow Garfinkel and Sokobin (2006) and construct three measures of investor disagreement based on unexpected trading volume. We estimate their main proxy for disagreement, standardized unexplained volume (*SUV*) as follows. For each earnings announcement, we first estimate a regression of volume for firm *i* on day *t* in the 50-day window ending five days prior to the announcement [-54, -5].

$$Volume_{it} = \alpha_i + \beta_1 |R_{it}|^+ + \beta_2 |R_{it}|^- \tag{6}$$

This model treats positive and negative return days as independent events; as illustrated by the + and – superscripts in the regression equation. 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). We then use the estimated coefficients to calculate unexpected volume of day *t* (*t*=-1 or 0) as the deviation from expected volume:

$$UV_{it} = Volume_{it} - (\hat{\alpha}_i + \hat{\beta}_1 |R_{it}|^+ + \hat{\beta}_2 |R_{it}|^-) \tag{7}$$

Finally, we calculate standardize unexpected volume *(SUV)* as below:

$$SUV = \frac{1}{\sqrt{2}} \sum_{t=-1}^{t=0} \frac{UV_{it}}{\sigma_{it}} \tag{8}$$

where σ_{it} is the standard deviation of the residuals from regression (7) in the [-54,-5] window.

We also follow Garfinkel and Sokobin (2006) to calculate their second measure of investor disagreement, market adjusted turnover, as below:

$$MATO = \frac{\sum_{t=-1}^{t=0} \left[\left(\frac{Vol_{i,t}}{Shares_{i,t}} \right)_{firm} - \left(\frac{Vol_t}{Shares_t} \right)_{market} \right]}{2} \tag{9}$$

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*. Finally, we calculate Garfinkel and Sokobin's (2006) third measure of disagreement, change in market adjusted turnover (*ΔTO*), as *MATO* minus the average daily market adjusted turnover in the [-54,-5] window.

4. Empirical Results

4.1 Summary Statistics and Correlations of Disagreement Measures

Table 2 presents summary statistics for the sample earnings announcements. The announcing firms vary in size, with a median market capitalization of $800.22 million. Our sample stocks tend to be larger than the CRSP universe, as a stock must capture the attention of Stocktwits posters to enter into our sample. Reflecting their size, our sample stocks' *AGE* and turnover (*TURN*) are also relatively larger, but the mean income (*INCVOL*) and return (*RETVOL*) volatilities are comparable with previous studies. The sample announcing firms have a median two-day abnormal announcement return of -0.04%. Additionally, investor disagreement is observed prior to 46% of the sample earnings announcement.

Table 3 presents the correlations between the commonly used disagreement measures and our measure, *DIVOP*. Theoretically it is interesting that *DIVOP* has low correlations with the commonly used measures. *DIVOP* has a low 0.05 correlation with standardized unexpected volume (*SUV*) and a -0.01 correlation with analyst forecast dispersion (*DISP*). Abarbanell, Lanen and Verrecchia (1995) point out that in the seminal models of Kim and Verrecchia (1991a, 1991b) analyst dispersion is orthogonal to investor's private beliefs. Thus, the *DIVOP* measure, derived from investor sentiment, should be largely uncorrelated with *DISP* in the Kim and Verrecchia (1991a, 1991b) framework. Further, the correlations between *DIVOP* and the other measures are quite low as well, ranging from -0.05 to 0.20. The correlations between the Berkman et al. (2009) disagreement measures and Garfinkel and Sokobin's (2006) main measures (*SUV* and *ΔTO*) are also low, ranging between -0.10 to 0.35. It is interesting that the commonly used measures of investor disagreement in the literature have low correlations. The low correlations between these measures and *DIVOP* indicate that the results we find for *DIVOP* are unlikely to be explained by its correlation with the existing measures.

4.2 Investor Disagreement and Earnings Announcement Returns

Miller (1977) gives rise to *Hypothesis 1* which predicts that investor disagreement about earnings causes overvaluation because prices will be set by the most optimistic investors. However, once a firm releases 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.

We first divide the sample earnings announcements into two groups based on the *DIVOP* measure, and examine earnings announcement returns. We measure earnings announcement returns as cumulative abnormal returns in the two-day window [0,1] where day 0 is the announcement day. To control for the effects of market risk, size, book-to-market ratio, and momentum on stock returns, we construct daily abnormal returns based on the four-factor model:

$$AR_{it} = R_{it} - \left(\hat{\alpha}_{it} + \hat{\beta}_{1i}(R_{mt} - r_f) + \hat{\beta}_{2i}SMB_t + \hat{\beta}_{3i}HML_t + \hat{\beta}_{4i}UMD_t\right) \quad (10)$$

where the daily factor returns R_{mt} -R_f, *SMB*, *HML*, and *UMD* are obtained from Kenneth French's data library, and the factor loadings are estimated in the 90-day window [-120,-31] prior to the earnings announcement.[11]

Figure 2 plots the cumulative abnormal returns in the [-5,10] window around earnings announcements for the disagreement and agreement groups. For the agreement group, stock prices move up on the announcement day and the day after, and remain relatively stable afterwards. While the disagreement group exhibits a material downward drift from the earnings announcement day. Consistent with this return divergence, Panel A of Table 4 further reports that the earnings announcement return in the [0,1] window is 0.33% for the agreement group (*DIVOP*=0) but -0.05% for the disagreement group (*DIVOP*=1). The return spread is 0.37% (t-stat 3.40), both economically and statistically significant. For robustness, we examine three-day announcement returns in Panel A and observe similar pattern. These

[11] We thank Kenneth French for making the data available.

results indicate that investor disagreement is associated with lower earnings announcement returns, a finding supportive of *Hypothesis 1*.

We then test *Hypothesis 2* by dividing the sample announcements into four groups according to the changes in disagreement around earnings announcements. Specifically, we compute the *DIVOP* measure in two-week window after earnings announcements and divide the sample announcements into four groups: i) Investors disagree before announcement but agree after announcement (*DA*), ii) Investors disagree both before and after announcement (*DD*), iii) Investors agree before announcement but disagree after announcement (*AD*), iv) Investors agree both before and after the announcement (*AA*).

Figure 3 plots cumulative abnormal returns in the [-5,10] window around earnings announcements for the four change groups. We observe that, consistent with *Hypothesis 2*, there is a strong upward drift for the *AD* group but a downward drift for the *DA* group on and after earnings announcements. In contrast, there is little difference in earnings announcement returns between *AA* and *DD* groups. Panel B of Table 4 confirms the pattern in Figure 3. Specifically, we observe a -0.16% announcement returns for the *DA* group but a 0.65% announcement return for the *AD* group, with a difference of 0.81% (t-stat 4.48), twice as large as the overall disagreement effect in Panel A of Table 4. In contrast, there is almost no difference in announcement returns (0.07%) between the *DD* group and the *AA* group. The results with three-day announcement returns are also similar. This new evidence supports *Hypothesis 2* and Miller's theoretical predictions.

If *DIVOP* captures investor disagreement and its effects on earnings announcement returns result from the mechanism in Miller's (1977) model, then we expect the effects of *DIVOP* to be stronger among stocks with tighter short sale constraints. We follow the existing literature (e.g., Nagel, 2005) and measure short sale constraints using institutional ownership, where lower ownership indicates tighter short sale constraints. We obtain institutional holdings from Thomson Reuters 13f database and calculate institutional ownership for a sample announcement as the shares of the announcing firm held by institutions divided by the firm's total shares outstanding, measured at the end of the quarter prior to the announcement.

In Panel C of Table 4, we classify the sample announcements into two groups of institutional ownership and examine the effect of *DIVOP* on earnings announcement return for both groups. The spread in earnings announcement return between the disagreement and the agreement groups is 0.64% (t-stat 3.90) for the low ownership stocks but only 0.18% (t-stat 1.22) for the high ownership stocks. In Panel D of Table 4, the return spread between the *AD* and *DA* groups is 1.11% (t-stat 4.05) for the low ownership stocks, much larger than the 0.50% (t-stat 2.14) for the high ownership stocks. These results show that the effect of *DIVOP* on announcement returns is much stronger when short sale constraints bind, which lends strong support to Miller's (1977) theory and the validity of *DIVOP* as a measure of investor disagreement.

For robustness, we also repeat the test using the disagreement within Twitter posts instead of that between Twitter posts and news articles. We require at least two Twitter posts for an announcement and calculate standard deviation of the sentiment of the posts for each announcement. We then classify an announcement into the disagreement (agreement) group if the standard deviation of sentiment is in the top (bottom) half of all announcements.[12] The results are not reported for brevity but they exhibit the same patterns as those using the DIVOP measure in Panels A and B of Table 4. For example, the two-day earnings announcement return (CAR[0,1]) is 0.05% for the disagreement group but -0.51% for the disagreement group. The spread of 0.56% is both economically and statistically significant (t-stat 2.31).

We further examine the effects of disagreement on earnings announcement returns in a multivariate regression setting. Clearly, return surrounding an earnings announcement depends on the content of the announcement itself. We measure the information content using standardized unexpected

[12] For an alternative approach, we also try classifying an announcement into the agreement (disagreement) group if the sentiment of all posts are equal (not equal), and the results are similar.

earning (*SUE*). We thus control for the effects that unexpected positive or negative news events have on the returns across *DIVOP* groups. We construct *SUE* as below:

$$SUE = \frac{Actual - Expected}{P} \tag{11}$$

where *Actual* is actual earnings, *Expected* is the median analyst forecasts prior to earnings announcements, and *P* is the stock price at the end of the fiscal quarter. To control for outliers, we winsorize *SUE* at the 1 percent and 99 percent cutoffs. Section A4 of the Appendix describes of the details of the construction of *SUE*.

Table 5 presents the regressions of cumulative abnormal returns in the [0,1] announcement window. We report Driscoll and Kraay (1998) robust t-statistics, which control for time-series and cross-sectional correlations. The key independent variable is the *DIVOP* measure. In the univariate regression in Model 1, the coefficient on *DIVOP* is negative -0.369 and significant (t-stat -3.21). The magnitude of this coefficient is consistent with the return spread in Panel A of Table 4. In Model 2, we include *SUE*, the *IMPACT* and *NEWS* variables, and ten lagged daily returns to control for short-term reversals. The coefficient on *DIVOP* is similar to that in Model 2. In model 3, we further include firm fixed effects to control for all the firm-level differences, an extremely strict control because its inclusion eliminates firm-level differences in investor disagreement. The coefficient on *DIVOP* in this specification is slightly smaller and significant at the 10% level (t-stat=-1.72). Regarding the control variables, the coefficient on *SUE* is significantly positive while the coefficients on *NEWS* and *IMPACT* are insignificant across most of the models. These results are consistent with the sorting analysis in Table 4.

To relate our analysis to the current literature, we add the five disagreement proxies used in Berkman et al. (2009) into models 4-8. We do not include firm fixed effects in these models because the majority of these proxies are not time-varying by construction. The coefficient on *DIVOP* is significantly negative in all models, suggesting that the effect of *DIVOP* is robust after controlling for the existing disagreement measures.

4.3 Investor Disagreement and Post-Earnings Announcement Returns

Garfinkel and Sokobin (2006) find that their measures of investor disagreement are *positively* related to returns in the post-earnings announcement period. In this subsection we re-examine the effect of investor disagreement on post-earnings announcement returns.

Figure 4 plots the cumulative abnormal returns in the post-announcement period [1,60] for disagreement and agreement groups. The disagreement group exhibits an almost monotonic downward drift in the post-announcement period, which does not exist in the agreement group. We report the magnitude of this return divergence in Table 6. Panel A of Table 6 shows that cumulative returns in the ten-, twenty- and sixty-day windows after earnings announcement are all much lower in the disagreement group than in the agreement group. For example, the cumulative abnormal return in the [2, 20] window is 0.45% lower in the disagreement group than in the agreement group. This return spread is both economically and statistically significant, with a t-statistic of 2.90. Panel B shows that the return difference between *AD* and *DA* groups is larger than the difference in Panel A, especially for the ten- and twenty-day windows. These results suggest that Miller's (1977) theory is a closer approximation of reality when his assumption that uncertainty is resolved is maintained.

In Table 7 we conduct multivariate regressions of the post-earnings announcement returns in which the dependent variable is the cumulative abnormal return in the [2,60] announcement window. The key independent variable is the *DIVOP* measure. We start our analysis with a simple univariate regression in Model 1, and then control for *SUE, IMPACT, NEWS*, and lagged returns in Model 2, and further control for firm fixed effects in Model 3. In all models the coefficient on *DIVOP* is significantly negative. The magnitude is also consistent with the sorting analysis in Table 6. For example, in Model 3 the coefficient is -1.470, suggesting that the presence of investor disagreement reduces *CAR* in the [2,60] window by 1.47%, which is consistent with the corresponding 1.41% return difference in the Panel A of Table 6. Models 4 and 5 show that the effect of *DIVOP* on post-earnings announcement returns is robust when we examine returns over ten- and twenty-day windows.

