Bulls and Bears: Disagreement and Trading Volume around News Announcements†

Adam Booker
University of Arkansas
abooker@walton.uark.edu

Asher Curtis
University of Washington
abcurtis@uw.edu

Vernon J. Richardson
University of Arkansas
vrichardson@walton.uark.edu

May 3, 2018

ABSTRACT: We examine the association between disagreement and trading volume around news events using a novel measure of disagreement that overcomes two challenges Bamber et al. (2011) identify as facing earlier measures. Specifically, we measure disagreement based on heterogenous opinions about firm value of StockTwits users. We find strong results that both pre-existing disagreement and the change in disagreement following the earnings announcement are both associated with trading volume. We next provide novel evidence on the differential effects of attention and disagreement by examining the impact of more influential users in the StockTwits network. We find that disagreement between influential investors is associated with incrementally higher trading volume. Our measure of disagreement generalizes to other news events, consistent with the measure capturing disagreement about firm value, not just the short-term earnings prospects of the firm. Our results provide evidence consistent with the importance of both pre-existing heterogeneity between users of financial statement information and preliminary evidence that disagreement between individuals about financial information can be affected by online social networks.

Keywords: *Social Media; Disagreement; Volume; Earnings Announcements.*

JEL Codes: M40.

Data Available: *All data are available from the sources described in the text.*

---

† We appreciate comments from Pierce Crosby from StockTwits, Kris Allee, Mark Bradshaw, Remington Curtis, Ed deHaan, Sam Melessa, Sarah McVay, Derek Oler, Roy Schmardebeck and seminar and workshop participants at BYU Accounting Symposium, the University of Arkansas, and the University of Miami. We thank StockTwits for providing academic access to their data. We thank Arkansas High Performance Computing Center for advice and access to their computers. Asher Curtis thanks the Herbert O. Whitten Endowment for financial support. All errors remain our responsibility.

Address for correspondence:
Asher Curtis, Foster School of Business, University of Washington, Box 353226, Seattle, WA 98195-3200, USA. Phone: (206) 685-2813, email: abcurtis@uw.edu.

1. Introduction

Using a novel measure, we re-examine the role of disagreement in explaining the association between trading volume and earnings announcements. Our measure, using a combination of natural language processing and supervised machine learning, estimates disagreement from variation in users' posts on the stock-market-focused social media website StockTwits.

Beaver (1968) suggests trading volume following a news event captures changes in the expectations of individual investors. Thus, volume studies allow insight into how earnings and other information announcements are used differentially by investors. Since Beaver (1968), a large body of literature has examined the association between earnings announcements and trading volume focusing on analyst-forecast-based measures of disagreement (see Bamber, Barron, and Stevens 2011, for a review). Bamber et al. (2011, 464), however, note two important limitations with using analysts' forecasts as a proxy for disagreement. First, both discount rate and cash flow news are used to value a firm, and thus disagreement can be about more than the analyst forecast of one-year-ahead earnings. Second, analyst forecasts likely reflect a more homogenous set of opinions that understate the heterogeneity in the opinions of financial statement users.

Recent changes in information technology, including real-time online communication platforms that increase the speed and breadth of information dissemination, allow accounting disclosures to reach a broader audience (Drake, Roulstone, and Thornock 2012; Blankespoor, Miller, and White 2014a; Blankespoor, Miller, and White 2014b; Blankespoor, deHaan, and Zhu 2018; Bartov, Faurel, and Mohanram 2018). These changes increase the expectation of disagreement in the same way priors, private information searches, and interpretations of new information will likely vary more widely, and in turn exacerbate the issues raised by Bamber et al. (2011). For example, the beliefs of a fundamental trader who focuses on traditional inputs to firm value will likely differ from a momentum trader who generally undertakes their trading strategies with no weight on accounting information. Differences in the priors about the usefulness of accounting information are important, as theory suggests that disagreement is expected to vary around earnings and other news announcements either due to investors revising differences in prior beliefs, or reordering beliefs which can be due to differential interpretations of the news (Karpoff 1986; Kim and Verrecchia 1991; Kandel and Pearson 1995; Bamber et al. 2011).

Thus, given the recent technological changes of the internet and social media, we re-examine whether the link between earnings news and trading volume is increasing in disagreement about the news event. We develop a novel measure of disagreement that is based on general opinions posted to social media of firm value, drawing on the assessments of a heterogenous group of potential financial statement users. Specifically, we estimate measures of disagreement based on the posts of users on the investor-focused social media website StockTwits. StockTwits represents users with differing beliefs about what drives stock prices (for example users can self-identify as momentum or fundamental based traders) and users that differ in terms of their market sophistication with both novices and experienced individuals contributing opinions to the website. Using this data, we estimate disagreement by combining explicit user recommendations (posts that are flagged with "bull" or "bear" indicators) and the tone of the posts on StockTwits around financial disclosures. We combine these measures using a supervised machine learning approach which learns from the text in posts of users who accompany their posts with an explicit bullish (positive) or an explicit bearish (negative) recommendation. This approach allows us to combine insight drawn from textual analysis, using measures of tone similar to Antweiler and Frank (2004), and extend their approach by using a classification scheme that is based on the context of the posts.

When used together, our approach provides an intuitive measure of investor disagreement based on the discussions of the users of StockTwits that can be calculated within minutes of the release of new information and does not require analyst following.[1] We use our measure to identify both the level of investor disagreement prior to an information event and the change in disagreement following various information events.

The first information event we examine is earnings announcements. When we examine total trading volume, our results are generally in line with theory which suggests that trade following informative news, including earnings announcements, can result from either prior disagreement or the reordering of beliefs following the disclosure (e.g., Karpoff 1986). Specifically, around earnings announcements we find evidence that both the level of existing disagreement and the

---

[1] Our study differs from contemporaneous working papers who examine disagreement in social media as our measure is continuous rather than dichotomous. We examine disagreement on an intra-day basis centered on the time of the event. We examine news announcements in addition to earnings announcements, and we directly compare our social media based disagreement measure to analyst-based measures of disagreement (Giannini, Irvine, and Shu 2015; Cookson and Niessner 2016).

change in disagreement following the earnings announcement are positively and significantly associated with trading volume. When we examine adjusted trading volume, which is expected to control for liquidity trading (Kim and Verrecchia 1991), we continue to find evidence of positive associations between adjusted trading volume with both levels and changes of disagreement.

Our primary results are based on a broad sample of firm-announcements that include smaller firms and many firms with limited analyst coverage. To assess whether certain types of firms are associated with increased or attenuated evidence of a link between disagreement and market volume, we next provide cross-sectional estimates of the association between disagreement and trading volume. Specifically, we sort firms into portfolios based on differences in firm characteristics, including media attention, book-to-market, analyst following, and market capitalization. Consistent with our main findings, we observe that both prior disagreement and the change in disagreement are positively and significantly associated with trading volume across all portfolios, but the tests also reveal, however, an increased association between disagreement and trading volume when media attention is lower, analyst following is lower, and market capitalization is lower. These results are consistent with earnings news changing the expectations of individual investors when earnings news is a more prominent source of information about the firm.

We next compare our disagreement proxy with an analyst-based measure of disagreement. For a subsample of our firms with at least five analysts following the firm and social media coverage to calculate disagreement, we follow Bamber, Barron, and Stober (1997) and examine how measures of analyst disagreement are associated with trading volume. Specifically, we calculate measures of the prior dispersion in analyst forecasts, the change in dispersion after an earnings announcement, and a measure of the jumbling of analysts' beliefs following Barron (1995). In contrast to these earlier findings, we find inconsistent evidence in our time-period. We note that analyst dispersion measures are consistently and significantly correlated with volume, but not for abnormal volume at conventional levels.[2] We do find, however, that the social media measures of disagreement are consistently and positively associated with both volume and abnormal volume around the time of the earnings announcement. These results suggest that

[2] Our sample period differs significantly from these earlier studies because our period of investigation is much more recent. One reason that we may fail to find strong evidence for analyst-based measures of disagreement could be due to the increasing impact of online trading on the reaction to earnings news (Ahmed, Schneible, and Stevens 2003).

analysts' disagreement does not subsume the association between social media disagreement and trading volume and implies that disagreement around earnings announcements is not limited to disagreement about future earnings.

We next extend prior studies by investigating differences between attention and disagreement that are made possible by our measure. First, we examine whether the attention and disagreement of the most influential users in the social media network have incremental explanatory power over trading volume relative to disagreement alone. Understanding social influences on investor behavior is still relatively unexplored in the literature (Hirshleifer 2015). Shiller and Pound (1989) provide early survey evidence that influence of an investor over their peers might be possible due to word-of-mouth effects. We construct indicator variables for 1) a post from one of the top 1000 most influential users, and 2) disagreement between one or more influential users (measuring influence as centrality in the network or number of followers). We find evidence that posts made by influential users and disagreement between influential users are both associated with increased trading volume. Second, we examine heterogeneity between individuals, using a measure of information entropy. Information entropy is a commonly used measure in information science and in economics starting with Theil (1967) that measures the degree of uniformity within a sample. We find evidence that heterogeneity, or the lack of uniformity of the distributions of StockTwits users' backgrounds, may be one of the underlying factors leading to disagreement between investors.

Finally, we provide two tests of the robustness of our findings. First, we use the RavenPack database to examine disagreement around many diverse news events. RavenPack provides a classification of news events by topic at various levels of aggregation. We begin by examining all news events pooled together. Then, we examine the news events based on whether they are scheduled or unscheduled according to the RavenPack classification, and whether the news events relate to business or society topics. We find that the role of disagreement in all five of these samples is consistent with the role found around earnings announcements. That is, both prior disagreement and changes in disagreement are positively and significantly associated with trading volume following these news events.[3] Second, we investigate how disagreement and attention are

[3] For 145 of the events identified in the RavenPack taxonomy, we obtained a sufficient sample size to examine the responsiveness to each of these events separately. We find significant associations between disagreement and volume for many individual classifications of news events (not tabulated here).

associated with trading volume over shorter periods of time around an earnings announcement. Specifically, in our main analysis we examine investor disagreement over 24-hour windows around the earnings announcement (i.e., the prior level of disagreement is measured over the 12 hours preceding the earnings announcement, and the change uses the subsequent 12 hours of disagreement). We find similar results for attention when examining shorter windows (for example, ½ hour, 1 hour, 3 hours) around the earnings announcement, and find consistent evidence of disagreement being associated with trading volume using these shorter windows.[4]

Our results extend prior research that examines investor disagreement using the dispersion of analysts' forecasts (e.g., Barron 1995; Bamber et al. 1997; Diether, Malloy, and Scherbina 2002) more recent textual analysis of the tone of Yahoo! Message board posts (Antweiler and Frank 2004), and concurrent working papers examining social media (Giannini et al. 2015; Cookson and Niessner 2016). Each of these approaches has strengths and weaknesses. Whereas dispersion in analysts' forecasts measures explicit disagreement about future earnings outcomes, Bamber et al. (2011) and Dinh and Gajewski (2015) highlight that the dispersion in analysts' forecasts underrepresents the level of disagreement expected from the heterogeneous investor population, and that earnings-based measures do not capture disagreement on other aspects of stock valuation (such as disagreement about expected returns). Our results confirm that measures of disagreement based on more heterogeneous user groups provide incremental explanation of the theoretical link between investor disagreement and trading volume around earnings announcements. Our measure also provides insight into disagreement around other disclosures and news events.

Our joint measurement of user disagreement and user tone dispersion provides two extensions of prior measures. Relative to the dispersion in analysts' forecasts, our measure is much timelier (being an intra-day measure) than analysts' forecasts as currently provided (which are observed to update much less frequently in commonly available data sources). Specifically, the

[4] In untabulated results where we modify the window length over which disagreement is measured using longer windows such as 5 and 30-day windows, we continue to find similar associations for user disagreement and user tone dispersion. Shorter windows, such as those over 1 minute, are consistent but some of the t-statistics fall below levels of conventional significance. For the small sample of firms with sufficient data to calculate user disagreement in the minute before an earnings announcement, we find that disagreement is negatively associated with volume, but measuring disagreement over a 60-minute window we find a positive association with trading volume. We conjecture that these results suggest that there is likely a dynamic role of disagreement, or that one minute prior to the earnings announcement is too short of a period to measure disagreement. We leave to future research the more detailed investigation of this possible dynamic disagreement.

nature of our measure allows us to isolate the changes in disagreement using a short event window (i.e., in 24 hours relative to the 30-day windows used with analysts' forecasts). This removes the confounding effects of other events that are close in time to the earnings announcement and allows us to discern sources of disagreement. Second, our measure provides a better link between the tone of a post and the sentiment of the post as it would be interpreted by other social media users. This is important, as it helps allay concerns with measuring tone using common language dictionaries that may not capture the use of language in social media settings.

