THE ACCOUNTING REVIEW
Vol. 98, No. 1
January 2023
pp. 109–137

American Accounting Association
DOI: 10.2308/TAR-2018-0259

Investor Disagreement, Disclosure Processing Costs, and Trading Volume Evidence from Social Media

Adam Booker
University of Denver

Asher Curtis
University of Washington

Vernon J. Richardson
University of Arkansas
Baruch College–CUNY

ABSTRACT: We use posts on the investor-focused StockTwits social media network to generate new insights regarding investor disagreement, disclosure processing costs, and trading volume around earnings announcements. Using social media-based measures of disagreement, we find that both preannouncement disagreement and increases in disagreement around an earnings announcement are positively associated with trading volume. Drawing upon the disclosure processing costs literature, we provide evidence that the effects of disagreement increase when disclosure processing costs are lower. Our social media measures of disagreement remain significant after including traditional analyst earnings estimate measures of disagreement in the model. Our study provides new evidence on the importance of disclosure processing costs and is consistent with lower disclosure processing costs amplifying both the resolution of preannouncement disagreement and new disagreement about earnings information.

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

Keywords: disagreement; trading volume; social media; disclosure processing costs.

I. INTRODUCTION

The linkage between trading volume and earnings announcements provides the most direct evidence that information in earnings disclosures changes individual investors' expectations, leading them to trade (see Bamber, Barron, and Stevens 2011, for a review). In turn, increased trade has been understood to result from investor disagreement, either from differing priors or conflicting interpretations of information perceived to affect discount rates

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

Adam Booker, University of Denver, Daniels College of Business, Accounting, Denver, CO, USA; Asher Curtis, University of Washington, Foster School of Business, Accounting, Seattle, WA, USA; Vernon J. Richardson, University of Arkansas, Sam M. Walton College of Business, William Dillard Department of Accounting, Fayetteville, AR and Baruch College–CUNY, Zicklin School of Business, Stan Ross Department of Accountancy, New York, NY, USA.

Editor's note: Accepted by Mark T. Bradshaw, under the Senior Editorship of Mary E. Barth.

Submitted: May 2018
Accepted: May 2022
Early Access: May 2022

and cash flows.[1] Importantly, the literature suggests that earnings announcements resolve differences in investor priors (Ahmed, Schneible, and Stevens 2003; Barron, Schneible, and Stevens 2018) and thus that accounting disclosures help "level the playing field" between investors with differential access to information before an earnings announcement.

Recent theories have increased our understanding of individual investor behavior by suggesting that investors face processing costs as new information arrives (Hirshleifer and Teoh 2003; Veldkamp 2011). Within the accounting literature, Blankespoor, deHaan, and Marinovic (2020) define processing costs as the costs of awareness, acquisition, and integration of new information. No model of which we are aware, however, has examined whether disclosure processing costs interact with investor disagreement. We expect processing costs to impede the resolution of differences in priors and lower the extent to which earnings announcements will "level the playing field" between differentially informed investors. Thus, understanding the potential interaction between disclosure processing costs and disagreement is important as it provides insight into the informativeness of earnings to diverse groups of investors.

Recently, Cookson and Niessner (2020, 215) provided preliminary evidence that the daily level of disagreement on social media has a higher positive association with trading volume on days when a firm announces earnings relative to the four weeks around the earnings announcement. Although they find evidence of an increased association, they find that the increased trading volume around earnings announcements is not explained by disagreement and conclude that more work is needed to explain trading volume around earnings announcements. Using intraday StockTwits data, we extend Cookson and Niessner (2020) by investigating how processing costs affect the link between trading volume and preannouncement disagreement and changes in disagreement around earnings announcements. We find strong evidence consistent with lower processing costs amplifying disagreement.