In Table 8, we include all the investor disagreement measures used in previous studies (e.g., Berkman et al., 2009; Garfinkel and Sokobin 2006) in the regressions of post-earnings announcement returns in the [2,60] window. Specifically, we test whether the effect of *DIVOP* remains significant in the presence of the alternative proxies for investor disagreement. We do not control for firm fixed effects as many of the alternative measures are not time-varying by construction. In all eight models the coefficient on *DIVOP* remains significantly negative.

To summarize, using the *DIVOP* measure, we observe a consistently negative relation between investor disagreement and returns on and after earnings announcements. The additional results on convergence of opinion and divergence of opinion show that Miller's (1977) theory holds provided his assumption that the earnings announcement resolves disagreement is true.

4.4 Investor Disagreement and Trading Volume

Investor disagreement is a common theoretical explanation for volume (Kim and Verrecchia, 1991a,b). As outlined in *Hypothesis 3*, we expect investor disagreement to be positively related to the volume reaction to earnings announcements. We test *Hypothesis 3* by examining abnormal trading volume around earnings announcements across investor disagreement groups.

We calculate cumulative abnormal volume by constructing an unexpected volume measure for every stock-day around earnings announcements. The methodology of the volume event study is very similar to a standard event study. To calculate abnormal trading volume we follow the methodology of Campbell and Wasley (1996) by defining $V_{it} = \frac{volume_{it}}{S_{it}}$, where S is the total number of shares outstanding. Next we estimate a market model of volume over the [-245,-45] 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.

Figure 5 plots cumulative abnormal volume in the [-5,20] window around earnings announcement for the disagreement group and the agreement group. While both groups experience a rise in trading volume on earnings announcement, the disagreement group has much higher volume than the agreement group. Panel A of Table 9 reports the cumulative abnormal volume in various windows from the earnings announcement day to twenty days after earnings announcement. In all windows, abnormal trading volume is much higher for the disagreement group than for the agreement group. For example, the cumulative abnormal volume in the [0,20] window is 2.18% for the agreement group but 4.03% for the disagreement group. The difference of 1.85% (t-stat 6.53) is both economically and statistically significant. These results are consistent with *Hypothesis 3* and the previous studies (e.g., Atiase and Bamber, 1994; Bamber et al., 2011).

As an alternative perspective about volume generation, Banerjee and Kremer (2010) contend that the release of earnings information can have dual effects on investor disagreement. Specifically, volume at an earnings announcement could reflect both convergence of opinion from diverse priors as well as differential interpretation of the earnings news. The timeliness of our measure of investor disagreement allows us to effectively test *Hypotheses 4a* and *4b* about the effects of both convergence and divergence of opinion on trading volume.

Figure 6 plots cumulative abnormal volume in the [-5,20] window around earnings announcement for the four groups based on the change in investor disagreement (*DA*, *DD*, *AD*, *AA*). In the pre-earnings announcement period, there are no dramatic differences in volume across the four categories. Interestingly, while abnormal volume increases for all four categories at the time of the earnings announcement, the dramatically different patterns of abnormal volume appear for the four categories.

In Figure 6 the *AA* (Agree→Agree) group has the lowest level of abnormal volume. What is striking is that the pattern of abnormal volume changes after the earnings announcement coincides with both divergence and convergence of investors' opinions generating trading volume. Specifically, as investors' opinions move from agreement to disagreement (*AD* group) around the earnings announcement period, abnormal volume spikes up and continues to increase above the *AA* group. This result suggests that divergence of opinion generates abnormal trading volume. Additionally, as investors' opinions move from disagreement to agreement (*DA* group) around the earnings announcement period, we also observe significantly higher abnormal volume than the *AA* group. This result confirms that convergence of opinion also generates abnormal trading volume. The *DD* (Disagree→Disagree) group has the highest abnormal volume throughout ending the quarter with almost 5% cumulative volume increase, suggesting, naturally, that investor disagreement generates trading volume.

Panel B of Table 9 further reports cumulative abnormal volume for the four categories of change in disagreement. For all windows, abnormal trading volume is much higher for the *AD* and *DA* groups than for the *AA* group. For example, compared to the *AA* group, the cumulative abnormal volume in the [0,20] window is 2.84% (t-stat 12.01) higher for the *AD* group and 2.13% higher (t-stat 9.43) for the *DA* group.

We proceed to investigate the effects of investor disagreement in a multivariate regression framework. In Table 10 we use cumulative abnormal volume *(CAV),* measured across various event windows, as the dependent variable. In Models 1 to 4 the dependent variable is *CAV* in the [0,1] announcement window. In Models 1 to 3, the key independent variable is *DIVOP*. We start with a univariate regression of *CAV* on *DIVOP* in Model 1, and then, in Model 2, control for *SUE, IMPACT, NEWS,* and ten lagged daily volumes. We add firm fixed effects to the regression specification in Model 3. In all models the coefficient on *DIVOP* is significantly positive, consistent with the sorting analysis in Panel A of Table 9. In Model 4, we include the three binary variables *AD*, *DA*, and *DD*. Therefore the coefficients represent how much abnormal volume is associated with these groups relative to the *AA* group. The coefficients on all three binary variables are significantly positive, consistent with the sub-group analysis. Models 5 to 7 further show that the patterns are similar when we examine abnormal volumes in the twenty-day post-earnings announcement window.

To summarize, our findings support *Hypotheses 4a* and *4b* that earnings announcements can produce trading volume through both divergence and convergence of opinion. This result is important in the context of interpreting Garfinkel's (2010) proxy of *SUV* (unexpected volume) for disagreement at the time of the announcement. As unusually high levels of volume can be generated by both convergence and divergence trades (Banerjee and Kremer, 2010), it is not necessarily the case that *SUV* at the time of the earnings announcement is associated with higher levels of disagreement, rather *SUV* could be high because of convergence trading. As these convergence trades represent decreasing, not increasing, disagreement, the volume generated by these trades does not reflect Varian's (1985) theoretical positive relation between uncertainty and returns. Thus, the positive relation between announcement-period volume (*SUV*) and future returns, as documented by Garfinkel and Sokobin (2006), should be qualified as a noisy proxy because the abnormal volume can be driven by both divergence and convergence trades. The continuous nature of our measure can help ongoing research identify specific convergent and divergent episodes around earnings which permits sharper empirical tests of existing theories relating disagreement to returns.

4.5. Classifying Sentiment in Twitter Posts using the NB Approach

To corroborate the results using the Maximum Entropy (ME) approach, we repeat the analysis but classify sentiment of the Twitter posts using a Naïve Bayesian (NB) approach rather than the ME approach. The existing literature (Li, 2010) shows that the NB approach is superior to dictionary and word count methods for predicting the sentiment of forward-looking statements in corporate filings. Similar to ME estimation, the NB classifier is a maximum likelihood application of Bayes' rule to a set of document features but it makes the simplifying assumption that all features are independent. In the field of machine learning, NB is a popular technique to use as a baseline approach for document classification or spam filtering. We follow a similar NB approach to Li (2010) to classify information, where the only difference is that Li uses four categories (positive, negative, neutral, and uncertain) whereas we maintain the three categories (positive, negative, neutral) to be consistent with our main tests.[13]

The NB approach is similar to the ME approach in that it relies on a training set of hand-classified posts to determine the probability that each word reflects a positive, negative or neutral sentiment. Each word is a feature used to classify any document in the full data set. The classification method uses the same equations: (1) and (2). Practically, NB is a constrained ME technique; the constraint being that the NB algorithm can only use single words, and not word combinations, to classify sentiment. Using the NB approach allows us to relate our study to the classification approach used in the existing literature and provides a robustness test of the findings in this paper.

After classifying the information in sample tweets, we repeat the main analyses in the paper. Panels A and B of Table 11 repeat the analyses of earnings announcement returns as in the Panels A and B of Table 4, and Panels C and D of Table 11 repeat the analyses of announcement trading volume as in the Panels A and B of Table 9. Compared to the results using the Maximum Entropy approach, the results using the Naïve Bayesian approach are similar, although sometimes slightly weaker. Overall, Table 11 shows that our findings are robust to the alternative approach of classifying sentiment.

5. What Factors Drive the Divergence and Convergence of Opinion?

In this section we ask what factors that are commonly available can be used to identify divergence and convergence of opinion around earnings announcements. This question is important because its answer not only helps us conceptually understand the formation and evolution of investment disagreement upon information releases, but also helps us empirically distinguish the different regimes of investor disagreement. It is worth noting that there has been little empirical research in this regard with the sole exception being Bamber, Barron, and Stober (1999). They show that earnings announcements with high

[13] Li's (2010) uncertain category refers to specific words in his dictionary that referred to uncertainty, rather than being uncertain as to how to classify the document. The Naïve Bayesian approach is also used by the recent study of Huang, Zang, and Zheng (2014).

trading volume but low absolute announcement returns are associated with divergence of opinions, a finding supportive of Kandel and Pearson's (1995) argument that the release of public information can, in some cases, lead to divergence of opinion. In this section, we conduct comprehensive analyses of this question using our high-frequency measure of investor disagreement that allows accurate measurement of the divergence and convergence of opinion. By identifying variables that are associated with divergence or convergence of opinion, we hope to assist future research that will not necessarily have access to a twitter feed to identify whether abnormal volume is driven by the divergence or convergence of opinion.

We first examine the factors driving the *divergence* of opinion, for which we estimate logistic regressions using the subsample of earnings announcements in the *AA* group (Agree→Agree) and *AD* group (Agree→Disagree).[14] The dependent variable is a dummy for divergence of opinion that equals 1 for the *AD* group and zero for the *AA* group. We then include independent variables that are potentially associated with divergence of opinion. We first follow Bamber et al. (1999) and construct a dummy variable *Low|Ret|/HighVol* which equals one if an earnings announcement is in the bottom quintile of the absolute earnings announcement return (CAR[-1,1]) and the top half of the earnings announcement trading volume (CAV[-1,1]), and zero otherwise. A significantly positive coefficient on this dummy variable will support Bamber et al.'s (1999) finding that low-return/high-volume earnings announcements indicates divergence of opinion.

We also include a number of additional independent variables including firm size, book-to-market ratio, and return momentum. We include these variables as they are commonly examined firm characteristics and thus will indicate if identifiable firm types tend to be associated with divergence of opinion. We also examine the effect of idiosyncratic return volatility constructed using market model of daily returns in the previous year. Higher idiosyncratic volatility can indicate divergence of opinions because it is associated with more firm-specific information and greater information asymmetry. We

[14] We also repeat the tests using probit regressions and observe similar results.

further examine breadth of ownership because the larger the number of shareholders of an announcing firm, the more likely that their priors contain difference pieces of information about the firm, which can make them respond differently to the firm's earnings news ("belief jumbling", Bamber et al., 1997). Since the data on individual shareholders are not available, we use the number of institutional shareholders before an earnings announcement. Finally, we examine the relation between earnings quality and divergence of opinion, because poor earnings quality reflects a noisy information environment about earnings, which may lead investors to interpret earnings information differently. Since there are a number of measures of earnings quality, we examine two commonly used measures including earnings persistence and absolute accruals (Dechow, Ge, and Schrand, 2010). Earnings persistence of an announcing firm is measured by the coefficient in regressions of annual earnings on lagged annual earnings over the past eight years. Absolute accrual of an announcing firm is the absolute value of the announcing firm's previous-year total accruals. Section A5 of the Appendix provides more detail on the construction of these variables.

Panel A of Table 12 presents the results of the logistic regressions for divergence of opinion. In Model 1, the independent variable is the low-return/high-volume dummy. Model 2 further includes change in turnover and change in return volatility around earnings announcement. Both turnover and return volatility are commonly used disagreement measures in the current literature.[15] Model 3 further includes the additional factors discussed above.

Several interesting results emerge in Panel A of Table 12. First, the coefficient on the low-return/high-volume dummy is significantly positive across all models. That is, using our unique measure of investor disagreement, we confirm Bamber et al.'s (1999) finding that small earnings-announcement return accompanied by high trading volume is an indication of divergence of opinion. Second, the coefficients on change in turnover and change in return volatility are both insignificant. This result is not

[15] The other three measures used by Berkman et al. (2009) are firm age, income volatility, and analyst forecast dispersion. Firm age and income volatility are time-invariant. Calculating change in analyst forecast dispersion requires at least two recent analyst forecasts of quarter q+1 earnings both before and after the announcement of q, which are unavailable for the majority of sample announcements.

surprising given the low correlations between these measures and *DIVOP*. Third, as we expected, the coefficient on idiosyncratic volatility is significantly positive, suggesting that divergence of opinion is more likely to occur in the firms with higher idiosyncratic volatility. Fourth, we observe that, consistent with "belief jumbling" (Bamber et al., 1997), firms with greater breadth of ownership are more likely to experience divergence of opinion. The other independent variables are insignificant except return momentum which is significantly positive, indicating that past winner stocks are more likely to experience divergence of opinion.