Our study contributes towards understanding how individual expectations change following the release of earnings and other news based on the theory that disagreement induces trade. Using an intuitive measure of disagreement available for a broad cross-section of firms, we find that the market reaction to earnings announcements and other disclosures and news events are associated with disagreement. Our results provide further evidence on the importance of both pre-existing heterogeneity between users of financial information and the importance of differences in the processing of accounting and other financial information. We also contribute a new measure of investor disagreement that is intuitive and applies to broader cross-sections of stocks and information events than analyst-based measures.[5] Specifically, we contribute a measure that combines both explicit indicators of user disagreement ("the bulls" versus "the bears") and the dispersion in user tone of the posts.

2. Institutional background and hypothesis development

2.1. Institutional background

StockTwits is an investor-focused social media website that is analogous to "Twitter for investors." For example, StockTwits and Twitter both restrict the posts to a maximum of 140-characters. Only around 11 percent of the users on StockTwits, however, share their ideas simultaneously on Twitter and StockTwits. In 2010, StockTwits was voted one of Time

[5] Note that we do not rely on, nor rule out, the possibility that investors that are external to the StockTwits community use this information in their trading decisions. Instead, we maintain the assumption that the bull versus bear recommendations of StockTwits users reflect their beliefs and that dispersion in those beliefs provides an empirical proxy for measuring the dispersion of beliefs of the heterogeneous investor (and potential investor) community reflected in trading volume.

magazine's 50 Best Websites. According to their website, StockTwits has over 300,000 active users whose content is viewed by 40 million people. StockTwits allows developers to connect their apps to their website, which enables additional content viewing. The volume of ideas, or posts, on StockTwits has been growing at an exponential rate since 2009. StockTwits is a place where investors, analysts, companies, and potential investors can learn, discuss, and participate in financial conversations.

StockTwits founders invented cashtags, a dollar sign and the ticker symbol (for example, $AAPL for Apple Inc.), which allow StockTwits users to reference specific stocks and provide an efficient method for aggregating ideas about a given stock. As of the writing of this paper, StockTwits' content aggregations are generally a top result when google searching a given cashtag.[6] Twitter, which had previously used hashtags to aggregate user contributions adopted the cashtag in 2012, three years after StockTwits. StockTwits also allows linking to Facebook, but only a fraction of a percent of the ideas in 2014 were linked to Facebook. StockTwits users can indicate their recommendation about a company on a given post as either "bullish" or "bearish" (see the Appendix for an example), combined with the cashtags, the bullish and bearish tags make identification of an individuals' stated opinion about a stock possible.

2.2. Hypothesis

According to Bamber et al. (1997), the idea of a relationship between investor disagreement and trading volume has been around since Louis Bachelier's thesis in 1900, and arguably since the work of Adam Smith in the late 1700s. In the empirical accounting literature, Beaver (1968) provides the first examination of the link between trading volume and earnings announcements and finds that companies that had more than 20 articles in the *Wall Street Journal* averaged about 1.5 times their 16-week-centered average volume in the week of an earnings announcement.

A large analytical literature considers how the disclosure of accounting information can cause increased trading volume. Kim and Verrecchia (1991) highlight three ways that investor information drives trading volume: 1) the noise of prior information causes differences in beliefs

[6] StockTwits' cashtag made it easier for search engines to find stock information as well. For example, at the time this study was conducted, a google search for the ticker "A" will not lead to the Agilent Technologies, but a search for $A leads to the latest StockTwits user posts about Agilent. StockTwits and Yahoo! Finance are often among the top results when searching on Google.com for ticker symbols.

prior to an earnings event to converge to a more precise equilibrium with new information, 2) new information causes investors to reorder, or jumble per Karpoff (1986), their beliefs, and 3) new information causes confusion and the dispersion of beliefs among investors increases. Bamber et al. (1997) provide empirical support for these predictions and use analysts' earnings forecasts to measure three categories of investor belief revision (they call it disagreement). Specifically, they look at 1) the dispersion of annual analyst forecasts around interim earnings announcements, 2) the change in dispersion of annual analyst forecasts around interim earnings announcements, and 3) the jumbling of annual analyst forecasts around interim earnings announcements.

Technological advances in internet and mobile technologies over the past few decades have made it easier for investors to connect and discuss earnings information. Antweiler and Frank (2004) show the tone dispersion of the text of messages about companies on internet message boards is related to trading volume. Even more recently, an average investor can use their phone to learn about and discuss trading ideas and also make trades in real time. Chen, De, Hu, and Hwang (2014) find that comments on the financial website *SeekingAlpha* are predictive of returns and earnings surprises. Bartov et al. (2018) find that Twitter sentiment prior to an earnings announcement is associated with the subsequent earnings surprise and market reaction to the earnings announcement. These papers suggest that the discussions of individuals on various online forums provide information that is correlated with actual market outcomes.

The content of the discussions, i.e., the online messages, have lexical and syntactic elements (titles, important sentences, etc.) that provide insight and opinion on current news stories that can be classified and coded using textual sentiment analysis. Importantly, similar to the difficulties with interpreting written text (Loughran and McDonald 2011; Guay, Samuels, and Taylor 2015), sentiment derived from textual analysis of user posts may not be indicative of a user's beliefs. The interpretation of the news and the news itself are often in a single message. It is possible for a user to have a bullish prediction while relaying a bad news story, which would cause the dispersion of textual tone to be less accurate about the beliefs of investors. The StockTwits dataset allows us to examine user-entered sentiment (explicit bullish or bearish user opinions). We predict that due to the lower levels of ambiguity in measuring the opinions of individuals by using their explicit bullish and bearish opinions allow for a proxy that better identifies investor disagreement.

Based on this literature, we hypothesize the following, stated in the alternative form:

H$_1$: *The prior dispersion of user opinions and the change in the dispersion of user opinions are positively associated with trading volume.*

There are several reasons why we may fail to find support for our hypothesis. First, if StockTwits users do not systematically reflect the wider market opinions of those investors' trading following an earnings announcement, we may not find support for our hypothesis. Second, if trading volume is only due to disagreement between investors prior to the announcement (similar to the conjecture in Karpoff 1986), then we will only find an association between prior dispersion of opinions, and not the change in the dispersion following the announcement. More importantly, we may not be able to change priors about investor disagreement, if the dispersion in user opinion simply reflects a previously identified proxy for investor disagreement, such as the negative tone of the posts (Antweiler and Frank 2004) or the dispersion of the analysts' forecasts (Bamber et al. 1997).

3. Variable measurement

3.1. Measurement of trading volume

We are interested in whether a measure of social media disagreement is associated with trading volume. We consider two approaches to measuring trading volume, vol_{jt}, for the earnings announcement of firm j at time t.

First, we calculate turnover, $turn_{t-1,t+1}$, as the natural log of the cumulative turnover in the three-days around an earnings announcement (days [-1,1]) divided by the number of shares outstanding. We use the natural log of turnover because trading volume is positively skewed (Ajinkya and Jain 1989). We use the three-day window because prior research has shown that a majority of the trading volume reaction to earnings occurs within this time. For each firm j we measure turnover as:

$$turn_{t-1,t+1} = \log\left(\Sigma_{\mathrm{j}=-1}^{\mathrm{j}=1}\frac{vol_{t+j}}{shrout_{t+j}}\right) \qquad (1)$$

where, t is the date of the first day following the earnings announcement such that for earnings announcements which occur after the market closes, t is the first trading day after the calendar date of the earnings announcement.[7] Note that by measuring the level of turnover, we control for cross-sectional differences in the number of shares available to trade.

Second, we calculate an adjusted turnover measure to control for cross-sectional differences in average turnover. Specifically, we calculate $adj.turn_{t-1,t+1}$ as the mean-adjusted shares traded in the three days around an earnings announcement divided by the number of shares outstanding:

$$adj.turn_{t-1,t+1} = \log\left(\Sigma_{\text{j}=-1}^{\text{j}=1} \frac{vol_{t+j}}{shrout_{t+j}} - \frac{\overline{vol_{\tau}}}{\overline{shrout_{\tau}}}\right) \tag{2}$$

where, we calculate $\overline{vol_{\tau}}$ as the median three-day share-weighted sum $\left(\frac{vol_{\tau-1}}{shrout_{\tau-1}} + \frac{vol_{\tau}}{shrout_{\tau}} + \frac{vol_{\tau+1}}{shrout_{\tau+1}}\right)$ of volume for each of the prior 249 trading days.

3.2 Measurement of sentiment

To measure disagreement from the text of discussions on StockTwits, we use a machine learning approach that uses paragraph vectors as described in Le and Mikolov (2014). The paragraph vector learning model incorporates a neural network to map the individual words used in a passage of text as well as the sum of the words used in a passage of text to semantic vector spaces. For example, a user posts the following: “i am bullish about earnings”. The words “am”, “bullish”, “about”, and “earnings” are used to train the word vector for “i” while simultaneously a paragraph vector for the whole phrase is being trained. The words “i”, “bullish”, “about”, and “earnings” are used to train the word vector for “am” at the same time the paragraph vector is being further trained. Conceptually, the paragraph vector acts as a memory of the words used together in a body of text. Figure 1 provides a graphical representation of the neural network that is used as the basis of this method.

Our method uses a simple neural network with one hidden layer that has a linear activation function. We restrict the vocabulary, or dictionary, used to train the model, to words that are used at least 50 times in our sample. After parsing out punctuation, converting proper nouns, removing cashtags, links, and mentions, and converting the posts to a lower case, we had

[7] Our classification of pre and post earnings announcements uses I/B/E/S where incorrect time-stamps in the I/B/E/S database could induce noise into our measures. However, deHaan, Shevlin, and Thornock (2015, Online Appendix Table IA.4) provide evidence that in recent periods the maximum error rate for using I/B/E/S to classify earnings announcements into before, during, or after hours is well below 10% and has been declining over time.

a vocabulary size of 65,500 words.[8] We use the continuous bag of words implementation (CBOW) of the paragraph vector method, which is appropriate for smaller passages of text like those found on StockTwits. Concatenation of CBOW and distributed memory (DM) did not significantly affect classification accuracy.[9] We gave the model an initial learning rate of 0.025 and decreased the learning rate by 5 percent on each training epoch. We experimented with different samples of posts: 1) training the full set of 37,037,180 ideas, 2) only ideas that used cashtags, and 3) only ideas with cashtags in the CRSP-I\B\E\S merged data that were posted in the 12 hours around earnings and did not find significant differences in accuracy. We used the full set classification for this study.

We used 100 dimensions for each paragraph vector. Each of the dimensions is a continuous number between 0 and 1. The paragraph vectors were used to parameterize a logistic regression model where the independent variables were the paragraph vectors and the dependent variables were the StockTwits user-identified sentiment (i.e., bull and bear flags). We use balanced logistic regression to account for the difference in the number of bullish and bearish training posts (there are many more post flagged as bulls than bears). We calculate the level of bullishness with the parameterized model as follows:

$$Bullishness = \Pr(Bullish) - \Pr(Bearish)$$

Using a continuous measure of sentiment helps to identify the statements that are neutral as neutral rather than as opposed to being classified as either bull or bear posts. We use a 50% threshold in cross validation tests on the training data (*Bullishness* > 0 is 1, *Bullishness* < 0 is -1). We code posts that are user classified as bullish with 1 (representing 100% chance of the user being bullish) and those classified as bearish with -1 (representing 100% chance of the user being bearish). Even if the user is using neutral language the explicit classification acts as an indicator in addition to the words used. That is, someone reading StockTwits posts identify the

[8] We experimented with lower thresholds for inclusion into the vocabulary but did not find a significant difference in classification accuracy.
[9] The distributed memory implementation of paragraph vectors is the inverse of the CBOW implementation. In the CBOW implementation, the model predicts a word given its context. In the DM implementation, the model predicts a context given a word.

intent of the user-identified posts based on the bull or bear flag. Consistent with other users not perfectly knowing the intent of unidentified posts, we use the *Bullishness* for these posts.

Panel A of Table 1 displays the cross-validation metrics for our classification method. The accuracy consistently averaged around 65% for the levels of cross validation that we used (2, 5, 10, 25, 50 groups). We use a binary average to measure the f1 score, precision and accuracy, with bullish posts representing the positive results. The precision was consistently around 87% while the recall was around 65.5%, which is an artifact of the cross-validation sample having more bullish posts than bearish posts. Figure 2 shows the distribution of bullishness for the non-classified sample of posts. The training sample is consistent with the paragraph-vector-classified sample as there are more bullish than bearish posts.[10]

As a further validation of our model, we randomly selected 20,000 posts from the StockTwits dataset and had users on Mechanical Turk classify these messages to further investigate the source of error. Workers on Mechanical Turk categorized each of the 20,000 randomly selected posts as either bullish, bearish, or neutral. 10,000 posts were StockTwits user classified as bullish and 10,000 as bullish. We had 5 different workers categorize each post. In Panel C of Table 1 we examine human categorizations in two ways: first we calculate the aggregate classification rates for all worker 100,000 classifications (full sample), and second, we examine the classifications for which a more then 50% majority chose a given classification (group majority). The full sample classification provides insight about the full sample while the group majority classification provides insight into how the posts would commonly be interpreted. Prior literature has adopted the group majority method for classifying sentiment (Peng and Park 2004).