Theory suggests that trading volume around earnings announcements reflects both the resolution of preannouncement disagreement and new disagreement about the interpretation of earnings news. We measure disagreement between individual StockTwits users. Users on StockTwits regularly express their opinions about stocks in short messages, or posts, with some users explicitly stating that they are positive (flagging posts with a "bullish" icon) or negative (flagging posts with a "bearish" icon). These posts occur in real-time, allowing direct measures of disagreement and changes in disagreement around earnings announcements. We measure preannouncement disagreement by calculating the standard deviation of the sentiment of users' posts in the 24-hour window before the announcement and the change in disagreement as the difference between the standard deviations of sentiment in the 24-hour windows before and after the earnings announcement (with positive changes interpreted as increases in disagreement).[2] Our results are consistent with prior studies and suggest that trading volume around earnings announcements increases with both preannouncement disagreement and increases in disagreement around the announcement (Bamber, Barron, and Stober 1997; Giannini, Irvine, and Shu 2019).[3]

We expect that disclosure processing costs impede the ability of investors to process earnings disclosures (Hirshleifer and Teoh 2003; Blankespoor, deHaan, and Zhu 2018). Based on the causal evidence in Blankespoor et al. (2018), we expect processing costs to have a significant main effect on trading volume around earnings announcements. We also predict that disclosure processing costs interact with disagreement, as these costs impede both the resolution of preannouncement disagreement and new disagreement about the interpretation of earnings news.

We consider two measures of disclosure processing costs for StockTwits users, investor attention and investor heterogeneity. We measure attention, an inverse proxy for disclosure processing costs based on limited attention, as the number of active users posting on StockTwits around the earnings announcement. Although prior work in social media settings often uses the number of posts to measure attention, the theory assumes that attention differs between individuals.[4] In addition, as we use intraday StockTwits data, we measure preannouncement attention and changes in attention around earnings announcements. As expected, we find evidence that both preannouncement attention and the change in attention around the earnings announcement are associated with increased trading volume. We also find that the interactive effects between attention and disagreement are consistently positive and significant around earnings announcements, whereas the main effects of

[1] Although the assumptions about investor behavior differ, both rational theories and behavioral theories support these expectations. Bamber et al. (2011) provide a detailed analysis of analytical models using adaptive expectations, noisy rational expectations, and differences of opinion. The authors highlight that these different classes of models commonly predict trading volume increases due to differences in prior information and differences in interpretation of the news.

[2] Due to the nature of language used in short social media posts, or tweets, relative to more formal writing such as seen in the 10-K, we provide multiple estimates of disagreement based on different machine-learning calculations of sentiment. We generally find that models with more complexity do better in our validation checks consistent with these models producing better out-of-sample estimates of the sentiment of each post. When we aggregate the individual post sentiment measures to disagreement, however, we find that the more complex models only perform marginally better in our regression analysis (see Section IV).

[3] Other studies have examined disagreement around different news events, including the disclosure of management forecasts (Cho and Kwon 2014), macroeconomic news events (Bollerslev, Li, and Xue 2018), or between different individuals such as mortgage brokers (Carlin, Longstaff, and Matoba 2014).

[4] Antweiler and Frank (2004) use the logged number of posts, Cookson and Niessner (2020) use the unlogged number of posts and Curtis et al. (2016) use a proprietary measure of attention from Market IQ.

disagreement vary and are often subsumed in models that account for the interactive effects. We corroborate our findings that lower disclosure processing costs amplify disagreement with a placebo test using a random sample of dates for which earnings are not announced and find that the interactive effects between attention and disagreement are inconsistent in these models.

We next examine whether disagreement interacts with investor heterogeneity, another proxy for disclosure processing costs based on rational inattention to earnings information (e.g., Veldkamp 2011). The content of StockTwits posts varies by individual, with some users posting fundamental information (such as earnings components and expectations) and other users posting information relating to technical trading strategies (such as short-term moving averages). Intuitively, posts from a population of investors with more diverse backgrounds would lead to higher disagreement since these posts include a greater diversity of prior beliefs linked to different trading strategies. Like Cookson and Niessner (2020), we measure investor heterogeneity using the self-reported user characteristics of users posting around the earnings event.[5] Our investor heterogeneity measure increases when the users posting at the earnings announcement have different backgrounds.[6] We find a positive association with trading volume around earnings announcements for interactive effects between heterogeneity and disagreement, consistent with rational inattention to earnings.