Next, we examine the drivers of *convergence* of opinion, where we also estimate logistic regressions but to identify convergence events we contrast the subsample of earnings announcements in the *DD* group (Disagree→Disagree) and *DA* group (Disagree→Agree). The dependent variable is a dummy variable that equals 1 for the *DA* group and zero for the *DD* group. We examine the same set of variables as discussed above to test whether these variables have opposite effects on convergence of opinion than they do in the case of divergence of opinion.

Panel B of Table 12 presents the results. Convergence of opinion is negatively related to low-return/high-volume dummy, although its coefficient becomes insignificant in Model 3 which uses a smaller sample restricted due to data availability. This result is in line with the observed positive relation between divergence of opinion and the low-return/high volume dummy. Additionally, convergence of opinion is *negatively* related to idiosyncratic volatility and breadth of ownership, whereas the relation is significantly *positive* for divergence of opinion. Among the other variables, the coefficient on firm size is significantly negative, suggesting that investor disagreement is more difficult to resolve for larger firms.

Overall, the results in this section provide interesting and new evidence on the factors that drive the divergence and convergence of investors' opinions around earnings announcements. Earnings announcements are more likely to resolve investor disagreement if idiosyncratic volatility and breadth of ownership are low and the firm is smaller. Earnings announcements are more likely to exacerbate disagreement when idiosyncratic volatility and breadth of ownership are high, when there are low announcement returns accompanied by high volume, or when the firm has had high returns over the past year. These results can help identify the source of abnormal volume at the time of an earnings announcement when direct measures of investors' opinions are not available.

6. Conclusion

In this paper, we use a unique set of data collected from the social network site Stocktwits to measure investor disagreement. We examine the impact of disagreement on the price and volume reactions to earnings announcements. We explicitly measure investor's opinions rather than relying on financial market data or analyst data to examine the effects of investor disagreement on stock returns and trading volume. By measuring the disagreement between public news and social network communication we find support for the theoretical predictions of Miller (1977) and Banerjee and Kremer (2010). We show that investor disagreement prior to the earnings announcement is generally uncorrelated with financial market proxies for investor disagreement.

We find that investor disagreement before earnings announcements leads to significantly lower returns on and after an earnings announcement. We further extend our tests by measuring changes to the level of disagreement from the pre- to post-earnings announcement period. We find that earnings announcements where opinions change from prior agreement to disagreement produce positive abnormal returns, a pattern that markedly differs from the negative abnormal returns for the sample announcements where prior disagreement changes into agreement. This test illustrates that while earnings announcements can sometimes increase consensus among investors, they can also produce the opposite effect, a divergence of opinion after the announcement.

We further examine the trading volume reaction to earnings announcements. Our findings show that investor disagreement prior to earnings announcement is associated with greater volume reaction to earnings announcement. More importantly, both divergence of opinion (Kim and Verrecchia, 1994) and convergence of opinion (Banerjee and Kremer, 2010) generate abnormal trading volume on and after earnings announcements.

References

Abarbanell, J., W. Lanen, and R. Verrecchia, 1995. Analysts forecasts as proxies for investor beliefs in empirical research. *Journal of Accounting and Economics* 20, 31-60.

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.

Atiase, R., and L. Bamber, 1994. Trading volume reactions to annual earnings announcements: The incremental role of information asymmetry. *Journal of Accounting and Economics* 17, 309-329.

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.

Bamber, L., O. Barron and T. Stober, 1997. Trading volume and different aspects of disagreement coincident with earnings announcements. *Accounting Review* 72, 575-597.

Bamber, L., O. Barron and T. Stober, 1999. Differential interpretations and trading volume. *Journal of Financial and Quantitative Analysis,* 34 (3), 369-386.

Bamber, L., O. Barron, and D. Stevens, 2011. Trading volume around earnings announcements and other financial reports: Theory, research design, empirical evidence and directions for future research. *Contemporary Accounting Research* 28(2), 431-471.

Banerjee, S. and I. Kremer, 2010. Disagreement and learning: Dynamic patterns of trade. *Journal of Finance* 65(4): 1269-1302.

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.

Bird, S., E. Klein and E. Loper, 2009. Natural language programming with Python, O'Reilly Media, Inc.

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.

Carlin, B., F. Longstaff, and K. Matoba, 2014, Disagreement and asset prices, *Journal of Financial Economics,* 114, 226-228.

Chen, H., B-H. Hwang, and B. Liu, 2013, The economic consequences of having social executives, Working paper, City University of Hong Kong.

Chen, H., P. De, Y. Hu and B-H. Hwang, 2014, Wisdom of crowds: The value of stock opinions transmitted through social media, *Review of Financial Studies*, forthcoming.

Diether, K., C. Malloy and A. Scherbina, 2002. Differences of opinion and the cross-section of stock returns, *Journal of Finance* 57(5), 2113-2141.

Dechow, P., W. Ge, and C. Schrand, 2010, Understanding earnings quality: A review of the proxies, their determinants and their consequences, *Journal of Accounting and Economics* 50, 344–401.

DeMarzo, P., D. Vayanos, and J. Zwiebel, 2003. Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics* 118(3), 909-968.

Doukas, J., C. Kim, and C. Pantzalis, 2006. Divergence of opinion and equity returns. *Journal of Financial and Quantitative Analysis* 41, 573–606.

Driscoll, J. and A. Kraay, 1998. Consistent covariance estimation with spatially dependent panel data. *Review of Economics and Statistics* 80, 549-560.

Eggers, B., 2014. More investors using social media to make decisions, Reuters.com, June, 21.

Engelberg, J., 2008. Costly information processing: Evidence from earnings announcements. Working paper, University of California, San Diego.

Garfinkel, J., 2010. Measuring investors' opinion divergence. *Journal of Accounting Research* 47(5), 1317-1348.

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.

Huang, A., A. Zang, and R. Zheng, 2014, Evidence on the information content of text in analyst reports, *Accounting Review* 89, 2151-2180.

Jegadeesh, N. and D. Wu, 2013. Word power: A new approach for content analysis, *Journal of Financial Economics* 110, 712-729.

Karpoff, J., 1986. A theory of trading volume, *Journal of Finance* 41, 1069-1087.

Karpoff, J., 1987. The relationship between price changes and trading volume: A survey. *Journal of Finance and Quantitative Analysis* 22, 109-26.

Kandel, E., and N. Pearson, 1995. Differential interpretations or information and trade in speculative markets. *Journal of Political Economy* 103(4), 831-872.

Kim, O. and R. Verrecchia, 1991a. Trading volume and price reactions to public announcements. *Journal of Accounting Research* 29(2), 302-21.

Kim, O. and R. Verrecchia, 1991b. Market reaction to anticipated announcements. *Journal of Financial Economics* 30(2), 273-209.

Kim, O. and R. Verrecchia, 1994. Market liquidity and volume around earnings announcements. *Journal of Accounting and Economics* 17, 41-67.

Lamont, O. and J. Stein, 2004. Aggregate short interest and market valuation. *American Economic Review* 94(2), 29-32.

Loughran, T., and B. McDonald, 2011, When is a liability not a liability? *Journal of Finance* 66, 35-65.

Miller, E., 1977. Risk, Uncertainty, and Divergence of Opinion. *Journal of Finance,* 32, 1151-1168.

Nagel, S., 2005, Short sales, institutional investors and the cross-section of stock returns, *Journal of Financial Economics* 78, 277–309.

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.

Tetlock, P., 2007. Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance* 62(3), 1139–1168.

Varian, H., 1985. Divergence of opinion in complete markets: A note. *Journal of Finance* 40, 309-317.

Appendix

A1. Matching Tickers to PERMNOs

We use PERMNOs to identify sample firms to assist the merging between data sets. Since both the Stocktwits messages and the news articles are based on stock tickers, we create a linking file that assigns PERMNO to a TICKER-date during 2009 to 2011. We first download CRSP daily stock file from January 2009 to December 2011, and identify the first and last dates of each PERMNO-ticker pair. Then, for each calendar day from January 2009 to December 2011, we assign the corresponding PERMNO to a ticker as long as the day is between the first and the last days of the PERMNO-ticker pair. We then examine the resulting matches and find that while most of the PERMNO-ticker pairs are one-one matches for a given day, there are a very small number multiple matches between PERMNO and ticker on a day. We address these multiple matches as follows:

1) One PERMNO matched to two tickers: Two PERMNOs 90469 and 91501 are each matched to two tickers on some days. This is due to the change in tickers during an interim period. For example, PERMNO 90469's ticker is ARBX for most of the time during our sample period, but for the one-month interim period from June 14, 2010 to July 12, 2010, its ticker changes to ARBXD. Therefore our procedure of using the start and end dates assigns both the tickers ARBX and ARBXD to the PERMNO 90469 for this one-month period. We address this issue by keep only the valid tickers (ARBXD in the case of PERMNO 90469) for these two tickers in the sub-periods.
2) One ticker matched to two PERMNOs: During 2009 to 2011, there are 52 tickers each matched to two PERMNOs for either the whole period or a sub-period. We find that these cases are due to a firm issuing shares of two classes which correspond to two different PERMNOs (e.g., shares with voting power vs. shares without voting power). To address this issue, for each of these 52 tickers, we calculate the total share volume for two PERMNOs separately during 2009 to 2011, and keep the PERMNO with the larger share volume. In most cases, the share volume of one PERMNO is much larger than the other.

A2. Collection of News Stories

The news search is based on the tickers and firm names. For each of the 5,927 stock tickers covered by Stocktwits, we collect the corresponding firm name (names) from the CRSP monthly stock file during the sample period. We then search the news stories from Dow Jones Newswire, Reuters News, and PR Newswire from July 10, 2009 and June 10, 2011. When we search a firm, we first enter the ticker, and then pick a name from Factiva's suggested list of firm names that matches the firm's name in CRSP. We also eliminate the duplicates of news stories for a given firm.

We collected in total 640,283 news articles using the tickers and company names. We then matched the articles to PERMNOs using the approach described in Section A1. There are 615,637 articles (96.2%) matched to PERMNOs. The unmatched articles are outside the date ranges of CRSP for the corresponding tickers. This happens because even when a firm is not traded in the exchange, it can still have news coverage. For example, General Motor (PERMNO 12079) stopped trading on June 1, 2009 and resumed trading on November 18, 2010 with a new PERMNO of 12369. GM's news articles during this interim period are therefore unmatched to a PERMNO.

A3. Construction of the Proxies of Investor Disagreement in Berkman et al. (2009)

We follow the approaches in Berkman et al. (2009) and construct the following proxies of investor disagreement:

- Analyst Dispersion (*DISP*): We first obtain the quarterly IBES detail file for quarters ending (FPEDATS) during 2009 to 2011. We then assign PERMNO to a firm-quarter in IBES by matching IBES' CUSIP to CRSP's historical CUSIP (NCUSIP) on the quarter-end date. If the IBES quarter-end date is a non-trading day, then we match to CRSP's PERMNO on the next trading day. We then merge the firm-quarter to the quarterly Compustat data by PERMNO and quarter-end date. We use the Compustat's earnings announcement date (RDQ) to calculate analyst dispersion. We take all the analyst forecasts in the 45-day window ending three days prior to earnings announcement [-47,-3], and calculate analyst dispersion as standard deviation of the forecasts divided by the absolute value of mean forecast. We require at least two forecasts in the estimation window to calculate the dispersion measure.
- Income Volatility (*INCVOL*): We take quarterly CRSP-Compustat merged data, and calculate the quarterly income as QIBDPQ divided by the average of this quarter and previous quarter's total assets (ATQ). We then calculate seasonal income difference for a quarter as quarterly income minus income of the same quarter of last year. Then for an earnings announcement, we take seasonally differenced income in the 20 quarters prior to the announcement quarter, and calculate the *INCVOL* measure as standard deviation of the seasonally differenced incomes. We require a firm to have available income data for at least eight quarters in the estimation window.
- Return Volatility (*RETVOL*): We first obtain daily stock returns from the CRSP daily file. For an earnings announcement, we calculate the *RETVOL* measure as standard deviation of excess

daily returns in the 45-day window ending 11 days prior to the earnings announcement [-55,-11], where excess daily return is raw return minus CRSP value-weighted return.