Workers were 4 times more likely to misclassify a bearish post as bullish than a bullish post as bearish. Bearish posts had a majority classification of being neutral. Our measure of *bullishness* is a continuous measure of the sentiment of a post. As seen in Figure 2, many posts are classified as being near neutral. This is also consistent with the worker classification in which 33% for the full sample is classified as neutral.

[10] The large spike in the distribution at zero reflects posts with non-text features such as URLs and images that are removed from the database.

3.2. Measurement of investor disagreement

Our primary independent variable of interest is the dispersion of user opinions. We base our first measure of dispersion of opinion on the StockTwits bull versus bear flag. Not all posts on StockTwits are tagged with a bull or bear flag, thus for each post, i, we code posts where the user selects bullish as $bull_{it} = 1$ and those selected as bearish as $bull_{it} = -1$, and those that do not select a bull or bear are classified using an estimate of sentiment based on a machine learning algorithm that uses a paragraph vector approach (Le and Mikolov 2014). The algorithm uses a neural network to map the combination of words used in each post to a numerical vector representing the content of the post. We use these vectors with the user-classified set of posts to set the parameters of a logistic regression model. We then use these parameters to derive the level of bullishness or bearishness of each post. The outputs range between +1 and -1. We measure the standard deviation of explicit bull versus bear opinions over the 12-hour period preceding the earnings announcement as follows:

$$\sigma(bull)_{t-12,t} = \sqrt{\frac{\Sigma_i\left(bull_{it} - \overline{bull_{it}}\right)^2}{n-1}} \qquad (3)$$

where t is the time of the earnings announcement, and n is the total number of unique users that posted in the 12-hour window. To remove the effect of individuals who post multiple times, we use the average for each unique user each time-period instead of the dispersion between posts to better align with the construct of disagreement between investors. We refer to this measure as the prior level of user dispersion as it measures the level of opinion dispersion among users in the 12 hours leading up to the earnings announcement. We measure the change in the dispersion of user opinions by measuring the dispersion in bull versus bear flags in the 12-hour window following the earnings announcement and take the difference of opinions, $\Delta(bull)_{t-12,t+12}$, as user dispersion after the earnings announcement less that before the announcement:

$$\Delta(bull)_{t-12,t+12} = \sigma(bull)_{t,t+12} - \sigma(bull)_{t-12,t} \qquad (4)$$

3.3. Measurement of social media attention

We also measure the attention of users on StockTwits to each news event. The level of attention, or posting activity about a firm increases significantly around the earnings announcement (Curtis, Richardson, and Schmardebeck 2016). We measure attention as the log of the number of users posting about a company on StockTwits in the 12 hours prior to the announcement (#users) and the log of the change in the number of users posting around the announcement (Δ#users).

3.4. Measurement of entropy

Prior literature has found a positive relationship between investor attention and trading volume (Antweiler and Frank 2004; Curtis et al. 2016; Blankespoor et al. 2018). For information-based trading to be increasing in the level of investor attention, investors must have different reactions to the same information. Barron, Schneible, and Stevens (2017) find evidence that the diversity of institutional shareholders is positively associated with trading volume, and Cookson and Niessner (2016) find evidence that the diversity of the posters based on differences in self-reported investor type on StockTwits is positively associated with trading volume. We add to this literature by examining the roles of investor diversity and attention.

StockTwits users volunteer information about their trading strategy, experience, and holding periods. We use this information to construct a measure of the diversity of the contributing users' perspectives and backgrounds. Specifically, we use entropy[11] to construct a composite measure that is increasing in the uniformity of the distribution of users' backgrounds. That is, entropy increases as equally-weighted diversity increases.

The measure of entropy that we use was developed in Shannon (1948) as a calculation of the expected number of bits to encode information given a sample. In our study, the sample is the set of the backgrounds of users' that post around a given event (for example, the category of users' level of experience could be encoded as novice, intermediate, professional, or undeclared). Our

---

[11] This measure is also known as Shannon Entropy or information entropy. For simplicity we will refer to this as entropy in this section of the paper.

entropy measure is increasing in the uniformity of the distribution of users' backgrounds and is calculated as:

$$H = \sum_{i=0}^{n} p_i * \log_2(p_i) \qquad (5)$$

where H is the entropy and p_i is the prior probability of a user having the background classification i. In this study, we use three different categories of StockTwits users' background to construct the measure of the heterogeneity of users that are paying attention to an announcement. The categories are the following: 1) trading strategy ($H_{trading\ strategy}$), 2) experience ($H_{experience}$), and 3) holding period ($H_{holding\ period}$). StockTwits users have the option to volunteer information about their background in these categories. For example, to construct the measure of entropy for users' level of experience, we use the sample of unique users to construct the prior probability of a post being from a user with a professional, intermediate, novice, or undeclared level of experience. We use the sum of the entropies for each of the categories of StockTwits users' backgrounds to serve as a proxy for the diversity of users that pay attention to a given announcement:

$$H_{total} = H_{trading\ strategy} + H_{experience} + H_{holding\ period} \qquad (6)$$

3.5. Measurement of control variables

To control for macroeconomic trends that lead to increased volume, we include the turnover of the shares traded in the market for the given company on the day of the earnings announcement (log(market turn)). To control for level of information in the earnings surprise, we include the absolute value of the difference between the consensus analyst forecast in the 45 days before the announcement and the actual earnings per share (log(|sue|)). To control for changes in price, we include the absolute value of the return in the days around the announcement (log(|return|)). To control for the size of the firm, we include the market value of the stock of the company multiplied by the shares outstanding (log(size)). Following Bamber et al. (1997), we use the log of these variables in our regression models in order to reduce the effects of skewness (we refer to them as log(market turn), log(|sue|), log(|return|), log(size)).

3.5. Measurement of analyst dispersion

For a subset of our analysis, to control for disagreement between analysts we follow Bamber et al. (1997). We calculate the natural log of the standard deviation of analyst forecasts (σ(analyst)).

Following Bamber et al. (1997), we measure the jumbling of forecasts (jumbling) as the natural log of 1.1 less the Pearson correlation coefficient for analysts' estimates in the 45 days prior and the 30 days post announcement, intuitively jumbling measures the proportion of forecasts of individual analysts that change relative to the distribution of forecast changes. We also calculate the change in the standard deviation of analyst estimates by taking the standard deviation in the period after earnings less the standard deviation of analyst estimates in the period before ($\Delta\sigma$(analyst)). We require each firm-quarter observation to have five or more analyst earnings forecasts in the prior and post periods to be included in these disagreement measures, which lowers the number of observations included in this subsample significantly.

4. Empirical analysis

4.1. Data and sample

We collect financial information and social media posts for a sample of public firms with quarterly earnings announcements between July 2010 and 2015. We begin in July 2010 as StockTwits did not have a mechanism for users to explicitly declare their sentiment before that time, which we require to classify the text of the posts in the same time periods of the user-classified posts. We end on December 31, 2015, as it is the final year with this data available.[12] We collect daily market data (returns, volume, and prices) from CRSP (The Center for Research on Securities), earnings data from I/B/E/S (Thompson Reuters Institutional Brokers Estimate System), and social media data from StockTwits.[13] We make use of both the posts and the related metadata on StockTwits for all posts made in our sample period. Each record includes an optional user-defined forecast (a bull or bear flag) and the cashtags that are mentioned in the body of the idea that the user contributes to StockTwits.

We discuss our sample selection criteria in Table 2. In Panel A, we report the sample selection in terms of the number of firm-quarters included in the sample. We identify 91,589 unique firm-quarters with earnings announcements in the I/B/E/S Detail History database between July 1, 2010, and December 31, 2015. Of these observations we can use 21,187 US-based firm-announcements in our main analyses. The loss in sample size is due to the requirement of 249 days

[12] Using data over the period July 9, 2009, to December 31, 2015, does not alter the conclusions drawn from the study.
[13] See http://stocktwits.com/developers/docs/start for a description of the raw data provided by StockTwits.

of trading data available on CRSP (5,741 observations) and due to our requirement that there is activity on StockTwits for the 24 hours surrounding the earnings announcement.[14]

In Panel B, we outline the underlying number of ideas shared on StockTwits (often labeled "tweets" or posts). The StockTwits dataset has 37,037,180 total user ideas between July 1, 2010, and December 31, 2015, approximately 15,654 posts per day, with the majority occurring during trading hours. We use cashtags and the time-stamps for each idea to merge ideas to firm-quarter earnings announcements. We require that a firm has at least 2 posts in the 12-hour window before the earnings announcement and at least 2 posts after the earnings announcement. This results in a sample of 140,384 posts in the 12-hour window prior to the earnings announcement, and 393,043 posts in the 12 hours following the earnings announcement. This suggests that the 24 hours surrounding earnings announcements accounts for approximately 1.44% of all posts on StockTwits, with a substantial increase of posts following the earnings announcement. Visually, we confirm that idea activity is higher following an earnings announcement in Figure 3. Specifically, we display the minute-by-minute quantity of posts on StockTwits for the two hours surrounding the earnings announcement. The figure displays prominent spike in the aggregate amount of social media volume at the time of the earnings announcement, followed by a reduction in volume in the hour after the earnings announcement. This figure is similar to that in Curtis et al. (2016, 219) who use MarketIQ's smart velocity measure, a measure of social media buzz.

4.2. Descriptive statistics

We present descriptive statistics in Table 3. As expected, median-adjusted volume is significantly positive around earnings announcements, on average, and displays considerable variation, consistent with prior literature. Both the prior user opinion dispersion and the change in user opinion dispersion have means and medians greater than zero, consistent with increased discussion among a heterogenous set of financial statement users. Part of the reason for this is that there is wide variation across the observations in the sample. Several firms display a reduction in

---

[14] Using alternative sample definitions, we find 90,938 I/B/E/S firm-announcements can be merged with CRSP and StockTwits. Of these firm-announcements, 24,194 have sufficient social media data, and 12,239 have sufficient analyst data. In the social media dataset, 18,943 have insufficient analysts forecast data to estimate analyst variables (78.3%), leaving 5,251 firm-announcement observations for which we can calculate dispersion measures for both social media and analysts. Of the 12,239 that have sufficient analyst data, 6,988 do not have sufficient social media data (57.1%). We note, however, that most of the observations missing social media data are concentrated in the early years of StockTwits (untabulated).

user opinion dispersion following an earnings announcement (e.g., the 10th percentile is -0.17), and other firms display an increase in user opinion dispersion (e.g., the 90th percentile is 0.24). Variation both in the prior user opinions and opinions following the earnings announcement is expected to vary considerably given that we observe a large heterogeneity among StockTwits users.[15] We also provide a graphical measure of the degree of self-described trading horizon, strategy and experience in Figure 4. StockTwits users have the option to choose to publicly display these self-descriptions on their profile. Whereas Figure 4 suggests that many users do not answer these descriptions, those that do have heterogeneous trading strategies, horizons and experience.

4.3. Tests of hypothesis

Our hypothesis predicts that our user opinion proxy allows us to identify investor disagreement on the day of the earnings announcement. Following theory, we anticipate that dispersion in user opinions is positively associated with trading volume. As such, we consider the following regression model:

$$vol_{jt} = \alpha + \beta_1 \sigma(bull)_{t-12,t} + \beta_2 \Delta(bull)_{t-12,t+12} + \beta_3 \#users\ pre + \beta_4 \Delta\#users + \sum Controls_j \quad (7)$$

where, vol_{jt} is trading volume measured as either *turn* or *adj. turn* for firm j measured over the three days around an earnings announcement at time t. We are most interested in the investor disagreement variables based on the explicit opinions of users controlling for the tone of their posts. Following H_1, we expect that $\beta_1 > 0$ and $\beta_2 > 0$, and based on the results in Antweiler and Frank (2004), we expect $\beta_3 > 0$ and $\beta_4 > 0$.

We report estimates of Equation (7) in Table 4 based on our cross-sectional sample with 21,647 firm-quarter observations. In Column (1), with the dependent variable of *turn*, the coefficient on prior user disagreement, is positive and significant, $\beta_1 = 0.229$ ($p < 0.001$) as is

[15] StockTwits has users self-report various characteristics about themselves including three levels of experience: novice, intermediate, and professional; and 6 different trading styles: fundamental, technical, momentum, growth, global macro, and value. Thus, we observe a large degree of stated differences in trading styles that could affect how information is processed differently between the users of StockTwits. We also examine analyst EPS dispersion as a measure of the disagreement among analysts about expected earnings per share. The variation in analysts EPS is high suggesting considerable variation between analysts, as well as significant change following an earnings announcement. We find that the correlation between the StockTwits user opinion- and analyst-based measures of disagreement to have a low positive correlation, around 0.3, suggesting that they are distinct proxies.

the change in user disagreement, $\beta_2 = 0.262$ ($p < 0.001$). Similar results are seen for the model in Column (3) where we examine median-adjusted volume. Specifically, when regressed on mean-adjusted volume the coefficient on prior user dispersion, is positive and significant, $\beta_1 = 0.015$ ($p < 0.001$), and the change in user dispersion, $\beta_2 = 0.018$ ($p > 0.001$), is significant at conventional levels. In Columns (2) and (4) we report models that include firm-fixed effects, and find similar results. However, the coefficients are notably smaller, consistent with part of the association between disagreement and trading volume being driven by other firm characteristics.

The control variables included in the model are generally consistent with predictions based on prior literature. Change in attention around the earnings announcement is positively associated with both measures of vol_{jt}. As expected, the level of market-wide trading volume is positively associated with trading volume around earnings announcements, consistent with variation in how market-wide, or macroeconomic news, can affect the level of trading volume.[16] Similarly, the coefficients on the absolute value of the returns and the unexpected earnings are positive and significant. The log of firm size is negative in models all models without firm-fixed effects, which suggests that conditional on the level of attention to the earnings announcement (which is positively correlated with size) larger firms have incrementally lower trading volume.[17]

The coefficient on the log of the absolute value of SUE changes from a significantly negative association (Column 1, with a coefficient of -0.019) with volume to a significantly positive association (Column 2, with a coefficient of 0.013) when firm fixed effects are included. The use of fixed effects calculates within firm estimates of the relation between SUE and returns. Thus, some of these firms will have much higher SUEs that are not surprising relative to a small SUE for another firm that is much lower but more surprising. The firm fixed effects remove between firm differences, so the surprise must be higher than the average for that firm to be considered a

[16] The inclusion of market-wide volume plays a role like date fixed effects. A date fixed effect would identify the average volume of firms reporting earnings on the same day. The difference between the fixed effect and the time-varying level of market volume is that macro-economic news is expected to have an incremental effect on the level of trading volume for all firms, not just those with earnings announcements occurring on the same day.

[17] On the surface this coefficient appears to be opposite to expectations (Beaver 1968). In more recent periods, however, post Reg-FD, many managers are expected to provide all material disclosures to the market, which may lower the need to rely on earnings announcements for small firms when gathering information about the firm. Our conjecture is that in our sample, the increased trading volume for large firms might relate to greater levels of attention paid to large cap stocks in modern markets.

surprise. This suggests that there may be omitted firm-earnings characteristics that confound the earnings-return relation which can be controlled for using firm fixed effects.

To summarize, we find results that are consistent with the predictions of theory on the role of investor disagreement in creating trading volume. Specifically, we present a proxy for investor disagreement based on the opinions of users on StockTwits and find that it is positively associated with median-adjusted, a measure of abnormal trading volume, around earnings announcements.

4.4. Cross-sectional analysis

In this section we examine cross-sectional differences between the disagreement-volume relation across portfolios of firms with different characteristics. In Table 5, we divide the sample into quintiles of aggregate event volume (*AEV*), analyst following, and marker capitalization. The disagreement-volume relation remains statistically significant for all portfolios we report in our cross-sectional tests. We compare the differences in the cross-sectional coefficients using the formula for the z-statistic derived in Clogg, Petkova, and Haritou (1995, 1276).[18] We test the difference between the 1st and 5th quintiles and find that the coefficient on *σ(bull)* is significantly greater for smaller firms, firms with lower analyst following, and firms with lower *AEV*. The coefficient on *Δσ(bull)* is greater for smaller firms and firms with lower analyst following.

Panel A of Table 5 displays the cross-sectional results of model 5 divided by quintile of *AEV*. *AEV* is the RavenPack rolling-window measure of the cumulative number of relevant events in a rolling 91-day window. We limit the full sample to observations with an AEV observation between 1 and 5 days prior to their earnings announcement. The coefficient on *σ(bull)* is significantly greater in the 1st quintile relative to the 5th quintile of the cross-sectional tests (p=0.015), but not significantly different for *Δσ(bull)*. It is likely that firms with higher AEVs have disclosed information that preempts earnings. As such it would make sense for both of our measures of disagreement to be different across extreme quintiles. The coefficient on *Δσ(bull)* in the 1st quintile is greater than the coefficient in the 5th with a p-value of 0.15, an insignificant difference based on traditional levels of significance.

[18] Specifically, we follow Clogg et al. (1995) estimate: $\frac{\beta_1-\beta_2}{\sqrt{se_1^2+se_2^2}}$.

We find evidence of significant differences between the top and bottom quintiles of analyst following for both of our measures of disagreement in Panel B ($p < 0.01$). Cashtags on StockTwits are often top results in google searches. This result provides evidence that crowdsourcing financial statement information and understanding among individuals has a more prominent role when traditional information intermediaries are less available. We see a similar finding in Panel C which is sorted on market capitalization. The smallest firms have the largest disagreement-volume relation and the large firms have the smallest. The differences between quintiles 1 and 5 are highly significant in the Panel C cross-sectional results.

5. Further analysis

5.1. Comparison to dispersion in analysts' forecasts

Prior literature has used characteristics of annual analyst earnings forecasts in the months around an earnings announcement to proxy for investor disagreement. The length of this timespan allows analysts to impound feedback from managers and the market into their revised earnings forecasts. It is possible, although unlikely, that analyst forecast dispersion could confound our analysis.[19]

Nonetheless, if we are unable to present evidence of incremental association of our measure over that of the existing literature, our results are of limited interest. Thus, to control for the elements of investor disagreement that are explained by analyst forecasts we use the model of investor disagreement in Bamber et al. (1997), which constructs measures of the dispersion of annual analyst forecasts in the 45 days before and 30 days after a quarterly earnings announcement. Including these measures in the model yields:

$$\begin{aligned} vol_j = \alpha + \beta_1 \sigma(\text{bull})_j + \beta_2 \Delta\sigma(\text{bull})_j + \beta_3 \sigma(\text{analyst})_j + \beta_4 \Delta\sigma(\text{analyst})_j \\ + \beta_5 \text{jumbling}_j + \sum Controls_j \end{aligned} \tag{8}$$

[19] In untabulated analysis, we find that the number of analysts that revise their forecasts on the same day increases significantly after the earnings announcement but is more spread out prior to the earnings announcement. As such, it is possible that during our sample period, the speed at which analysts revise their forecasts allows for timelier measures of changes in analyst dispersion following an earnings announcement, at least relative to the samples used in prior research.

where the subscript j represents an earnings announcement from firm j.[20]

We report estimates of Equation (8) in Table 6. We limit this analysis to the subsample of our firms with at least five analysts following the firm to maintain consistency with Bamber et al. (1997) to examine whether including measures of analyst disagreement explain part or all of the association between the dispersion in user opinion and trading volume.[21] As discussed above following Bamber et al. (1997), we calculate measures of the prior dispersion in analyst forecasts, the change in dispersion after an earnings announcement, and a measure of the jumbling of analysts' beliefs following Barron (1995). In Column (1) and (4) we reconcile to prior literature by examining the sample of firms with available analyst variables and find results that are consistent with prior literature. In Column (2) we require both social media and analyst variables and find that only prior analyst dispersion is associated with trading volume (the change in analyst disagreement and jumbling are no longer significant at conventional levels).

In Column (3) we report the sample of firms with both analyst and social media measures with turnover as the dependent variable, and then in Column (5) we report estimates with median-adjusted volume as the dependent variable. We find that the inclusion of the analyst forecast-based measures of dispersion do not subsume the association between the social media-based measures of disagreement with turnover in Column (3), but in contrast, in Column (5) we find that the negative coefficients on disagreement are no longer significant at conventional levels. For example, in Column (3) when regressed on mean-adjusted volume the coefficient on prior user disagreement is positive and significant, $\beta_1 = 1.844\ (p < 0.001)$ and the coefficient on the change in user disagreement is also positive and significant, $\beta_2 = 1.303\ (p < 0.001)$. Note, however, that the coefficients have declined marginally from the earlier evidence presented in Table 4, which could be due either to the smaller sample size that requires greater levels of analyst coverage, or due to the inclusion of the analyst dispersion measures. Similarly, the prior level of

[20] Both LVOL and MDAJVOL are centered on the next trading day for announcements made after the market closes.
[21] In untabulated analysis, we find that the correlation between analyst-forecast-based and our measure of disagreement have a small positive correlation. The sample of firms for which we can measure both changes in social media disagreement and changes in analyst forecast dispersion is relatively small due for two main reasons: StockTiwts discussions cover many more smaller firms in recent years, but StockTwits coverage is much smaller than I/B/E/S in the early years.

dispersion in the tone of the user posts and the change in the dispersion of the tone of user posts continue to have positive and significant coefficient estimates.

5.2. Weighting investor disagreement by user influence

In Table 7, we examine whether user influence appears to moderate the association between disagreement and trading volume by including indicator variables for 1) a post from one of the top 1,000 most influential users, and 2) disagreement between 1 or more of the top 1,000 most influential users. Investor influence may affect the reaction to news if other investors are persuaded by the opinions of the influential investor. We consider two user-specific characteristics that could serve as a proxy for influence, 1) the centrality of the user to the network and 2) the number of followers in the network. We report the results in Table 7. We find evidence that our measures of influence are positively associated with trading volume. In all columns, we continue to find both prior disagreement, and changes in disagreement have positive and significant associations with trading volume. As such, our results appear robust to the possibility that there is heterogeneity in the influence of users.

5.3. Investor disagreement by user influence

In Table 8, we examine the importance of user heterogeneity in explaining the association between disagreement and attention by including a measure of entropy and interactions between entropy and attention. In columns 1-4 of Table 8, we find a positive and significant relationship between heterogeneity and adj. turn. In column 5, when we include our measure of dispersion, the coefficients on the main effect flip negative. The correlation between the coefficients of σ(bull) and entropy is -0.68. The correlation between the coefficients of Δσ(bull) and Δentropy is -0.67. The F-test for the differential effects of heterogeneity in column 4 has a F-stat of 129.86. In column 5, when we add σ(bull) and Δσ(bull) in column 5, the F-stat decreases to 0.19. Our measures of disagreement subsume the differential effect of investor heterogeneity and volume, providing evidence that different perspectives may be related to the dispersion of opinions.

5.4. Information events and investor disagreement

In Table 9, we provide estimates of the association between the prior level and changes in investor disagreement around different news events. We use the RavenPack database to identify the dates of various news events classified in the RavenPack event taxonomy. RavenPack classifies news events in real-time using the headlines of news articles from traditional and web-based media sources. We define a firm-announcement at the daily level by including each unique firm-event observation in the RavenPack database.

In Column (1) of Table 9, we examine all news events, that is, all days in which a firm has an event coded in the RavenPack database and find that the role of investor disagreement is consistent in these broader set of news events with the role found around the narrower news event of an earnings announcement. In Column (2) we report the association between prior level and changes in investor disagreement around scheduled news events and in Column (3) unscheduled news events. The associations are similar for the disagreement proxies for both types of news events, however, the intercepts are significantly different, with scheduled news having an intercept indistinguishable from zero for scheduled news but a significantly positive intercept for unscheduled news. This difference is consistent with unscheduled news leading to greater levels of trading activity given the level of disagreement.

Finally, in Columns (4) and (5) we report the associations splitting news into the events classified by RavenPack as business news and those into society news. Note that to be included in our sample, the society news events must reference the firm. We find similar associations between these two news classifications. Overall, our results suggest that our measure of investor disagreement has theoretically consistent associations with trading volume around many important firm-specific news events.

5.5. The timeliness of the measurement window of investor disagreement

In Table 10 we examine the robustness of the association between our investor disagreement proxies and trading volume when altering the measurement window for the posts used in the disagreement proxies based on user opinions and the tone of user posts. We consider shorter window lengths (½ hour, 1 hour, 3 hours, and 6 hours) to examine robustness to our choice of a

12-hour window used in our main analysis.[22] We report the results in Table 9, reporting only the turnover models, and find consistent results for all windows around the earnings announcement.

6. Conclusion

We use a new measure of disagreement to test theories about how earnings announcements are associated with trading volume. Our measure is based on the explicit bull (positive) and bear (negative) recommendations made by participants in the investor-focused social media website StockTwits. Our measure is available for a larger sample than traditional disagreement proxies and outperforms existing disagreement measures based on variation in analysts' forecasts. Our results are consistent with theory which suggests that variation in trading volume following informative news can result from either disagreement about prior beliefs or the reordering of beliefs following the disclosure. Consistent with the predictions of theory, we find that disagreement between investors on social media is positively associated with abnormal trading volume following earnings announcements.