- Firm Age: We take the first year of a firm in CRSP as the firm's first year. Then for an earnings announcement, we calculate firm age as the number of years between the first year and the earnings announcement year.
- Turnover (*TURN*): We obtain daily share volume and shares outstanding from CRSP daily file, and calculate daily turnover as daily share volume divided by shares outstanding. Then for an earnings announcement, we calculate the *TURN* measure as average daily turnover in the 45-day window ending 11 days prior to the announcement date [-55,-11].

We follow Berkman et al. (2009) to winsorize the *INVOL*, *RETVOL*, *DISP*, and *TURN* measures at the 99 percent cutoffs.

A4. Construction of *SUE*

We obtain analyst forecasts and actual earnings data from quarterly IBES detail file for all firm-quarters (FPEDATS) during 2009 to 2011. We then assign PERMNO and match to Compustat earnings announcement date using the approaches described in Section A3. To calculate *SUE*, we first take the analyst forecasts of an announcement. If an analyst issues more than one forecasts during this period, we pick the latest forecast. We then calculate *SUE* = (*Actual* – *Median Forecasts*)/*Price*, where *Price* is the stock price of the quarter-end date (or next trading day if the quarter-end is a non-trading day). To control for outliers, we winsorize *SUE* at the 1 percent and 99 percent cutoffs.

A5. Construction of the Variables Driving the Convergence and Divergence of Opinion

We construct the independent variables in the regression analyses of factors driving the convergence and divergence of opinion (Section 5) as follows:

- *Low|Ret|/HighVol*: A dummy variable that equals one if an earnings announcement is in the bottom quintile of the absolute value of earnings announcement return (CAR[-1,1]) and the top half of the trading volume around earnings announcement (CAV[-1,1]), and zero otherwise. The ranks are based on the 19,751 sample earnings announcements used in our analyses.
- *⊿TURN*: Change in turnover (*TURN*) around earnings announcement. We follow Berkman et al. (2009) and calculate pre-announcement turnover (*TURN*) as average daily turnover measures over a 45-day period ending 10 days before the earnings announcement day, and

post-announcement turnover (*TURN*) as average daily turnover measures over a 30-day period starting 10 days after the earnings announcement day.

- *⊿RETVOL*: Change in return volatility (*RETVOL*) around earnings announcement. We follow Berkman et al. (2009) and calculate pre-announcement return volatility (*RETVOL*) as the standard deviation of the announcing firm's daily excess returns relative to a value-weighted market index, over a 45-day period ending 10 days before the earnings announcement. We also follow Berkman et al. (2009) and calculate post-announcement return volatility (*RETVOL*) as the standard deviation of the announcing firm's daily excess returns relative to a value-weighted market index, over a 30-day period starting 10 days after the earnings announcement.
- ln(*ME*): Natural log of the market capitalization of the announcing firm at the end of the previous year.
- ln(*B/M*): Natural log of the book-to-market ratio of the announcing firm. We calculate book-to-market ratio of an announcing firm of fiscal year *y* as its book value of common equity (Compustat annual item CEQ) divided by its market value of equity (Compustat annual item PRCC_F×CSHO), and apply the book-to-market ratio to its earnings announcements in the one-year period starting from July of year *y*+1.
- *Ret*[-12, -2]: Buy-and-hold stock return of the announcing firm from calendar month -12 to calendar month -2, where calendar month 0 is the calendar month of announcement.
- Idiosyncratic volatility: Standard deviation of the residuals of the market model of daily stock returns of the announcing firm in the one-year period ending in the month before earnings announcement. We require at least 100 daily return observations in the estimation year.
- ln(1+#*inst*): Natural log of 1 plus the number of institutional shareholders of the announcing firm at the end of the quarter before the earnings announcement. The number of institutional shareholders is calculated as the number of unique manager IDs (MGRNO) in Thomson Reuters's 13f institutional holdings data.
- Earnings persistence: For each firm-year, we calculate annual earnings as Compustat annual item IB divided by lagged Compustat annual item AT. We estimate earnings persistence as coefficient in the regression of the announcing firm's earnings on lagged earnings in the eight years ending in the year before the earnings announcement. We require earnings to be available in each of the estimation years.
- |Accrual|: Absolute value of the announcing firm's total accrual of the fiscal year before the earnings announcement. Total accrual is calculated as earnings minus operating cash flow, scaled by lagged total assets (Compustat annual item (IB - OANCF)/lagged AT)

Figure 1
Interface of www.Stocktwits.com

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

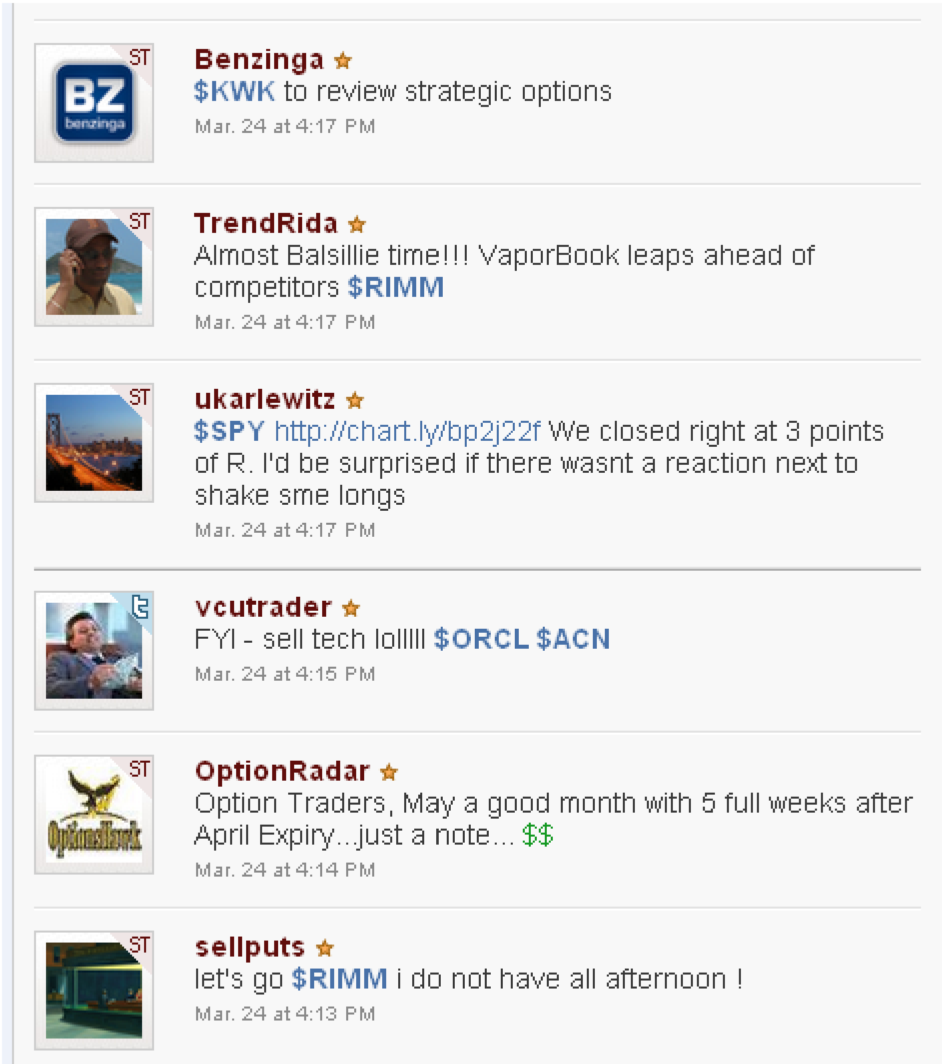

Figure 2: Cumulative Abnormal Returns around Earnings Announcement: [-5,10] Window

This figure depicts cumulative abnormal returns in the [-5,10] window around earnings announcements, where day 0 is the earnings announcement day. For an earnings announcement, we estimate daily abnormal returns using the four-factor model where the coefficients are estimated in the 90-day window [-120,-31] ending 31 days prior to the earnings announcement. We then sort announcements into two groups based on the *DIVOP* measures, which is a dummy variable that takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. We calculate average daily abnormal returns during [-5,10] for the two groups and plot cumulative abnormal returns. For the ease of comparison, we set the abnormal returns of day -6 to zero.

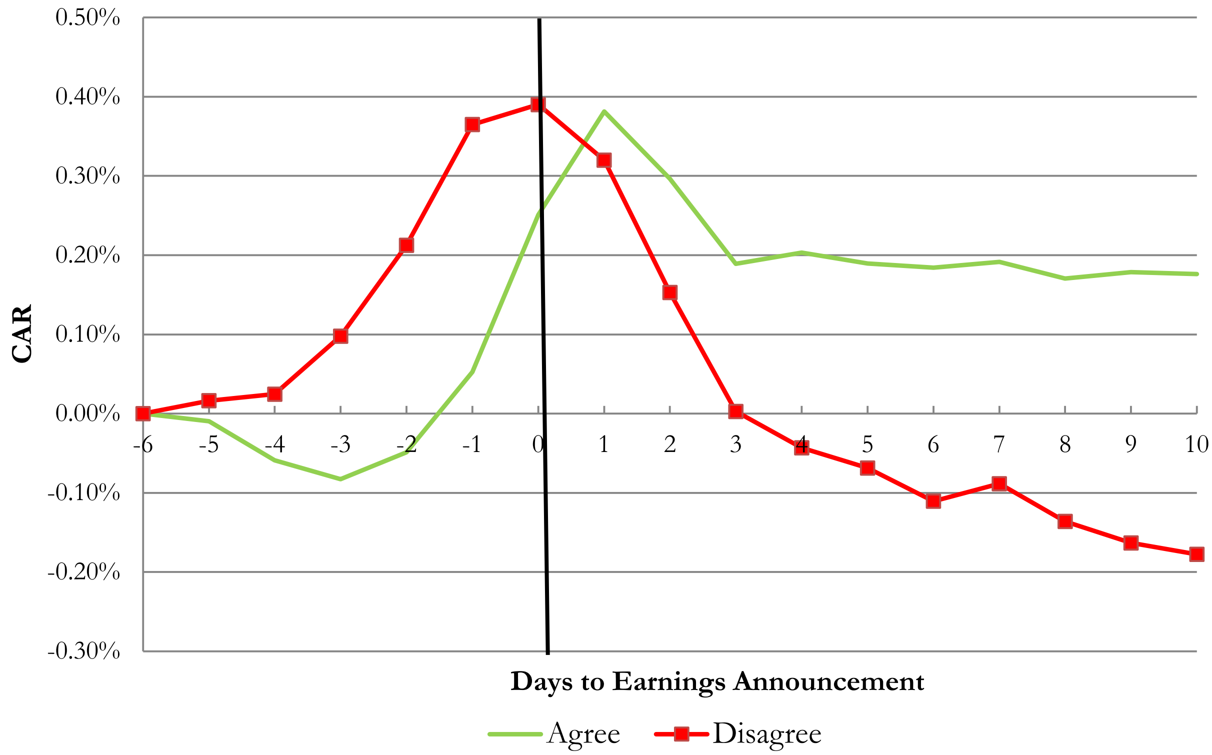

Figure 3: Cumulative Abnormal Returns around Announcements: Change in Disagreement

This figure depicts cumulative abnormal returns in the [-5,10] window around earnings announcements, where day 0 is the earnings announcement day. For an earnings announcement, we estimate daily abnormal returns using the four-factor model where the coefficients are estimated in the 90-day window [-120,-31] ending 31 days prior to the earnings announcement. Then for each announcement, we first calculate the pre-announcement *DIVOP* which is a dummy variable that takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. We further calculate the post-announcement *DIVOP* over the two weeks after earnings announcement. We assign announcements into four groups based on the values of the pre- and post-announcement *DIVOP* measures. The "Agree → Agree" group contains the announcements where the *DIVOPs* equal 0 in both pre- and post-announcement periods. The "Agree → Disagree" group contains the announcements where *DIVOP* equals 0 in the pre-announcement period but 1 in the post-announcement period. The "Disagree → Agree" group contains the announcements where *DIVOP* equals 1 in the pre-announcement period but 0 in the post-announcement period. The "Disagree → Disagree" group contains the announcements where the *DIVOPs* equal 1 in both pre- and post-announcement periods. We then calculate average daily abnormal returns during [1,60] for each group and plot cumulative abnormal returns. For the ease of comparison, we set the abnormal returns of day -6 to zero.