In addition to providing a useful measure of investor disagreement for future research, we also provide preliminary evidence on the role of influence within a social network. We find that when two or more prominent members of the StockTwits network post comments that disagree after the earnings announcement, trading volume is incrementally higher on average. Future research could examine the role of investor influence in how accounting information is interpreted and re-interpreted. For example, we do not know whether influential investors affect the dissemination, or even the interpretation, of accounting disclosures.

[22] Figure 1 suggests that social media activity peaks just following the time of the earnings announcement. In untabulated results we examined both much shorter (1 minute, 2 minutes, 5 minutes) and much longer windows (1 day, 5 days, 30 days) and find positive associations between trading volume around earnings announcements and our two disagreement measures, the statistical power, and statistical significance is lower for much shorter windows.

References

Ahmed, A. S., R. A. Schneible, and D. E. Stevens. 2003. An Empirical Analysis of the Effects of Online Trading on Stock Price and Trading Volume Reactions to Earnings Announcements. *Contemporary Accounting Research* 20 (3):413-439.

Ajinkya, B. B., and P. C. Jain. 1989. The behavior of daily stock market trading volume. *Journal of Accounting and Economics* 11 (4):331-359.

Antweiler, W., and M. Z. Frank. 2004. Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards. *The Journal of Finance* 59 (3):1259-1294.

Bamber, L. S., O. E. Barron, and D. E. 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.

Bamber, L. S., O. E. Barron, and T. L. Stober. 1997. Trading Volume and Different Aspects of Disagreement Coincident with Earnings Announcements. *The Accounting Review* 72 (4):575-597.

Barron, O. E. 1995. Trading Volume and Belief Revisions That Differ among Individual Analysts. *The Accounting Review* 70 (4):581-597.

Barron, O. E., R. A. Schneible, and D. E. Stevens. 2017. The changing behavior of trading volume reactions to earnings announcements: Evidence of the increasing use of accounting earnings news by investors. *Forthcoming, Contemporary Accounting Research*.

Bartov, E., L. Faurel, and P. S. Mohanram. 2018. Can Twitter Help Predict Firm-Level Earnings and Stock Returns? *Forthcoming, The Accounting Review*.

Beaver, W. H. 1968. The Information Content of Annual Earnings Announcements. *Journal of Accounting Research* 6:67-92.

Blankespoor, E., E. deHaan, and C. Zhu. 2018. Capital market effects of media synthesis and dissemination: evidence from robo-journalism. *Review of Accounting Studies* 23 (1):1-36.

Blankespoor, E., B. P. Miller, and H. D. White. 2014a. Initial evidence on the market impact of the XBRL mandate. *Review of Accounting Studies* 19 (4):1468-1503.

Blankespoor, E., G. S. Miller, and H. D. White. 2014b. The Role of Dissemination in Market Liquidity: Evidence from Firms' Use of Twitter™. *The Accounting Review* 89 (1):79-112.

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* 27 (5):1367-1403.

Clogg, C. C., E. Petkova, and A. Haritou. 1995. Statistical Methods for Comparing Regression Coefficients Between Models. *American Journal of Sociology* 100 (5):1261-1293.

Cookson, J. A., and M. Niessner. 2016. Why Don't We Agree? Evidence from a Social Network of Investors. *Available at SSRN: https://ssrn.com/abstract=2754086*.

Curtis, A., V. J. Richardson, and R. Schmardebeck. 2016. Investor Attention and the Pricing of Earnings News. In *Handbook of Sentiment Analysis in Finance*, edited by G. Mitra and X. Yu. London: Optirisk Systems Ltd.

deHaan, E., T. Shevlin, and J. Thornock. 2015. Market (in)attention and the strategic scheduling and timing of earnings announcements. *Journal of Accounting and Economics* 60 (1):36-55.

Diether, K. B., C. J. Malloy, and A. Scherbina. 2002. Differences of Opinion and the Cross Section of Stock Returns. *The Journal of Finance* 57 (5):2113-2141.

Dinh, T. H., and J.-F. Gajewski. 2015. Trading Volume, Heterogeneous Expectations, and Earnings Announcements. *Journal of Behavioral Finance* 16 (4):327-343.

Drake, M. S., D. T. Roulstone, and J. R. Thornock. 2012. Investor Information Demand: Evidence from Google Searches Around Earnings Announcements. *Journal of Accounting Research* 50 (4):1001-1040.

Giannini, R. C., P. J. Irvine, and T. Shu. 2015. The convergence and divergence of investors' opinions around earnings news: Evidence from a social network.

Guay, W. R., D. Samuels, and D. J. Taylor. 2015. Guiding through the fog: Financial statement complexity and voluntary disclosure. *Available at SSRN: http://ssrn.com/abstract=2564350*.

Hirshleifer, D. 2015. Behavioral finance. *Annual Review of Financial Economics* 7:133-159.
Kandel, E., and N. D. Pearson. 1995. Differential Interpretation of Public Signals and Trade in Speculative Markets. *Journal of Political Economy* 103 (4):831-872.
Karpoff, J. M. 1986. A Theory of Trading Volume. *The Journal of Finance* 41 (5):1069-1087.
Kim, O., and R. E. Verrecchia. 1991. Market reaction to anticipated announcements. *Journal of Financial Economics* 30 (2):273-309.
Le, Q., and T. Mikolov. 2014. Distributed Representations of Sentences and Documents. Paper read at Proceedings of The 31st International Conference on Machine Learning.
Loughran, T. I. M., and B. McDonald. 2011. When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks. *The Journal of Finance* 66 (1):35-65.
Peng, W., and D. H. Park. 2004. Generate adjective sentiment dictionary for social media sentiment analysis using constrained nonnegative matrix factorization. *Urbana* 51:61801.
Shannon, C. E. 1948. A Mathematical Theory of Communication. *Bell System Technical Journal* 27 (3):379-423.
Shiller, R. J., and J. Pound. 1989. Survey evidence on diffusion of interest and information among investors. *Journal of Economic Behavior & Organization* 12 (1):47-66.
Theil, H. 1967. Economics and information theory.

Figure 1
Paragraph Vector Training Model

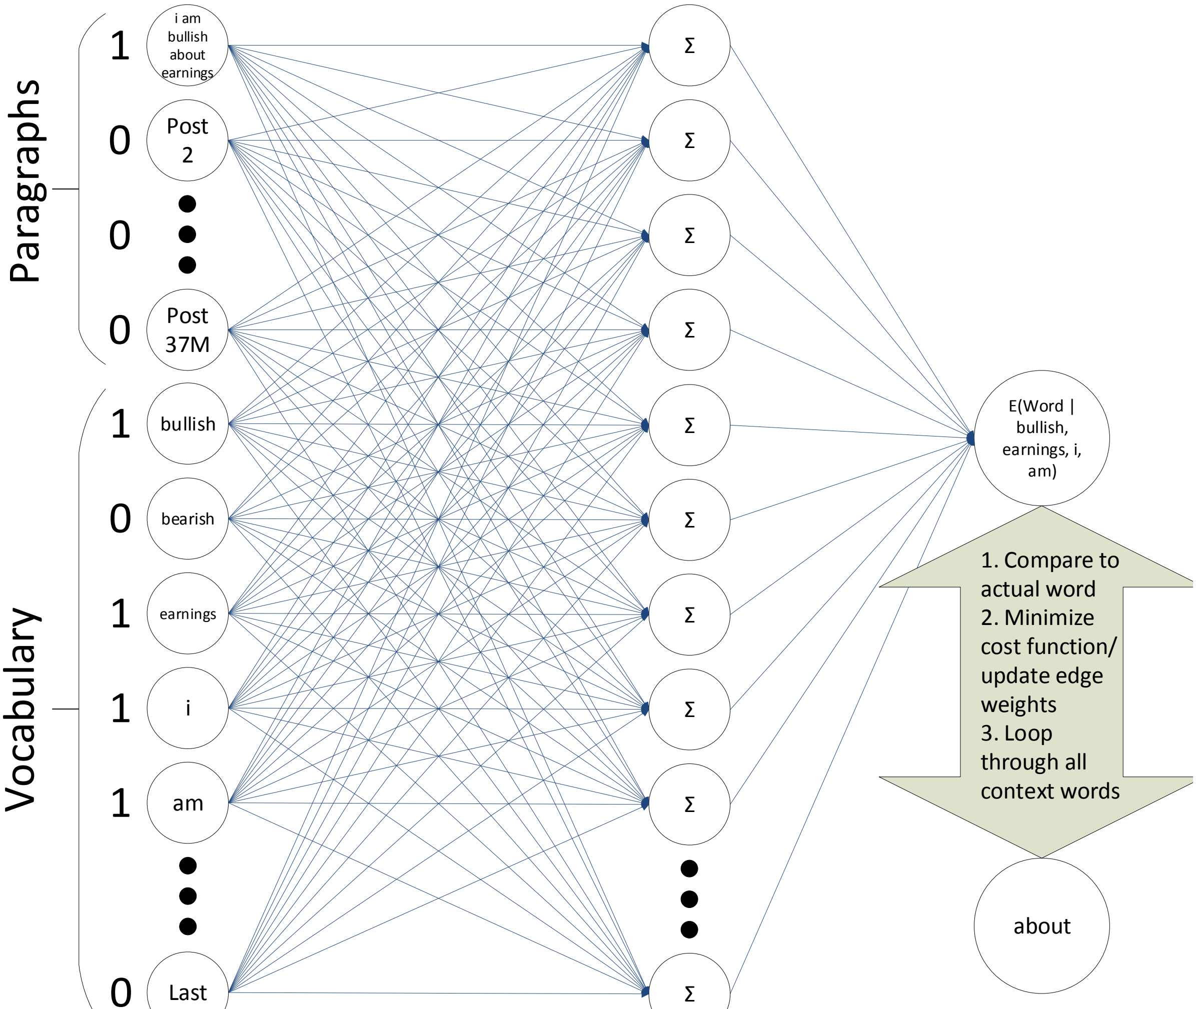

$$\begin{bmatrix} Edge\ Weight_{11} & \cdots & Edge\ Weight_{1p} \\ \vdots & \ddots & \vdots \\ Edge\ Weight_{v1} & \cdots & Edge\ Weight_{vp} \end{bmatrix}$$

Notes: This is an example of part of the training for the post "i am bullish about earnings." The edge weights for the first post in the network are updated as the weights for each word in the post are updated. The paragraph vectors that represent each post are the edge weights between the currently activated paragraph vector input and the hidden layer. Le and Mikolov (2014) note that this method allows each paragraph (StockTwits posts in the case of this study) to act as a memory of the words used. Each context is used to predict a given word. The paragraph is trained in all contexts.

Figure 2
Bullishness Distribution

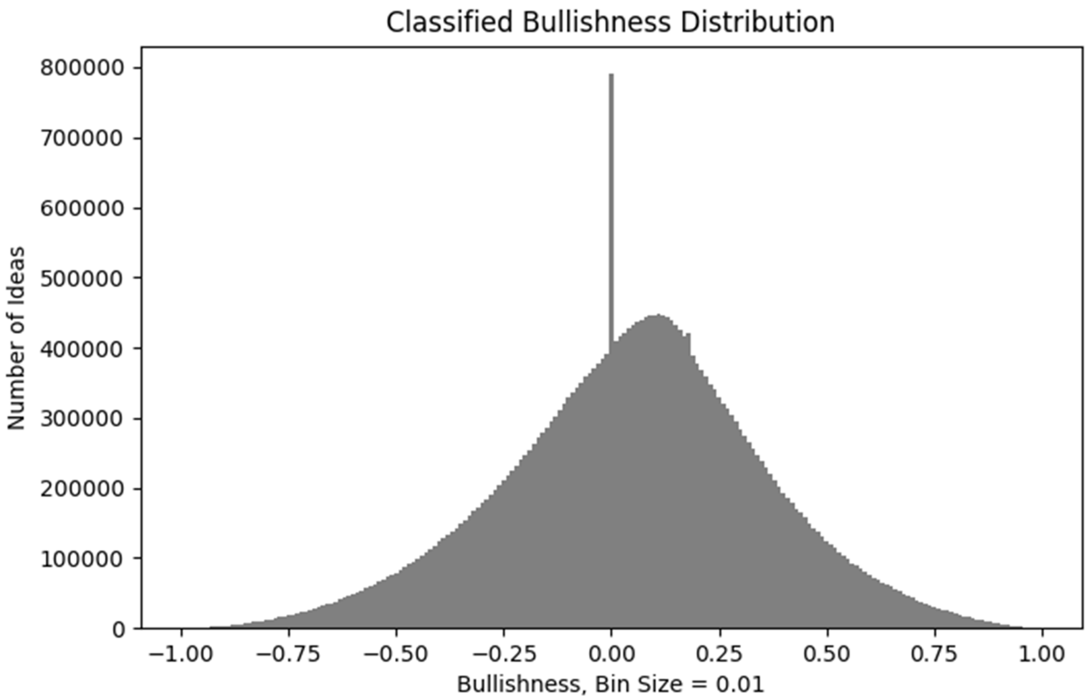

Notes: This figure is a histogram of our measure of *Bullishness*. This plot describes the sample of 30,959,238 ideas that were classified by the paragraph vector algorithm.

Figure 3
Social Media Volume around an Earnings Announcement

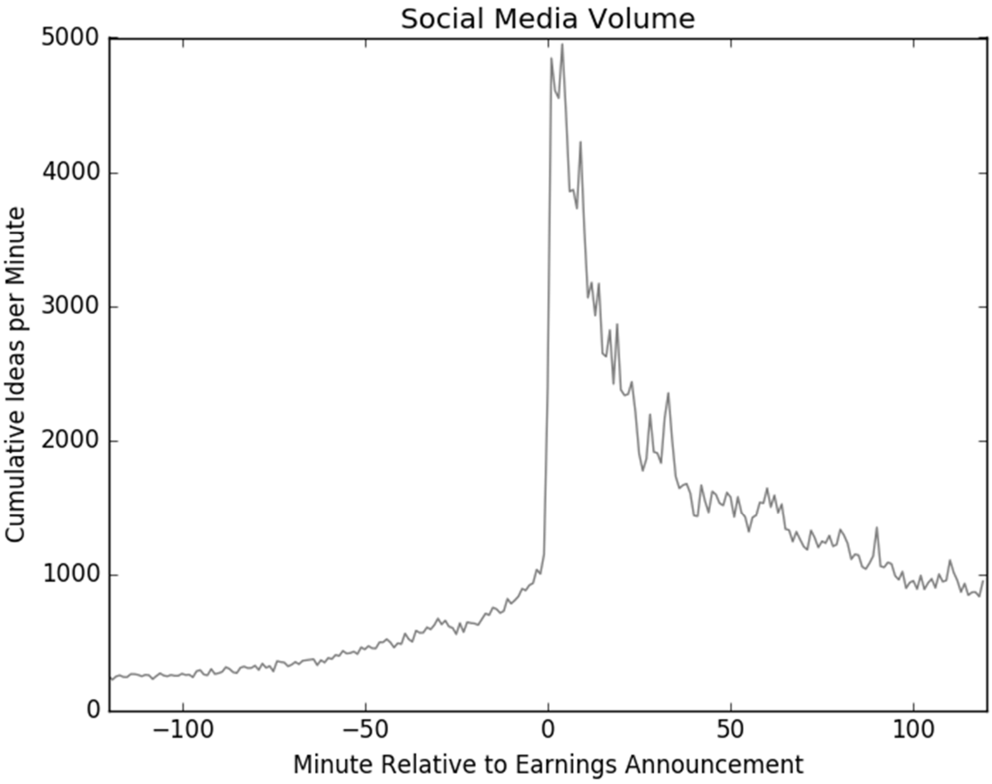

Notes: This figure plots aggregate post volume by minute around the earnings announcement. Time 0 is the I/B/E/S time of the earnings announcement.

Figure 4
Social Media Volume around an Earnings Announcement

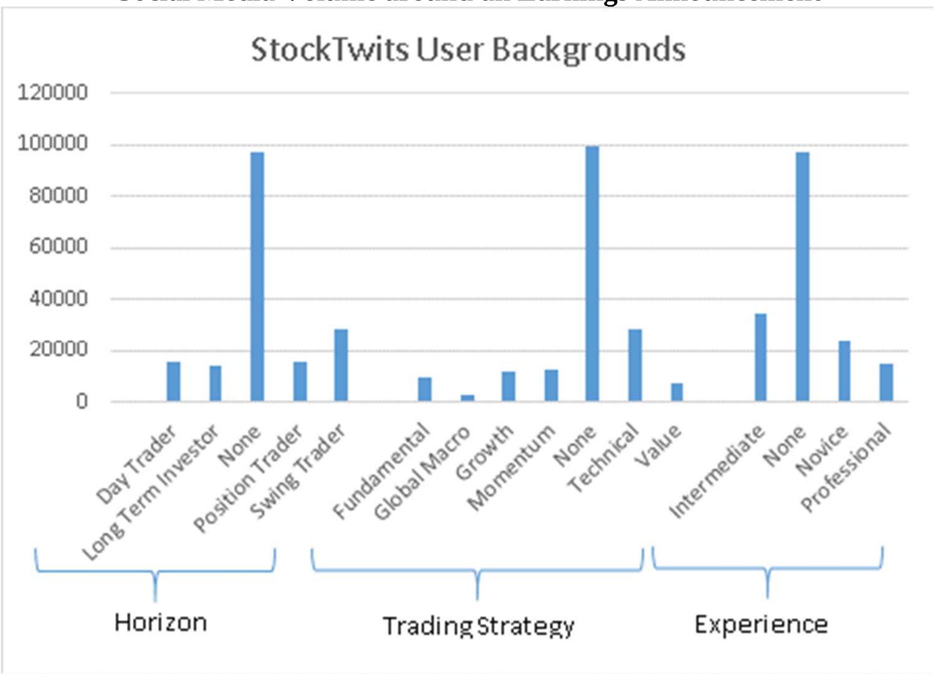

Notes: This figure plots heterogeneity of StockTwits users' self-described trading horizon, strategy and experience.

Table 1
Sentiment Analysis Validation

Panel A: Cross Validation

#folds	accuracy	f1 score	precision	recall
2	0.649823	0.750059	0.875753	0.655918
5	0.649799	0.750029	0.875777	0.655859
10	0.649862	0.750086	0.875792	0.655937
25	0.649846	0.750072	0.875785	0.65592
50	0.64983	0.750057	0.875782	0.6559

Panel B: 10-Fold Cross Validation

holdout group#	accuracy	f1 score	precision	recall
1	0.649639	0.749911	0.875296	0.655948
2	0.650019	0.75033	0.876231	0.656063
3	0.649581	0.749869	0.875948	0.655518
4	0.649379	0.749768	0.875245	0.655757
5	0.649994	0.750101	0.875964	0.655863
6	0.650931	0.750886	0.875365	0.657402
7	0.64964	0.749687	0.875316	0.655593
8	0.64982	0.750209	0.876994	0.655451
9	0.650116	0.750365	0.875984	0.656256
10	0.649503	0.74973	0.875572	0.655516

Panel C: Mechanical Turk Classification of 20,000 Random StockTwits Posts

StockTwits user identified	bullish			bearish		
Mechanical Turk classified	bullish	neutral	negative	bullish	neutral	negative
Cumulative classification	0.483	0.400	0.117	0.303	0.401	0.296
>50% of 5 workers threshold classification	0.469	0.334	0.045	0.183	0.335	0.218

Notes: Panel A shows 2, 5, 10, 25, and 50-fold cross validations for the full set of paragraph vectors. We use random sampling to constrict the groups for cross validation. The measures displayed in Panel A are an average of the number the cross validation for the given number of folds. Panel B shows the results from each individual cross validation for the 10-fold measurement in Panel A. Panel C displays the human classifications of StockTwits posts. The ratios are calculated with respect to the individual samples of bullish and bearish samples. For example, at 10,000 bullish posts from StockTwits with 5 workers classifying each post the number of workers that classified posts bullish that were also StockTwits user identified as bullish would be 10,000 * 0.483 * 5 = 24,150 out of a total 50,000 possible worker classifications for this portion of the sentiment set. The >50% qualification is only represented if the majority agrees on the classification. Each ratio is out of 10,000 in this row. The 100 less the sum of ratios in this column gives the percentage of posts for which workers did not agree on the classification.

Table 2
Sample selection

Panel A	
Unique I/B/E/S earnings announcements	91,589
Merge with CRSP (249 prior trading days not available)	(14,413)
Merge with Social Media (at least 2 conversations in 12 hours before announcement and posts from at least 2 unique users in the 12 hours after the announcement)	(55,989)
Total	21,187
Panel B	
Full dataset from July1, 2010 to December 31, 2015	
Social media observations	37,037,180
#Unique users	204,143
#User-identified bullish posts	5,160,405
#User-identified bearish posts	1,281,537
Posts and users in 24 hours centered on earnings	
Social media observations 12 hours before announcement	140,384
Social media observations 12 hours after announcement	393,043
#Unique users 12 hours before announcement	15,021
#Unique users 12 hours after announcement	22,018

Notes: This table presents the sample selection criteria and resulting sample size for the various analyses reported in this paper. The sample period is from July 1, 2010, to December 31, 2015. This is the period that the bullish or bearish indicators were available on StockTwits.

Table 3
Descriptive statistics

Variable	N	**mean**	**variance**	**p10**	**p50**	**p90**
$turn_{t-1,t+1}$	21,187	-3.03	0.89	-4.14	-3.05	-1.83
adj. $turn_{t-1,t+1}$	21,187	0.05	0.01	0.00	0.02	0.10
σ(bull)	21,187	0.23	0.02	0.06	0.20	0.46
Δσ(bull)	21,187	0.03	0.03	-0.17	0.03	0.24
#users	21,187	1.57	0.81	0.69	1.39	2.77
Δ#users	21,187	0.69	0.71	-0.38	0.69	1.79
log(market turn)	21,187	-3.65	0.04	-3.88	-3.65	-3.41
log(|return|)	21,187	-3.32	1.59	-4.92	-3.14	-1.93
log(size)	21,187	21.71	3.53	19.27	21.72	24.14
log(|sue|)	21,187	-6.40	2.90	-8.56	-6.51	-4.29

Notes: $turn_{t-1,t+1}$ is the log share turnover in the 3 trading days around an earnings announcement; adj.$turn_{t-1,t+1}$ is the turnover in the 3 trading days around an earnings announcement less the median 3 day moving average turnover in 249 preceding trading days; σ(bull) is the dispersion of user sentiment in the given period before an earnings announcement; Δσ(bull) is the change dispersion of user sentiment in the given period around an earnings announcement; #users is a count of the number of individuals posting during this time period. log(market turn) is the natural log of the volume in the primary exchange in which the stock for the given company is traded; log(|return|) is the absolute value of the return in the 3 trading days around an announcement; log(size) is the natural log of the market capitalization of the company; log(|sue|) is the natural log of the difference between quarterly consensus EPS forecasts and the actual EPS, using the I/B/E/S Summary database.

Table 4
Trading volume and investor disagreement around earnings announcements

$$vol_j = \alpha + \beta_1 \sigma(\text{bull})_j + \beta_2 \Delta\sigma(\text{bull})_j + \beta_3 \#\text{users}_j + \beta_4 \Delta\#\text{users}_j + \sum Controls_j$$

	$turn_{t-1,t+1}$		$adj.turn_{t-1,t+1}$	
	(1)	(2)	(3)	(4)
constant	2.170***	0.901**	0.481***	0.462***
	(8.69)	(2.49)	(16.97)	(5.79)
σ(bull)	0.229***	0.186***	0.015***	0.021***
	(28.71)	(32.75)	(13.20)	(16.11)
Δσ(bull)	0.262***	0.206***	0.018***	0.023***
	(32.89)	(35.87)	(17.32)	(18.06)
#users	0.295***	0.291***	0.029***	0.039***
	(21.43)	(36.10)	(10.59)	(14.30)
Δ#users	0.120***	0.109***	0.010***	0.013***
	(12.32)	(21.45)	(7.80)	(11.04)
log(market turn)	0.225***	0.525***	0.001	0.033***
	(4.08)	(17.48)	(0.16)	(5.84)
log(|return|)	0.184***	0.116***	0.014***	0.010***
	(34.34)	(33.90)	(15.01)	(17.79)
log(size)	-0.134***	-0.027*	-0.016***	-0.007**
	(-13.71)	(-1.66)	(-11.26)	(-2.05)
log(|sue|)	-0.019***	0.013***	-0.002***	0.000
	(-3.08)	(3.96)	(-3.09)	(0.30)
Year f.e.	Yes	Yes	Yes	Yes
Firm f.e.	No	Yes	No	Yes
Observations	21,187	21,187	21,187	21,187
Adj. R^2	0.398	0.311	0.194	0.147
Within R^2		0.311		0.147
Between R^2		0.379		0.109

Notes: $turn_{t-1,t+1}$ is the log share turnover in the 3 trading days around an earnings announcement; $adj.turn_{t-1,t+1}$ is the turnover in the 3 trading days around an earnings announcement less the median 3 day moving average turnover in 249 preceding trading days; σ(bull) is the dispersion of user sentiment in the given period before an earnings announcement; Δσ(bull) is the change dispersion of user sentiment in the given period around an earnings announcement; #users is a count of the number of individuals posting during this time period. log(market turn) is the natural log of the volume in the primary exchange in which the stock for the given company is traded; log(|return|) is the absolute value of the return in the 3 trading days around an announcement; log(size) is the natural log of the market capitalization of the company; log(|sue|) is the natural log of the difference between quarterly consensus EPS forecasts and the actual EPS, using the I/B/E/S Summary database. Intraday measures are split based on the earnings announcement time and cover a 12-hour window pre and 12-hour window post the time of the announcement, for announcements after 4 PM EST, measures using daily data are centered on the next trading day. We report *p*-values in parentheses based on standard errors that are clustered by firm. *p<0.1 **p<0.05 ***p<0.01.