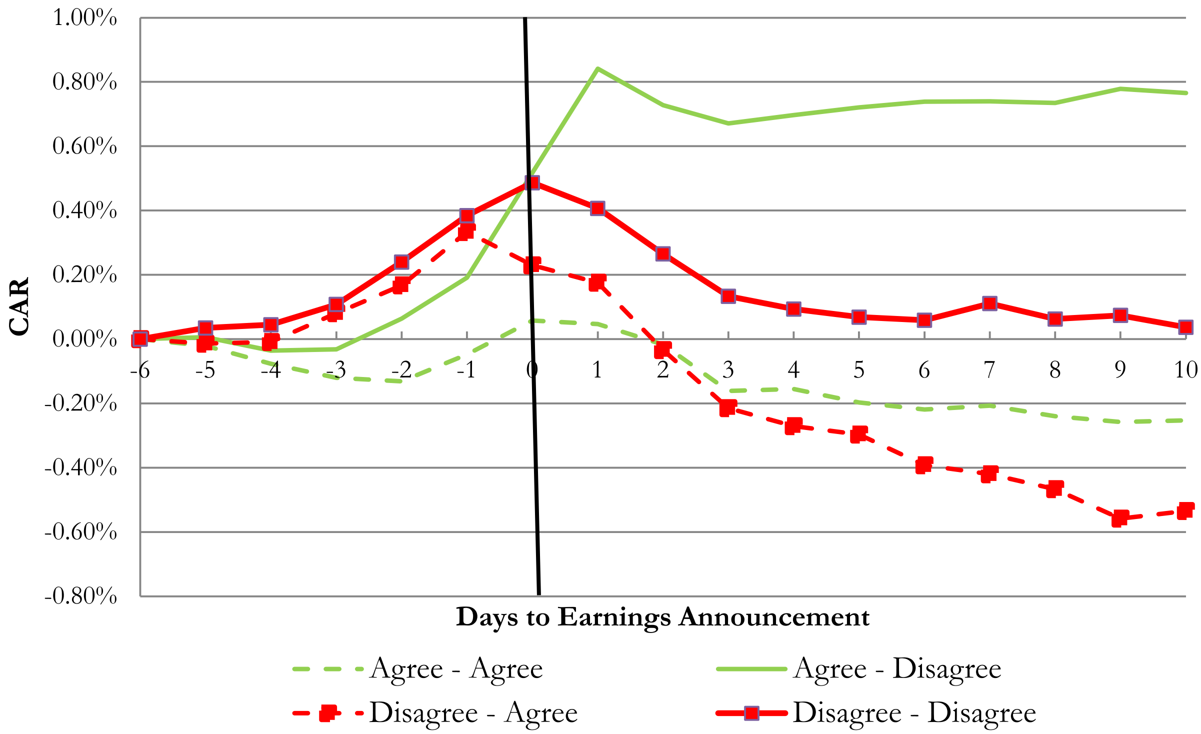

Figure 4: Cumulative Abnormal Returns in the Post-Announcement Period: [1,60] Window

This figure depicts cumulative abnormal returns in the post-earnings announcement window of [1,60], where day 0 is the earnings announcement day. For an earnings announcement, we estimate daily abnormal returns using the four-factor model where the coefficients are estimated in the 90-day window [-120,-31] ending 31 days prior to the earnings announcement. We first sort announcements into two groups based on the value of the *DIVOP* measure, which is a dummy variable that takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. We then calculate average daily abnormal returns during [1,60] for the two groups, and plot the cumulative abnormal returns. For the ease of comparison, we set the abnormal returns of day 0 to zero.

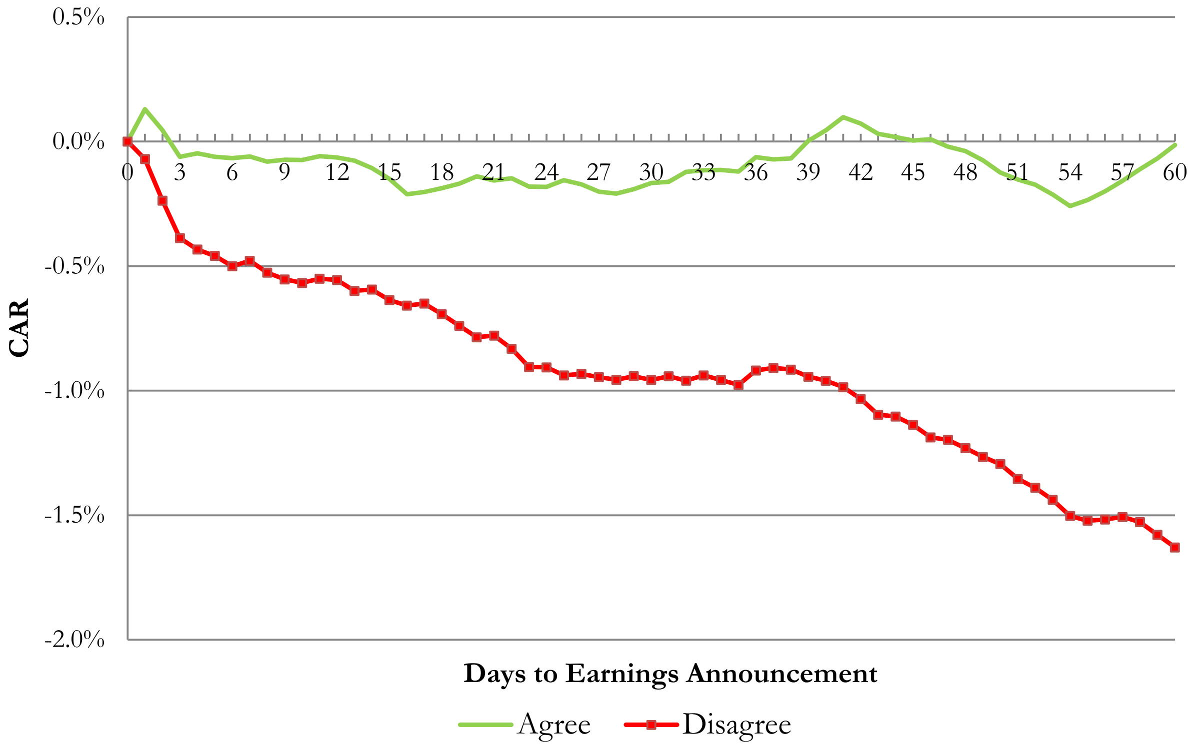

Figure 5: Cumulative Abnormal Trading Volume around Earnings Announcements

This figure depicts cumulative abnormal trading volume in the [-5,20] window around earnings announcements, where day 0 is the earnings announcement day. For an announcement, we use the methodology of Campbell and Wasley (1996) to calculate daily abnormal trading volume where the coefficients are estimated in the 200-day window ending 45 days prior to the earnings announcement [-245,-45]. We first sort announcements into two groups based on the value of the *DIVOP* measure, which is a dummy variable that takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. We then calculate average daily abnormal trading volume for the two groups and plot cumulative abnormal volumes. For the ease of comparison, we set the abnormal volume of day -6 to zero.

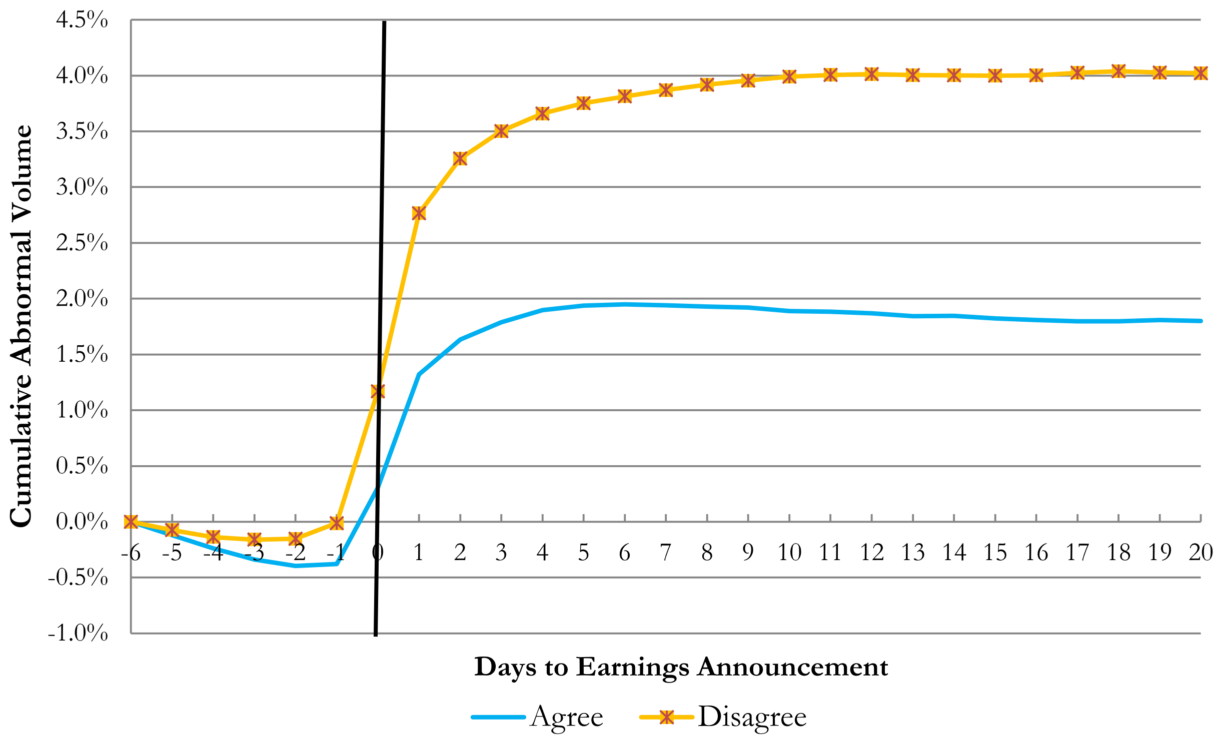

Figure 6: Cumulative Abnormal Trading Volume around Earnings Announcements: Change in Disagreement

This figure depicts cumulative abnormal trading volume in the [-5,20] window around earnings announcements. For an announcement, we use the methodology of Campbell and Wasley (1996) to calculate daily abnormal trading volume where the coefficients are estimated in the 200-day window ending 45 days prior to the earnings announcement [-245,-45]. For each announcement, we calculate the pre-announcement *DIVOP* measure, which is a dummy variable that takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. We further calculate the post-announcement *DIVOP* measure over the two weeks after earnings announcement. We then assign announcements into four groups based on the values of the pre- and post-announcement *DIVOP* measures. The "Agree → Agree" group contains the announcements where the *DIVOPs* equal 0 in both pre- and post-announcement periods. The "Agree → Disagree" group contains the announcements where *DIVOP* equals 0 in the pre-announcement period but 1 in the post-announcement period. The "Disagree → Agree" group contains the announcements where *DIVOP* equals 1 in the pre-announcement period but 0 in the post-announcement period. The "Disagree → Disagree" group contains the announcements where the *DIVOP* equals 1 in both pre- and post-announcement periods. We then calculate average daily abnormal trading volume during the [-5,20] window for each group, and plot the cumulative abnormal volume. For the ease of comparison, we set the abnormal volume of day -6 to zero.

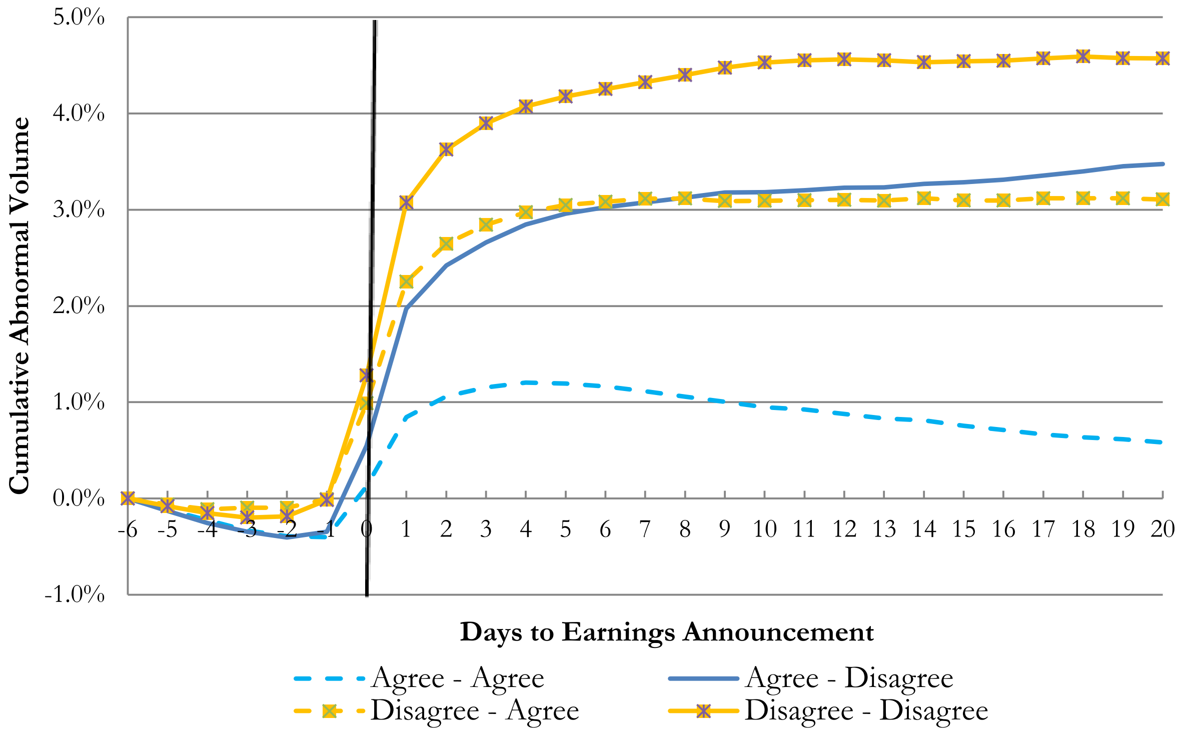

Table 1: The Twitter Corpus

This table presents examples of the Twitter dictionary constructed to classify posts. Panel A reports words with the highest frequencies from the positive and negative Twitter posts. *Frequency* is the percentage of an individual word representation in a given category. Panel B presents a sample of Features (words and phrases) that have strong odds of classifying posts into positive and negative categories. The *Positive Classification* panel presents the positive feature and the odds ratio relative to one other particular classification. The *Negative Classification* panel presents the negative feature and the odds ratio relative to one other particular classification. For example, the feature "strong" has odds of 12.2:1 that the post is positive rather than neutral.