Table 5
Trading volume and investor disagreement around news events: Cross-sectional analysis

$$adj.turn_{j,t-1,t+1} = \alpha + \beta_1 \sigma(\text{bull})_j + \beta_2 \Delta\sigma(\text{bull})_j + \beta_3 \#\text{users}_j + \beta_4 \Delta\#\text{users}_j + \sum Controls_j$$

Quintile	(1)	(2)	(3)	(4)	(5)
		Panel A: By Aggregate Event Volume			
average AEV	5.549	17.470	37.441	74.598	354.135
constant	0.360***	0.637***	0.299**	0.696***	1.021***
	(3.53)	(3.99)	(2.45)	(3.11)	(6.79)
σ(bull)	0.021***	0.022***	0.022***	0.023***	0.015***
	(7.94)	(7.14)	(5.50)	(7.54)	(8.69)
Δσ(bull)	0.022***	0.020***	0.024***	0.026***	0.017***
	(8.48)	(9.67)	(5.95)	(9.20)	(8.60)
#users	0.037***	0.038***	0.043***	0.042***	0.030***
	(8.40)	(6.70)	(5.29)	(8.08)	(9.05)
Δ#users	0.014***	0.011***	0.011***	0.018***	0.012***
	(5.21)	(2.98)	(4.06)	(6.18)	(7.93)
N	3,796	3,795	3,795	3,795	3,795
Adj. R^2	0.195	0.149	0.170	0.173	0.196
		Panel B: Analyst Following			
average #Analysts	2.6227	6.09087	10.1647	15.9537	25.6439
constant	0.245*	0.643*	0.530***	0.547***	0.897***
	(1.74)	(1.94)	(3.33)	(3.31)	(6.84)
σ(bull)	0.029***	0.027***	0.025***	0.015***	0.012***
	(6.93)	(5.98)	(8.47)	(7.75)	(10.18)
Δσ(bull)	0.029***	0.029***	0.024***	0.017***	0.015***
	(7.06)	(6.85)	(9.44)	(9.58)	(11.15)
#users	0.055***	0.052***	0.045***	0.026***	0.023***
	(5.76)	(5.54)	(8.34)	(8.84)	(12.09)
Δ#users	0.015***	0.010***	0.014***	0.011***	0.011***
	(3.97)	(3.65)	(4.46)	(7.15)	(8.67)
N	4,238	4,237	4,237	4,237	4,238
Adj. R^2	0.130	0.188	0.183	0.199	0.225
		Panel C: Market Capitalization			
average log(size)	19.052	20.715	21.714	22.724	24.340
constant	0.303	0.653***	0.578***	0.618***	0.612***
	(1.40)	(3.26)	(4.58)	(6.18)	(8.83)
σ(bull)	0.034***	0.028***	0.019***	0.013***	0.005***
	(8.76)	(8.24)	(12.85)	(10.75)	(10.28)
Δσ(bull)	0.037***	0.027***	0.019***	0.015***	0.006***
	(9.71)	(8.82)	(14.90)	(12.76)	(10.88)
#users	0.069***	0.050***	0.030***	0.024***	0.010***
	(7.61)	(7.26)	(12.37)	(11.97)	(9.27)
Δ#users	0.016***	0.016***	0.012***	0.011***	0.005***
	(3.81)	(5.31)	(9.13)	(8.60)	(9.39)
N	4,238	4,237	4,237	4,237	4,238
Adj. R2	0.159	0.170	0.254	0.221	0.184

Notes: $turn_{t-1,t+1}$ is the log share turnover in the 3 trading days around an earnings announcement; $adj.turn_{t-1,t+1}$ is the turnover in the 3 trading days around an earnings announcement less the median 3 day moving average turnover in 249 preceding trading days; σ(bull) is the dispersion of user sentiment in the given period before an earnings announcement; Δσ(bull) is the change dispersion of user sentiment in the given period around an earnings announcement; #users is a count of the number of individuals posting during this time period. Control variables are not reported for parsimony, all specifications include year fixed effects and firm fixed effects, we also include the log of market turnover, log of the absolute value of the prior firm return, log of market capitalization, and the log of the absolute value of the standardized unexpected earnings as controls. Intraday measures are split based on the earnings announcement time and cover a 12-hour window pre and 12-hour window post the time of the announcement, for announcements after 4 PM EST, measures using daily data are centered on the next trading day. We report *p*-values in parentheses based on standard errors that are clustered by firm. *p<0.1 **p<0.05 ***p<0.01.

Table 6
Social media sentiment and tone dispersion with analyst EPS control variables

$$vol_j = \alpha + \beta_1 \sigma(\text{bull})_j + \beta_2 \Delta\sigma(\text{bull})_j + \beta_3 \#\text{users}_j + \beta_4 \Delta\#\text{users}_j + \beta_5 \sigma(\text{analyst})_j + \beta_6 \Delta\sigma(\text{analyst})_j + \beta_7 \text{jumbling}_j + \sum Controls_j$$

	$turn_{t-1,t+1}$			$adj.turn_{t-1,t+1}$	
	(1)	(2)	(3)	(4)	(5)
constant	1.677***	5.183***	7.886***	0.266***	0.657***
	(5.08)	(12.16)	(22.00)	(16.03)	(17.69)
σ(bull)			0.186***		0.011***
			(16.37)		(9.73)
Δσ(bull)			0.214***		0.014***
			(20.85)		(14.26)
#users			0.283***		0.021***
			(23.10)		(9.20)
Δ#users			0.104***		0.010***
			(9.50)		(8.77)
σ(analyst)	0.067***	0.095***	0.056***	0.001	-0.001
	(4.84)	(5.50)	(4.69)	(0.68)	(-0.67)
Δσ(analyst)	0.044**	0.004	-0.019	0.001	-0.003
	(2.48)	(0.07)	(-0.42)	(1.00)	(-0.86)
jumbling	0.043**	0.009	0.034*	0.000	-0.001
	(2.25)	(0.33)	(1.71)	(0.12)	(-0.47)
log(market turn)	0.461***	0.510***	0.226***	0.025***	0.008
	(5.72)	(5.01)	(2.93)	(6.33)	(1.23)
log(|return|)	0.227***	0.213***	0.119***	0.014***	0.010***
	(25.73)	(19.28)	(16.51)	(18.62)	(13.29)
log(size)	-0.090***	-0.207***	-0.384***	-0.005***	-0.023***
	(-7.32)	(-13.88)	(-29.49)	(-8.27)	(-15.51)
log(|sue|)	0.007	-0.017	-0.031***	0.000	-0.003***
	(0.78)	(-1.26)	(-3.28)	(0.30)	(-2.70)
Year f.e.	Yes	Yes	Yes	Yes	Yes
N	11,555	4,813	4,813	11,555	4,813
Adj. R^2	0.234	0.427	0.638	0.149	0.413

Notes: $turn_{t-1,t+1}$ is the log share turnover in the 3 trading days around an earnings announcement; $adj.turn_{t-1,t+1}$ is the turnover in the 3 trading days around an earnings announcement less the median 3 day moving average turnover in 249 preceding trading days; σ(bull) is the dispersion of user sentiment in the given period before an earnings announcement; Δσ(bull) is the change dispersion of user sentiment in the given period around an earnings announcement; #users is a count of the number of individuals posting during this time period; σ(analyst) is the standard deviation of analysts EPS forecast prior to the earnings announcement; Δσ(analyst) is the standard deviation of analysts EPS forecast following the earnings announcement; jumbling is the natural log of the correlation between analyst forecasts in the periods before and after the earnings announcement; log(market turn) is the natural log of the volume in the primary exchange in which the stock for the given company is traded; log(|return|) is the absolute value of the return in the 3 trading days around an announcement; log(size) is the natural log of the market capitalization of the company; log(|sue|) is the natural log of the difference between quarterly consensus EPS forecasts and the actual EPS, using the I/B/E/S Summary database. Intraday measures are split based on the earnings announcement time and cover a 12-hour window pre and 12-hour window post the time of the announcement, for announcements after 4 PM EST, measures using daily data are centered on the next trading day. We report *p*-values in parentheses based on standard errors that are clustered by firm. *p<0.1 **p<0.05 ***p<0.01.

Table 7
Social media disagreement and user influence

$$vol_j = \alpha + \beta_1 \text{pre top 1000 user interest}_j + \beta_2 \text{pre top 1000 disagreement}_j + \beta_3 \text{post top 1000 user interest}_j + \beta_4 \text{post top 1000 disagreement}_j + \beta_5 \sigma(\text{bull})_j + \beta_6 \Delta\sigma(\text{bull})_j + \beta_7 \#\text{users}_j + \beta_8 \Delta\#\text{users}_j + \sum Controls_j$$

	$turn_{t-1,t+1}$		$adj.turn_{t-1,t+1}$	
	(1) Eigenvector centrality	(2) Number of followers	(3) Eigenvector centrality	(4) Number of followers
constant	0.705***	0.625**	0.387***	0.365***
	(2.87)	(2.56)	(15.54)	(15.10)
pre top 1000 user interest	0.128***	0.170***	0.008***	0.018***
	(8.23)	(4.80)	(5.93)	(5.57)
pre top 1000 user disagreement	0.182***	0.149***	0.018***	0.013***
	(10.71)	(10.45)	(6.73)	(7.27)
post top 1000 user interest	0.044	0.044	-0.004**	0.004
	(1.55)	(0.89)	(-2.08)	(0.74)
post top 1000 user disagreement	0.342***	0.337***	0.023***	0.020***
	(21.58)	(21.17)	(10.82)	(9.31)
σ(bull)	1.515***	1.821***	0.055***	0.081***
	(18.91)	(20.24)	(5.25)	(6.38)
Δσ(bull)	1.078***	1.230***	0.046***	0.060***
	(17.61)	(18.78)	(5.98)	(6.76)
#users	0.001	0.002	0.000	0.000
	(1.39)	(1.51)	(1.18)	(1.25)
Δ#users	0.003***	0.004***	0.001***	0.001***
	(7.62)	(7.93)	(4.19)	(4.38)
log(market turn)	0.241***	0.242***	0.002	0.002
	(4.53)	(4.67)	(0.35)	(0.32)
log(|return|)	0.200***	0.202***	0.015***	0.015***
	(37.02)	(37.23)	(14.91)	(14.87)
log(size)	-0.108***	-0.108***	-0.014***	-0.013***
	(-11.44)	(-11.34)	(-10.63)	(-10.47)
log(|sue|)	-0.012**	-0.008	-0.002**	-0.002**
	(-2.03)	(-1.30)	(-2.55)	(-2.11)
Year f.e.	Yes	Yes	Yes	Yes
Observations	21,187	21,187	21,187	21,187
Adj. R^2	0.363	0.356	0.178	0.174

Notes: $turn_{t-1,t+1}$ is the log share turnover in the 3 trading days around an earnings announcement; $adj.turn_{t-1,t+1}$ is the turnover in the 3 trading days around an earnings announcement less the median 3 day moving average turnover in 249 preceding trading days; σ(bull) is the dispersion of user sentiment in the given period before an earnings announcement; Δσ(bull) is the change dispersion of user sentiment in the given period around an earnings announcement; #users is a count of the number of individuals posting during this time period; Top 1000 User interest is an indictor variable set to one when a post from one of the top 1,000 most influential users occurs; Top 1000 User Disagreement is an indicator set to one when at least two of the top 1,000 most influential users posts disagree based on our bullishness measure. log(market turn) is the natural log of the volume in the primary exchange in which the stock for the given company is traded; log(|return|) is the absolute value of the return in the 3 trading days around an announcement; log(size) is the natural log of the market capitalization of the company; log(|sue|) is the natural log of the difference between quarterly consensus EPS forecasts and the actual EPS, using the I/B/E/S Summary database. Intraday measures are split based on the earnings announcement time and cover a 12-hour window pre and 12-hour window post the time of the announcement, for announcements after 4 PM EST, measures using daily data are centered on the next trading day. We report *p*-values in parentheses based on standard errors that are clustered by firm. *p<0.1 **p<0.05 ***p<0.01.

Table 8
Social media disagreement and user heterogeneity

$$vol_j = \alpha + \beta_1 \sigma(\text{bull})_j + \beta_2 \Delta\sigma(\text{bull})_j + \beta_3 \#\text{users}_j + \beta_4 \Delta\#\text{users}_j + \beta_5 \text{entropy}_j + \beta_6 \Delta\text{entropy}_j + \beta_7 \text{entropy}_j * \#\text{users}_j + \beta_8 \Delta\text{entropy}_j * \Delta\#\text{users}_j + \sum Controls_j$$

$$adj.turn_{t-1,t+1}$$

	(1)	(2)	(3)	(4)	(5)
Constant	-0.009**	0.090	0.215***	0.274***	0.506***
	(-2.23)	(1.09)	(2.76)	(3.64)	(6.24)
entropy	0.012***	0.016***	0.010***	0.014***	-0.004***
	(13.78)	(15.38)	(12.68)	(14.25)	(-2.86)
Δentropy	0.008***	0.010***	0.005***	0.006***	-0.005***
	(13.25)	(14.93)	(6.67)	(7.79)	(-5.18)
#users pre			0.018***	-0.020***	0.023***
			(9.07)	(-4.74)	(5.02)
Δ#users			0.008***	0.006***	0.014***
			(5.62)	(4.11)	(8.65)
entropy x #users				0.009***	0.005***
				(6.60)	(3.50)
Δentropy x Δ#users				0.004***	0.003***
				(7.23)	(6.13)
σ(bull)					0.021***
					(14.29)
Δσ(bull)					0.025***
					(16.90)
log(market turn)		0.023***	0.024***	0.027***	0.035***
		(4.03)	(4.26)	(4.76)	(6.07)
log(|$return_{t-1t+1}$|)		0.013***	0.013***	0.012***	0.010***
		(19.39)	(19.41)	(19.82)	(17.77)
log(size)		0.003	-0.001	-0.004	-0.009**
		(0.75)	(-0.28)	(-1.14)	(-2.43)
log(|sue|)		0.001***	0.001**	0.001	-0.000
		(2.63)	(1.99)	(1.29)	(-0.17)
Year f.e.	Yes	Yes	Yes	Yes	Yes
Firm f.e.	Yes	Yes	Yes	Yes	Yes
Observations	21,187	21,187	21,187	21,187	21,187
Adj. R^2	0.017	0.072	0.097	0.122	0.159
Within R^2	0.017	0.073	0.098	0.123	0.160
Between R^2	0.029	0.078	0.090	0.108	0.116

Notes: $turn_{t-1,t+1}$ is the log share turnover in the 3 trading days around an earnings announcement; $adj.turn_{t-1,t+1}$ is the turnover in the 3 trading days around an earnings announcement less the median 3 day moving average turnover in 249 preceding trading days; σ(bull) is the dispersion of user sentiment in the given period before an earnings announcement; Δσ(bull) is the change dispersion of user sentiment in the given period around an earnings announcement; #users is a count of the number of individuals posting during this time period; Entropy is based on the combination of three Shannon-Entropy calculations based on users' self-described backgrounds for trading strategy, horizon and experience; Top 1000 User interest is an indictor variable set to one when a post from one of the top 1,000 most influential users occurs; Top

1000 User Disagreement is an indicator set to one when at least two of the top 1,000 most influential users posts disagree based on our bullishness measure. log(market turn) is the natural log of the volume in the primary exchange in which the stock for the given company is traded; log(|return|) is the absolute value of the return in the 3 trading days around an announcement; log(size) is the natural log of the market capitalization of the company; log(|sue|) is the natural log of the difference between quarterly consensus EPS forecasts and the actual EPS, using the I/B/E/S Summary database. Intraday measures are split based on the earnings announcement time and cover a 12-hour window pre and 12-hour window post the time of the announcement, for announcements after 4 PM EST, measures using daily data are centered on the next trading day. We report *p*-values in parentheses based on standard errors that are clustered by firm. *p<0.1 **p<0.05 ***p<0.01.

Table 9
Trading volume and investor disagreement around news events

$$adj.turn_{j,t-1,t+1} = \alpha + \beta_1\sigma(\text{bull})_j + \beta_2\Delta\sigma(\text{bull})_j + \sum Controls_j$$

	(1) All news	(2) Scheduled	(3) Unscheduled	(4) Business	(5) Society
constant	0.356***	0.264***	0.387***	0.355***	0.402***
	(17.79)	(16.38)	(17.04)	(17.69)	(7.81)
σ(bull)	0.012***	0.011***	0.012***	0.012***	0.008***
	(10.69)	(10.33)	(10.36)	(10.86)	(4.09)
Δσ(bull)	0.009***	0.008***	0.009***	0.009***	0.007***
	(9.71)	(10.22)	(9.44)	(9.86)	(4.07)
#users	-0.003***	-0.001***	-0.004***	-0.003***	-0.001
	(-8.96)	(-2.98)	(-8.33)	(-9.15)	(-0.76)
Δ#users	0.018***	0.015***	0.019***	0.018***	0.009***
	(6.55)	(5.96)	(6.68)	(6.92)	(2.70)
log(market turn)	0.009**	0.007*	0.008**	0.009**	0.014
	(2.38)	(1.86)	(2.22)	(2.37)	(1.37)
log(|return|)	0.012***	0.010***	0.012***	0.012***	0.012***
	(19.77)	(16.77)	(19.05)	(19.74)	(6.24)
log(size)	-0.010***	-0.007***	-0.011***	-0.010***	-0.012***
	(-9.37)	(-10.23)	(-9.30)	(-9.34)	(-7.88)
Year f.e.	Yes	Yes	Yes	Yes	Yes
N	156,981	34,346	122,635	155,000	1,969
Adj. R^2	0.282	0.232	0.300	0.283	0.270

Notes: $turn_{t-1,t+1}$ is the log share turnover in the 3 trading days around an earnings announcement; $adj.turn_{t-1,t+1}$ is the turnover in the 3 trading days around an earnings announcement less the median 3 day moving average turnover in 249 preceding trading days; σ(bull) is the dispersion of user sentiment in the given period before an earnings announcement; Δσ(bull) is the change dispersion of user sentiment in the given period around an earnings announcement; #users is a count of the number of individuals posting during this time period. log(market turn) is the natural log of the volume in the primary exchange in which the stock for the given company is traded; log(|return|) is the absolute value of the return in the 3 trading days around an announcement; log(size) is the natural log of the market capitalization of the company; log(|sue|) is the natural log of the difference between quarterly consensus EPS forecasts and the actual EPS, using the I/B/E/S Summary database. Intraday measures are split based on the earnings announcement time and cover a 12-hour window pre and 12-hour window post the time of the announcement, for announcements after 4 PM EST, measures using daily data are centered on the next trading day. We report *p*-values in parentheses based on standard errors that are clustered by firm. *p<0.1 **p<0.05 ***p<0.01.

Table 10
Social media sentiment and tone dispersion by various time horizons

$$adj.turn_j = \alpha + \beta_1\sigma(\text{bull})_j + \beta_2\Delta\sigma(\text{bull})_j + \beta_3\sigma(\text{analyst})_j + \beta_4\Delta\sigma(\text{analyst})_j + \beta_5\text{jumbling}_j + \sum Controls_j$$

	$turn_{t-1,t+1}$				
	(1)	(2)	(3)	(4)	(5)
	30 minutes	1 hour	3 hours	6 hours	12 hours
constant	0.796***	0.645***	0.561***	0.529***	0.481***
	(8.41)	(9.79)	(12.15)	(14.06)	(16.97)
σ(bull)	0.021***	0.019***	0.016***	0.014***	0.015***
	(6.58)	(8.01)	(9.82)	(10.57)	(13.20)
Δσ(bull)	0.021***	0.022***	0.021***	0.020***	0.018***
	(5.86)	(8.55)	(12.33)	(14.23)	(17.32)
#users	0.028***	0.024***	0.022***	0.021***	0.022***
	(3.89)	(4.33)	(5.43)	(6.45)	(8.66)
Δ#users	0.016***	0.018***	0.022***	0.024***	0.021***
	(4.83)	(6.85)	(9.58)	(10.86)	(12.77)
log(market turn)	0.000	-0.001	-0.004	-0.002	0.001
	(0.01)	(-0.04)	(-0.41)	(-0.29)	(0.16)
log(|return|)	0.026***	0.021***	0.018***	0.015***	0.014***
	(5.72)	(7.26)	(10.05)	(11.56)	(15.01)
log(size)	-0.026***	-0.021***	-0.019***	-0.018***	-0.016***
	(-6.34)	(-7.07)	(-8.34)	(-9.35)	(-11.26)
log(|sue|)	0.000	-0.001	-0.002	-0.001	-0.002***
	(0.03)	(-0.42)	(-1.56)	(-1.55)	(-3.09)
Year f.e.	Yes	Yes	Yes	Yes	Yes
N	3,773	5,863	10,184	13,723	21,187
Adj. R^2	0.169	0.169	0.185	0.195	0.194

Notes: $turn_{t-1,t+1}$ is the log share turnover in the 3 trading days around an earnings announcement; $adj.turn_{t-1,t+1}$ is the turnover in the 3 trading days around an earnings announcement less the median 3 day moving average turnover in 249 preceding trading days; σ(bull) is the dispersion of user sentiment in the given period before an earnings announcement; Δσ(bull) is the change dispersion of user sentiment in the given period around an earnings announcement; #users is a count of the number of individuals posting during this time period. log(market turn) is the natural log of the volume in the primary exchange in which the stock for the given company is traded; log(|return|) is the absolute value of the return in the 3 trading days around an announcement; log(size) is the natural log of the market capitalization of the company; log(|sue|) is the natural log of the difference between quarterly consensus EPS forecasts and the actual EPS, using the I/B/E/S Summary database. Intraday measures are split based on the earnings announcement time and cover a window of the specified time in each column pre and equal window post the time of the announcement, for announcements after 4 PM EST, measures using daily data are centered on the next trading day. We report *p*-values in parentheses based on standard errors that are clustered by firm. *p<0.1 **p<0.05 ***p<0.01.

Appendix

Extract from StockTwits

Figure A.1
Example discussion on StockTwits for Fitbit with bullish and bearish tags ($FIT)

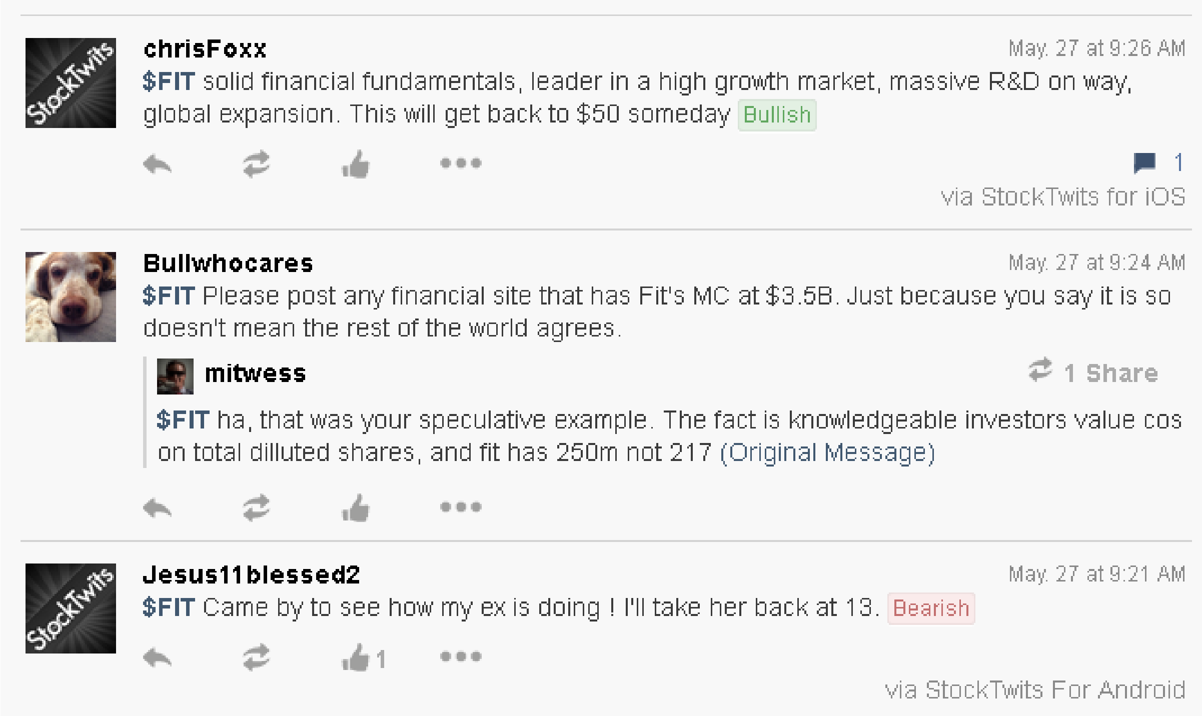

Notes: The green "Bullish" and red "Bearish" symbols are the source of disagreement for the bull-bear disagreement measure. The text from the posts are used to calculate textual analysis based measures of disagreement and are based on the positive and negative words in each post.