Panel A: High Frequency Words from Twitter Posts

Positive Corpus				Negative Corpus			
Word	Frequency	Word	Frequency	Word	Frequency	Word	Frequency
long	0.84%	go	0.30%	short	0.74%	get	0.28%
good	0.81%	new	0.30%	down	0.49%	downside	0.28%
buying	0.48%	above	0.27%	rt	0.46%	rally	0.25%
buy	0.45%	back	0.27%	good	0.40%	still	0.25%
bought	0.42%	strong	0.27%	new	0.40%	bad	0.25%
today	0.42%	week	0.27%	today	0.37%	sell	0.25%
get	0.36%	call	0.24%	last	0.31%	looks	0.22%
looking	0.36%	chart	0.24%	big	0.31%	see	0.22%
looks	0.36%	low	0.24%	just	0.31%	looking	0.22%
just	0.33%	play	0.24%	market	0.28%	one	0.22%

Panel B: High Weight (λ) Words from Twitter Posts

Positive Classification			Negative Classification		
Feature	λ	Relative to classification in:	Feature	λ	Relative to classification in:
strong	12.2:1	Neutral	rt	11.0:1	Positive
a good	10.7:1	Neutral	bad	7.4:1	Neutral
buying	10.7:1	Neutral	I don't	7.4:1	Neutral
sector	7.9:1	Neutral	lower	7.4:1	Neutral
question	7.9:1	Neutral	article	6.4:1	Neutral
for this	6.4:1	Neutral	saying	5.7:1	Neutral
small	6.4:1	Neutral	end	5.7:1	Neutral
long	5.8:1	Negative	short	5.6:1	Positive
above	5.6:1	Neutral	says	5.0:1	Positive
great	5.0:1	Neutral	twitpic	5.0:1	Positive

Table 2 Summary Statistics

This table presents summary statistics for the 19,751 sample earnings announcements during the sample period of July 10, 2009 to June 10, 2011. *Market Cap.* is the announcing firm's market capitalization measured at the beginning of the year of announcement. *CAR* [0,1] is the cumulative abnormal return in the [0,1] window where 0 is the earnings announcement day. Daily abnormal return is based on the four-factor model estimated in the 90-day window [-120,-31] ending 31 days prior to the earnings announcement. *DIVOP* is the dummy variable of investor disagreement constructed using Twitter posts and news articles in the two weeks prior to the earnings announcement. *DIVOP* takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. We also report summary statistics for the commonly used measures of investor disagreement. *DISP* is the dispersion of analyst forecasts in the 45-day period ending 3 days prior to the earnings announcement. *INCVOL* is the standard deviation of the quarterly operating income over the 20 quarters prior to the earnings announcement quarter. *RETVOL* is the standard deviation of the announcing firm's daily excess returns relative to a value-weighted market index, over a 45-day period ending 10 days before the earnings announcement. *AGE* is the number of years the announcing 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. *MATO* (market adjusted turnover) is the average daily abnormal turnover in the [0,1] window, where daily abnormal turnover is the firm's daily turnover minus the daily turnover of all CRSP firms. *ΔTO* is calculated as *MATO* minus the announcing firm's average daily abnormal turnover in the [-54,-5] window. Standardized unexpected volume, *SUV*, is the abnormal volume in the [0,1] window where abnormal volume is constructed based on a market model-style model regression of volume on absolute valued returns estimated in the 50-day window ending 5 days prior to the announcement.

Variable	Mean	Std Dev	P10	P25	P50	P75	P90
Market Cap. ($M)	4,605.75	17,247.46	165.80	306.83	800.22	2,574.08	8,520.79
CAR [0,1] (%)	0.16	7.69	-8.19	-3.63	-0.04	3.82	8.73
DIVOP	0.46	0.50	0.00	0.00	0.00	1.00	1.00
DISP	0.24	0.56	0.01	0.03	0.07	0.18	0.50
INCVOL	0.020	0.027	0.003	0.006	0.012	0.023	0.045
RETVOL	0.021	0.011	0.010	0.013	0.018	0.025	0.034
AGE	20.23	17.76	3.00	7.00	15.00	27.00	42.00
TURN	0.010	0.009	0.003	0.005	0.008	0.012	0.020
MATO	0.008	0.030	-0.011	-0.007	0.000	0.012	0.033
ΔTO	0.011	0.025	-0.002	0.001	0.005	0.013	0.028
SUV	2.980	6.045	-0.586	0.241	1.726	4.154	7.526

Table 3: Correlations between *DIVOP* and the other Disagreement Measures

This table presents the simple correlation coefficients between *DIVOP* and the other disagreement measures used in the literature. *DIVOP* is a dummy variable that takes the value of 1 if there is disagreement between firm's news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. We construct the *DISP*, *INCVOL*, *RETVOL*, Log(1/*AGE*), and *TURN* measures using the approaches in Berkman, Dimitrov, Jain, Koch and Tice (2009). *DISP* is the dispersion of analyst forecasts in the 45-day period ending 3 days prior to the earnings announcement. *INCVOL* is the standard deviation of the quarterly operating income over the 20 quarters prior to the earnings announcement quarter. *RETVOL* is the standard deviation of the announcing firm's daily excess returns relative to a value-weighted market index, over a 45-day period ending 10 days before the earnings announcement. *AGE* is the number of years the announcing 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 calculate the *SUV*, *MATO*, and *ΔTO* measures using the approaches in Garfinkel and Sokobin (2006). *MATO* (market adjusted turnover) is the average daily abnormal turnover in the [0,1] window, where daily abnormal turnover is the firm's daily turnover minus the daily turnover of all CRSP firms. *ΔTO* is calculated as *MATO* minus the announcing firm's average daily abnormal turnover in the [-54,-5] window. Standardized unexpected volume, *SUV*, is the abnormal volume in the [0,1] window where abnormal volume is constructed based on a market model-style model regression of volume on absolute valued returns estimated in the 50-day window ending 5 days prior to the announcement. We also report the correlations between the disagreement measures and *SUE*, where *SUE* is the standardized unexpected earnings calculated as the difference between the actual earnings and expected earnings (based on IBES median estimates) divided by stock price.

	DIVOP	DISP	INCVOL	RETVOL	log(1/AGE)	TURN	SUV	ΔTO	MATO
DISP	-0.0122								
INCVOL	0.0440	0.0681							
RETVOL	-0.0471	0.1854	0.3158						
log(1/AGE)	-0.0422	0.0593	0.2028	0.2895					
TURN	0.1986	0.1066	0.1608	0.3280	-0.0054				
SUV	0.0510	-0.0371	-0.0265	-0.0979	0.0026	0.0100			
ΔTO	0.1087	0.0146	0.0348	0.0610	0.0282	0.3489	0.6172		
MATO	0.1573	0.0425	0.0823	0.1537	0.0229	0.6030	0.5175	0.9475	
SUE	0.0019	-0.0402	0.0020	-0.0940	-0.0165	-0.0285	0.0067	-0.0108	-0.0256

Table 4: Earnings Announcement Returns across Disagreement Groups

Panel A presents earnings announcement returns across prior disagreement groups. We first sort announcements into two groups based on the value of the *DIVOP* measure, which is a dummy variable that takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. We then report average cumulative abnormal returns (CARs) in the [0,1] and [-1,1] window for the two groups, where day 0 is the announcement day. For an earnings announcement, we estimate daily abnormal returns using the four-factor model where the coefficients are estimated in the 90-day window [-120,-31] ending 31 days prior to the earnings announcement. We also report the return differences between the two groups and the associated two-sample t-statistics assuming unequal variances. Panel B presents earnings announcement returns for the four groups of changes in investor agreement. We calculate *DIVOP* over the two weeks after earnings announcements and divide the sample earnings announcements into four groups. The "Agree → Agree" group contains the announcements where the *DIVOPs* equal 0 in both pre- and post-announcement periods. The "Agree → Disagree" group contains the announcements where *DIVOP* equals 0 in the pre-announcement period but 1 in the post-announcement period. The "Disagree → Agree" group contains the announcements where *DIVOP* equals 1 in the pre-announcement period but 0 in the post-announcement period. The "Disagree-Disagree" group contains the announcements where *DIVOP* equals 1 in both pre- and post-announcement periods. We also report the differences in *CARs* between the "Disagree → Agree" and the "Agree → Disagree" groups and the associated t-statistics. Panels C and D repeats the tests in Panels A and B, respectively, but for the low institutional ownership subsample and the high institutional ownership subsample separately. We classify the sample earnings announcements into two subgroups according to institutional ownerships of the announcing firms, where institutional ownership for an announcing firm is the number of shares held by institutional investors divided by total shares outstanding, measured at the end of the quarter prior to the earnings announcement. ***, **, * represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

	#obs	CAR [0,1]	CAR [-1,1]
Panel A: Earnings Announcement Returns Across Groups of Disagreement (%)			
Agree	10,762	0.33	0.43
Disagree	8,989	-0.05	0.11
Agree – Disagree		-0.37***	-0.32***
t-stat		-3.40	-2.85
Panel B: Earnings Announcement Returns Across Groups of Change in Disagreement (%)			
Agree → Disagree	4,532	0.65	0.78
Agree → Agree	6,230	0.09	0.18
Disagree → Disagree	5,616	0.02	0.17
Disagree → Agree	3,373	-0.16	0.01
DA – AD		-0.81***	-0.77***
t-stat		-4.48	-4.14
Panel C: Earnings Announcement Returns Across Groups of Disagreement (%): Subgroups of Institutional Ownership			
Low Institutional Ownership			
Agree	5,784	0.25	0.39
Disagree	4,089	-0.38	-0.18
Disagree – Agree		-0.64***	-0.57***
t-stat		-3.90	-3.38

	#obs	CAR [0,1]	CAR [-1,1]
Panel C: Earnings Announcement Returns Across Groups of Disagreement (%): Subgroups of Institutional Ownership			
High Institutional Ownership			
Agree	4,978	0.42	0.48
Disagree	4,900	0.24	0.35
Disagree – Agree		-0.18	-0.13
t-stat		-1.22	-0.86
Panel D: Earnings Announcement Returns Across Change in Disagreement (%): Subgroups of Institutional Ownership			
Low Institutional Ownership			
Agree → Disagree	2,155	0.61	0.77
Agree → Agree	3,629	0.04	0.16
Disagree → Disagree	2,415	-0.30	-0.04
Disagree → Agree	1,674	-0.50	-0.39
DA – AD		-1.11***	-1.16***
t-stat		-4.05	-4.08
High Institutional Ownership			
Agree → Disagree	2,377	0.69	0.79
Agree → Agree	2,601	0.18	0.20
Disagree → Disagree	3,201	0.27	0.32
Disagree → Agree	1,699	0.19	0.40
DA – AD		-0.50***	-0.39
t-stat		-2.14	-1.60

Table 5: Regressions of Earnings Announcement Returns

This table presents the regressions of earnings announcement returns on the investor disagreement measures. The dependent variables are the two-day cumulative abnormal returns in the [0,1] window with respect to earnings announcement, where day 0 is the announcement day. For an earnings announcement, we estimate daily abnormal returns using the four-factor model where the coefficients are estimated in the 90-day window [-120,-31] ending 31 days prior to the earnings announcement. 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's news articles and Twitter sentiment over the two weeks prior to earnings announcement, and 0 if there is agreement. We also control for *IMPACT*, which measures the level of Twitter sentiment over the two weeks prior to earnings announcement; *NEWS*, which measures the sentiment of firm's news articles over the two weeks prior to earnings announcement; and *SUE*, which is the standardized unexpected earnings measured as the difference between the actual earnings and expected earnings (based on IBES median estimates) divided by stock price, measured in percentage. We also include disagreement proxies in Berkman et al. (2009), including *INCVOL, RETVOL*, *DISP*, log(1/*AGE*) and *TURN*. *INCVOL* is the standard deviation of the quarterly operating income over the 20 quarters prior to the earnings announcement quarter. *RETVOL* is the standard deviation of the announcing 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* is the dispersion of analyst forecasts in the 45-day period ending 3 days prior to the earnings announcement. *AGE* is the number of years the announcing 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 also include, but do not report, ten days of lagged abnormal returns up to two days prior to earnings announcement. We report in parentheses the Driscoll and Kraay (1998) robust t-statistics, which control for time series and cross-sectional correlation. ***, **, * represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively. To ease reading, we multiply the coefficients on *IMPACT* by 1,000 and divide the coefficients on *INCVOL* and *TURN* by 100.

	Dependent Variables: CAR [0,1]							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
DIVOP	-0.369***	-0.389***	-0.230*	-0.254**	-0.378***	-0.213*	-0.373***	-0.271**
	(-3.21)	(-3.33)	(-1.72)	(-2.05)	(-3.31)	(-1.69)	(-3.24)	(-2.42)
SUE		1.289***	1.312***	1.297***	1.289***	1.216***	1.290***	1.283***
		(18.97)	(20.04)	(16.88)	(18.55)	(11.18)	(18.92)	(18.86)
IMPACT		-0.032	0.136***	-0.024	-0.032	-0.021	-0.033	-0.012
		(-1.00)	(-3.45)	(-0.72)	(-1.00)	(-0.67)	(-1.02)	(-0.36)
NEWS		0.014	-0.012	0.0152	0.014	0.012	0.013	0.009
		(1.61)	(0.69)	(1.58)	(1.62)	(1.47)	(1.53)	(0.98)
INCVOL				0.071***				
				(-2.94)				
RETVOL					0.341			
					(0.05)			
DISP						-0.129		
						(-0.85)		
Ln(1/AGE)							0.072	
							(0.97)	
TURN								-0.364***
								(-3.84)
Lagged Returns	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Fixed Effects	No	No	Yes	No	No	No	No	No
Observations	19,750	18,584	18,584	16,158	18,584	8,957	18,584	18,584
R-square	0.002	0.054	0.061	0.061	0.061	0.052	0.061	0.062

Table 6: Post-Earnings Announcement Returns across Disagreement Groups

Panel A presents post-earnings announcement returns across the groups of investor disagreement. We first sort earnings announcements into two groups based on the values of the *DIVOP* measure, which is a dummy variable that takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. We then report the averages of cumulative abnormal returns (*CARs*) over the [2,10], [2,20], or [2,60] post-announcement windows for the two groups, where day 0 is the announcement day. For an earnings announcement, we estimate daily abnormal returns using the four-factor model where the coefficients are estimated in the 90-day window [-120,-31] ending 31 days prior to the earnings announcement. We also report the differences in *CARs*, and the associated two-sample t-statistics assuming unequal variances. Panel B presents the post-earnings earnings announcement returns for the two extreme groups of the changes in investor agreement. We calculate the *DIVOP* measure over the two weeks after earnings announcement. The "Agree → Disagree" group contains the announcements where *DIVOP* equals 0 in the pre-announcement period but 1 in the post-announcement period. The "Disagree → Agree" group contains the announcements where *DIVOP* equals 1 in the pre-announcement period but 0 in the post-announcement period. We also report the differences in *CARs* between the "Disagree → Agree" group and the "Agree → Disagree" groups and the associated t-statistics. ***, **, * represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Post-Earnings Announcement Returns Across Groups of Disagreement (%)				
	#obs	CAR [2,10]	CAR [2,20]	CAR [2,60]
Agree	10,760	-0.21	-0.27	-0.15
Disagree	8,987	-0.50	-0.72	-1.55
D–A		-0.29***	-0.45***	-1.41***
t-stat		-2.82	-2.90	-4.40
Panel B: Post-Earnings Announcement Returns Across Groups of Change in Disagreement (%)				
Agree → Disagree	4,532	-0.08	-0.03	0.23
Disagree → Agree	3,372	-0.71	-0.95	-1.46
DA – AD		-0.63***	-0.92***	-1.69***
t-stat		-3.91	-3.80	-3.39

Table 7: Regressions of Post-Earnings Announcement Returns

This presents the regressions of post-earnings announcement returns on the investor disagreement measure. The dependent variables are cumulative abnormal returns in the [2,60], [2,20], or [2,10] windows with respect to earnings announcement, where day 0 is the earnings announcement day. For an earnings announcement, we estimate daily abnormal returns using the four-factor model where the coefficients are estimated in the 90-day window [-120,-31] ending 31 days prior to the earnings announcement. 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's news articles and Twitter sentiment over the two weeks prior to earnings announcement, and 0 if there is agreement. We control for *IMPACT*, which measures the level of Twitter sentiment over the two weeks prior to earnings announcement; *NEWS*, which measures the sentiment of firm's news article over the two weeks prior to earnings announcement; and *SUE*, which is the standardized unexpected earnings measured as the difference between the actual earnings and expected earnings (based on IBES median estimates) divided by stock price, measured in percentage. We also include, but do not report, ten days of lagged abnormal returns up to two days prior to earnings announcement. We report in parentheses the Driscoll and Kraay (1998) robust t-statistics, which control for time series and cross-sectional correlation. ***, **, * represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively. To ease reading, we multiply the coefficients on *IMPACT* by 1,000.

	Dependent Variables: CAR [2,60]			**Dep: CAR[2,20]**	**Dep: CAR[2,10]**
	(1)	(2)	(3)	(4)	(5)
DIVOP	-1.795***	1.644***	-1.470***	-0.332*	-0.299***
	(-3.88)	(-3.52)	(-3.21)	(-1.90)	(-2.80)
SUE		-0.348**	-0.232	0.234**	0.247***
		(-1.99)	(-1.30)	(2.47)	(3.44)
IMPACT		-0.378***	-0.472***	-0.146***	-0.079***
		(-3.92)	(-4.73)	(-3.40)	(-2.82)
NEWS		-0.038	-0.095	-0.011	-0.013
		(-1.03)	(-1.59)	(-0.40)	(-0.57)
Lagged Returns	No	Yes	Yes	Yes	Yes
Firm Fixed Effects	No	No	Yes	Yes	Yes
Observations	19,746	18,582	18,582	18,582	18,582
R-square	0.073	0.077	0.080	0.064	0.032

Table 8: Regressions of Post-Earnings Announcement Returns: Control for Alternative Disagreement Measures

This presents the regressions of post-earnings announcement returns on the measures of investor disagreement. The dependent variables are cumulative abnormal returns in the [2, 60] window with respect to earnings announcement, where day 0 is the earnings announcement day. For an earnings announcement, we estimate daily abnormal returns using the four-factor model where the coefficients are estimated in the 90-day window [-120,-31] ending 31 days prior to the earnings announcement. 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's news articles and Twitter sentiment over the two weeks prior to earnings announcement, and 0 if there is agreement. We control for *IMPACT*, which measures the level of Twitter sentiment over the two weeks prior to earnings announcement; *NEWS*, which measures the sentiment of firm's news articles over the two weeks prior to earnings announcement; and *SUE*, which is the standardized unexpected earnings measured as the difference between the actual earnings and expected earnings (based on IBES median estimates) divided by stock price, measured in percentage. *INCVOL* is the standard deviation of the quarterly operating income over the 20 quarters prior to the earnings announcement quarter. *RETVOL* is the standard deviation of the announcing 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* is the dispersion of analyst forecasts in the 45-day period ending 3 days prior to the earnings announcement. *AGE* is the number of years the announcing 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 calculate the *SUV*, *MATO*, and *ΔTO* measures using the approaches in Garfinkel and Sokobin (2006). *MATO* (market adjusted turnover) is the average daily abnormal turnover in the [0,1] window, where daily abnormal turnover is the firm's daily turnover minus the daily turnover of all CRSP firms. *ΔTO* is calculated as *MATO* minus the announcing firm's average daily abnormal turnover in the [-54,-5] window. Standardized unexpected volume, *SUV*, is the abnormal volume in the [0,1] window where abnormal volume is constructed based on a market model-style model regression of volume on absolute valued returns estimated in the 50-day window ending 5 days prior to the announcement. We also include, but do not report, ten days of lagged abnormal returns up to two days prior to earnings announcement. We report in parentheses the Driscoll and Kraay (1998) robust t-statistics, which control for time series and cross-sectional correlation. ***, **, * represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively. To ease reading, we multiply the coefficients on *IMPACT* by 1,000 and divide the coefficients on *INCVOL*, *RETVOL*, *TURN, MATO* and *ΔTO* by 100.

Dependent Variables: CAR [2,60]								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
DIVOP	-1.774***	-1.697***	-1.363***	-1.724***	1.691***	-1.595***	-1.595***	-1.624***
	(-3.64)	(-3.73)	(-3.03)	(-3.75)	(-4.15)	(-3.58)	(-3.43)	(-3.47)
SUE	-0.442**	-0.386**	-0.492*	-0.359**	-0.345**	-0.351**	-0.350**	-0.347**
	(-2.16)	(-2.26)	(-1.87)	(-2.06)	(-1.98)	(-2.01)	(-2.00)	(-1.98)
IMPACT	-0.364***	-0.375***	-0.310***	-0.368***	-0.387***	-0.366***	-0.365***	-0.375***
	(-3.87)	(-3.94)	(-3.53)	(-3.83)	(-3.72)	(-3.75)	(-3.80)	(-3.90)
NEWS	-0.043	-0.036	-0.031	-0.024	-0.036	-0.039	-0.038	-0.038
	(-1.15)	(-0.97)	(-0.87)	(-0.64)	(-0.95)	(-1.05)	(-1.03)	(-1.02)
INCVOL	0.135							
	(1.36)							
RETVOL		-0.576						
		(-1.44)						
DISP			1.806***					
			(4.07)					
Ln(1/AGE)				-1.006***				
				(-3.46)				
TURN					-0.159			
					(-0.35)			
MATO						-0.058		
						(-0.61)		
ΔTO							-0.098	
							(-0.86)	
SUV								-0.035
								(-0.85)
Lagged Returns	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Fixed Effects	No	No	No	No	No	No	No	No
Observations	16,156	18,582	8,956	18,582	18,582	18,582	18,582	16,576
R-square	0.074	0.074	0.079	0.075	0.073	0.074	0.074	0.077

Table 9: Cumulative Abnormal Trading Volume across Groups of Disagreement

Panel A presents cumulative abnormal trading volumes (*CAV*) around earnings announcements across the groups of investor disagreement. We first sort announcements into two groups based on the value of the *DIVOP* measure, which is a dummy variable that takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is disagreement. We then report the averages of the cumulative abnormal volumes (*CAVs*) in the [0,1], [0,20], [2,10], and [2,20] windows for the two groups, where day 0 is the earnings announcement day. For an announcement, we use the methodology of Campbell and Wasley (1996) to daily abnormal trading volume where the coefficients are estimated in the 200-day window ending 45 days prior to the earnings announcement [-245,-45]. We also report the volume differences and the associated two-sample t-statistics assuming unequal variances. Panel B presents *CAVs* for the four groups of the changes in investor agreement. We calculate the *DIVOP* measures over the two weeks after earnings announcements. The "Agree → Agree" group contains the announcements where the *DIVOPs* equal 0 in both pre- and post-announcement periods. The "Agree → Disagree" group contains the announcements where *DIVOP* equals 0 in the pre-announcement period but 1 in the post-announcement period. The "Disagree → Agree" group contains the announcements where *DIVOP* equals 1 in the pre-announcement period but 0 in the post-announcement period. The "Disagree → Disagree" group contains the announcements where the *DIVOPs* equal 1 in both pre- and post-announcement periods. We also report the differences in *CAV* between the "Agree → Agree" group and the other three groups and the associated t-statistics. ***, **, * represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

	#obs	CAV [0,1]	CAV [2,10]	CAV [2,20]	CAV [0,20]
Panel A: Cumulative Abnormal Volume Across Groups of Disagreement (%)					
Agree	10,754	1.70	0.57	0.48	2.18
Disagree	8,977	2.78	1.22	1.25	4.03
D – A		1.08***	0.66***	0.77***	1.85***
t-stat		14.11	4.87	3.24	6.53
Panel B: Cumulative Abnormal Volume Across Groups of Change in Disagreement (%)					
(1) Agree → Agree	6,224	1.25	0.10	-0.26	0.99
(2) Agree → Disagree	4,530	2.32	1.21	1.50	3.82
(2) – (1)		1.07***	1.11***	1.77***	2.84***
t-stat		12.01	12.01	12.01	12.01
(3) Disagree → Agree	3,370	2.26	0.84	0.85	3.11
(3) – (1)		1.01***	0.74***	1.12***	2.13***
t-stat		9.43	9.43	9.43	9.43
(4) Disagree → Disagree	5,607	3.09	1.45	1.49	4.59
(4) – (1)		1.84***	1.35***	1.76***	3.60***
t-stat		19.24	19.24	19.24	19.24

Table 10: Regressions of Cumulative Abnormal Volume

This table presents regressions of cumulative abnormal volume after earnings announcement on the investor disagreement measures. The independent variables are cumulative abnormal volume over the [0,1], [2,10], [2,20], and [0,20] event windows, where day 0 is the earnings announcement day. For an announcement, we use the methodology of Campbell and Wasley (1996) to daily abnormal trading volume where the coefficients are estimated in the 200-day window ending 45 days prior to the earnings announcement [-245,-45]. *DIVOP* is the key independent variable in Models 1-3, which is a dummy that takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. In Models 4-8, the key independent variables are changes in investor disagreement. We calculate the *DIVOP* measure over the two weeks after earnings announcement. The "Agree → Disagree" group contains the announcements where *DIVOP* equals 0 in the pre-announcement period but 1 in the post-announcement period. The "Disagree → Agree" group contains the announcements where *DIVOP* equals 1 in the pre-announcement period but 0 in the post-announcement period. The "Disagree → Disagree" group contains the announcements where the *DIVOPs* equal 1 in both pre- and post-announcement periods. We control for *IMPACT*, which measures the level of Twitter sentiment over the two weeks prior to earnings announcement; *NEWS*, which measures the sentiment of firm's news articles over the two weeks prior to earnings announcement; and *SUE*, which is the standardized unexpected earnings measured as the difference between the actual earnings and expected earnings (based on IBES median estimates) divided by stock price, measured in percentage. We also include, but do not report, ten lagged daily abnormal volumes up to two days prior to earnings announcement. We report in parentheses the Driscoll and Kraay (1998) robust t-statistics, which control for time series and cross-sectional correlation. ***, **, * represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively. To ease reading, we multiply the coefficients on *IMPACT* by 1,000.

	Dependent Variables						
	CAV [0,1]				CAV[2,10]	CAV[2,20]	CAV[0,20]
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
DIVOP	1.080***	0.856***	0.187***				
	(11.14)	(9.78)	(2.96)				
Agree → Disagree				0.492***	0.996***	1.847***	2.339***
				(7.58)	(5.92)	(5.17)	(6.04)
Disagree → Agree				0.311***	0.447**	0.902**	1.212***
				(3.00)	(2.19)	(2.39)	(2.77)
Disagree → Disagree				0.519***	0.728***	1.255***	1.774***
				(6.11)	(3.68)	(2.90)	(3.74)
SUE		-0.054	-0.015	-0.013	-0.136	-0.284*	-0.298*
		(-0.96)	(-0.48)	(-0.42)	(-1.50)	(-1.87)	(-1.79)
IMPACT		0.239***	0.062**	0.062**	0.001	-0.029	0.033
		(5.52)	(2.26)	(2.27)	(0.14)	(-0.37)	(0.36)
NEWS		0.012	0.028**	0.027**	0.061**	0.131**	0.158***
		(1.64)	(2.48)	(2.39)	(2.03)	(2.29)	(2.58)
Lagged Volumes	No	Yes	Yes	Yes	Yes	Yes	Yes
Firm Fixed Effects	No	No	Yes	Yes	Yes	Yes	Yes
Observations	19,735	18,570	18,570	18,570	18,568	18,568	18,568
R-square	0.011	0.187	0.234	0.235	0.341	0.352	0.382

Table 11: Robustness Tests: Sentiment Measured Using the Naïve Bayesian (NB) Approach

This table presents robustness tests of earnings announcement return and trading volume using the Naïve Bayesian approach to measure the sentiment of Tweets and news instead of the Maximum Entropy approach. Panel A presents earnings announcement returns across prior disagreement groups. We first sort announcements into two groups based on the value of the *DIVOP* measure, which is a dummy variable that takes the value of 1 if there is disagreement between the firms' news articles and Twitter sentiment over the two weeks prior to earnings announcements, and 0 if there is agreement. We then report average cumulative abnormal returns (CARs) in the [0,1] and [-1,1] window for the two groups, where day 0 is the announcement day. For an earnings announcement, we estimate daily abnormal returns using the four-factor model where the coefficients are estimated in the 90-day window [-120,-31] ending 31 days prior to the earnings announcement. We also report the return differences between the two groups and the associated two-sample t-statistics assuming unequal variances. Panel B presents earnings announcement returns for the four groups of changes in investor agreement. We calculate *DIVOP* over the two weeks after earnings announcements and divide the sample earnings announcements into four groups. The "Agree → Agree" group contains the announcements where the *DIVOPs* equal 0 in both pre- and post-announcement periods. The "Agree → Disagree" group contains the announcements where *DIVOP* equals 0 in the pre-announcement period but 1 in the post-announcement period. The "Disagree → Agree" group contains the announcements where *DIVOP* equals 1 in the pre-announcement period but 0 in the post-announcement period. The "Disagree-Disagree" group contains the announcements where *DIVOP* equals 1 in both pre- and post-announcement periods. We also report the differences in *CARs* between the "Disagree → Agree" and the "Agree → Disagree" groups and the associated t-statistics. Panel C presents cumulative abnormal trading volumes (*CAV*) around earnings announcements across the groups of investor disagreement. We first sort announcements into two groups based on the value of the *DIVOP* measure. We then report the averages of the cumulative abnormal volumes (*CAVs*) in the [0,1], [0,20], [2,10], and [2,20] windows for the two groups, where day 0 is the earnings announcement day. For an announcement, we use the methodology of Campbell and Wasley (1996) to daily abnormal trading volume where the coefficients are estimated in the 200-day window ending 45 days prior to the earnings announcement [-245,-45]. We also report the volume differences and the associated two-sample t-statistics assuming unequal variances. Panel D presents *CAVs* for the four groups of the changes in investor agreement. We also report the differences in *CAV* between the "Agree → Agree" group and the other three groups and the associated t-statistics. ***, **, * represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

Panel A: Earnings Announcement Returns Across Groups of Disagreement (%)

	#obs	CAR [0,1]	CAR [-1,1]
Agree	11,008	0.28	0.39
Disagree	8,743	0.01	0.15
Disagree – Agree		-0.27**	-0.25**
t-stat		-2.49	-2.18

Panel B: Earnings Announcement Returns Across Groups of Change in Disagreement (%)

	#obs	CAR [0,1]	CAR [-1,1]
Agree → Disagree	4,570	0.53	0.68
Agree → Agree	6,438	0.10	0.19
Disagree → Disagree	5,305	0.15	0.26
Disagree → Agree	3,438	0.21	0.03
DA – AD		-0.74***	-0.71***
t-stat		-4.20	-3.90

Panel C: Cumulative Abnormal Volume Across Groups of Disagreement (%)

	#obs	CAV [0,1]	CAV [2,10]	CAV [2,20]	CAV [0,20]
Agree	11,008	1.78	0.61	0.54	2.32
Disagree	8,743	2.70	1.18	1.21	3.91
D – A		0.92***	0.51***	0.67***	1.59***
t-stat		11.97	4.18	2.79	5.56

Panel D: Cumulative Abnormal Volume Across Groups of Change in Disagreement (%)

	#obs	CAV [0,1]	CAV [2,10]	CAV [2,20]	CAV [0,20]
(1) Agree → Agree	6,438	1.35	0.17	-0.24	1.12
(2) Agree → Disagree	4,570	2.39	1.24	1.62	4.01
(2) – (1)		1.04***	1.08***	1.86***	2.90***
t-stat		10.81	7.87	6.56	8.61
(3) Disagree → Agree	3,438	2.16	0.61	0.32	2.47
(3) – (1)		0.81***	0.44**	0.55*	1.36***
t-stat		9.01	2.41	1.74	3.73
(4) Disagree → Disagree	5,305	3.06	1.56	1.78	4.84
(4) – (1)		1.70***	1.39***	2.02***	3.72***
t-stat		17.01	7.66	6.54	10.01

Table 12: Determinants of Divergence and Convergence of Opinion

This table presents logistic regressions using dummy variables for the divergence of opinion (Panel A) or convergence of opinion (Panel B) as dependent variables. In Panel A, the sample consists of earnings announcements in the "Agree → Agree" group and the "Agree → Disagree" group). The dependent variable is a dummy variable that equals 1 if an event is in the "Agree → Disagree" group, and 0 if an event is in the "Agree → Agree" group. For Panel B, the sample consists of earnings announcements in the "Disagree → Agree" group and the "Disagree → Disagree" group). The dependent variable is a dummy variable that equals 1 if an event is in the "Disagree → Agree" group, and 0 if an event is in the "Disagree → Disagree" group. The independent variables in Panels A and B include *Low|Ret|/HighVol*, which is a dummy variable that equals 1 if the earnings announcement is in the bottom quintile of the absolute earnings announcement return (CAR[-1,1]) and the top half of the earnings announcement trading volume (CAV[-1,1]), and 0 otherwise. *⊿TURN* is the change in turnover (*TURN*) around the earnings announcement, where we follow Berkman et al. (2009) and calculate pre-announcement turnover as average daily turnover measures over a 45-day period ending 10 days before the earnings announcement day, and post-announcement turnover as average daily turnover measures over a 30-day period starting 10 days after the earnings announcement day. *⊿RETVOL* is the change in return volatility around earnings announcement, where we calculate pre-announcement return volatility as the standard deviation of the announcing firm's daily excess returns relative to a value-weighted market index, over a 45-day period ending 10 days before the earnings announcement (Berkman et al. (2009)). We calculate post-announcement return volatility as the standard deviation of the announcing firm's daily excess returns relative to a value-weighted market index, over a 30-day period starting 10 days after the earnings announcement. ln(*ME*) is natural log of the market capitalization of the announcing firm at the end of the previous year. ln(*B/M*) is natural log of the book-to-market ratio of the announcing firm. Ret[-12, -2] is the buy-and-hold stock return of the announcing firm from month -12 to month -2, where month 0 is the month of announcement. Idiosyncratic volatility is standard deviation of the residuals of the market model of daily stock returns of the announcing firm in the one-year period ending in the month before earnings announcement. ln(1+#*inst*) is natural log of 1 plus the number of institutional shareholders of the announcing firm at the quarter-end before the earnings announcement. Earnings persistence is the coefficient in regressions of annual earnings on lagged annual earnings in the eight years ending in the year before the earnings announcement. |Accrual| is absolute value of the announcing firm's total accrual (earnings minus operating cash flows) of the year before the earnings announcement. We control for year fixed effects in all models. ***, **, * represent statistical significances at the 0.01, 0.05, and 0.10 levels, respectively.

	Panel A: Divergence of Opinions			Panel B: Convergence of Opinions		
	(1)	(2)	(3)	(6)	(7)	(8)
Low|Ret|/HighVol	0.532***	0.538***	0.289***	-0.168**	-0.160*	-0.073
	(5.96)	(6.00)	(2.59)	(-2.01)	(-1.91)	(-0.69)
ΔTURN		5.607	8.271		-1.116	-0.422
		(1.29)	(1.38)		(-0.27)	(-0.07)
ΔRETVOL		1.328	-0.544		4.093	3.775
		(0.53)	(-0.15)		(1.35)	(0.87)
ln(ME)			0.073			-0.266***
			(1.54)			(6.72)
ln(B/M)			-0.029			-0.043
			(-0.78)			(-1.17)
Ret[-12,-2]			0.084**			-0.014
			(2.15)			(-0.41)
Idiosyncratic Vol.			16.762***			-23.964***
			(5.83)			(-7.42)
ln(1+#Inst.)			0.607***			-0.193***
			(6.80)			(-2.72)
Earnings Persistence			0.063			-0.004
			(1.26)			(-0.06)
|Accrual|			0.295			-0.350
			(0.91)			(-0.96)
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,762	10,747	7,045	8,989	8,970	5,923
Pseudo-R^2	0.013	0.013	0.053	0.001	0.001	0.038