We then conduct further analysis, including robustness to key empirical choices. First, we compare our measure of social media disagreement with analyst-based disagreement measures from Bamber et al. (1997) and find that social media measures of disagreement are distinct from analyst-based measures of disagreement. Second, we provide evidence of important cross-sectional differences driven by variations in earnings news, existing information, and attention to the earnings announcement. We find that the role of disagreement in explaining trading volume around earnings announcements is larger in the presence of extreme earnings surprises, especially those containing positive news. Increased access to information before the earnings announcement (e.g., larger firms and firms with more media visibility) and limited attention to the announcement (e.g., busy earnings announcement days) reduce the role of disagreement in explaining trading volume around earnings announcements. Third, we analyze nonearnings announcement days as a placebo test and do not find that attention and disagreement have an interactive effect on trading volume, consistent with these days having lower processing costs than earnings announcements.

Our study contributes to the growing literature about disclosure processing costs, which is currently underexplored (Blankespoor et al. 2020). Our results contrast with the evidence of Cookson and Niessner (2020), who find that disagreement and attention are distinct for a sample of regular trading days. Instead, our results suggest that disagreement and attention are related and have an interactive effect when explaining trading volume around earnings announcements. One possible reason for the difference between the two studies, following the limited attention literature, is that the effects of attention are more likely to be observed around earnings announcements, both because those trading days include disclosures that require processing and because they tend to cluster together with the earnings announcements of other firms. We also find that cross-sectional differences in attention interact with cross-sectional differences in preannouncement disagreement and changes in disagreement around the announcement. These results are consistent with disclosure processing costs mitigating the degree to which investors resolve prior disagreements around information releases and influencing new disagreement about information released at the earnings announcement.

Our study also contributes to the literature by highlighting that earnings announcements are informative to a diverse set of investors, even in the information-rich setting of social media. Our results provide additional insight into the results of Barron et al. (2018), who find that there has been an increase in the trading volume reactions to earnings announcements that are coincident with the trend over time of greater investor diversity. They suggest that the increased diversity stems from increased participation from both increasingly diverse institutional investors as well as smaller traders over time. Our evidence contributes additional insights into the role of disagreement from a heterogeneous set of likely smaller, noninstitutional investors interacting on social media. As highlighted by Bamber et al. (2011), analysts represent relatively sophisticated investors that underrepresent the heterogeneity in the broader marketplace. Our setting provides new insights by incorporating measures of disagreement and disclosure processing costs that reflect a broad population of investors interacting on social media.

II. INSTITUTIONAL BACKGROUND AND HYPOTHESIS DEVELOPMENT

Institutional Background

StockTwits is an investor-focused social media network analogous to "Twitter for investors." StockTwits, like Twitter, restricts the length of the post an individual can make to a small number of characters (originally to a maximum

[5] Whereas Cookson and Niessner (2020) examine differences in trading strategy only, we use all the self-reported user characteristics, which include horizon, trading approach, investing horizon, and experience.

[6] We also consider two measures of investor heterogeneity based on the visibility of StockTwits users posting around earnings announcements (see the "Examination of Additional Predictions" section). The social finance literature suggests visible users could be more influential, resulting in either distraction or increased attention depending on the content of their posts (Shiller and Pound 1989; Hirshleifer 2015).

of 140 characters). StockTwits is a popular website for investors with two-million registered members, four-million monthly messages (i.e., posts), and an external audience of three-million monthly viewers.[7] The reach of the information on StockTwits is extended by the ability to obtain a license to ingest the messages in real-time via the application programming interface, or API, calls. Third parties who provide StockTwits data to their clients include several online trading platforms, including Bloomberg, Fidelity, and Charles Schwab. Anecdotally, larger sophisticated market participants also access this data via license and leverage the data in various trading strategies.

The founders of StockTwits invented cashtags, a dollar sign appended to the ticker symbol, to identify stocks (for example, $AAPL for Apple Inc.). This allows StockTwits users to reference specific stocks quickly and provides an efficient method for aggregating ideas about a given stock, with each cashtag linking to a webpage about that individual stock. As of the writing of this paper, StockTwits' content aggregations are generally a top result in internet search engines when searching by cashtag. Posts about multiple stocks are referenced using multiple cashtags (e.g., $AAPL $MSFT) and appear on all the individual StockTwits webpages of the stocks referenced (in our example, on both Apple, Inc. and Microsoft's StockTwits webpages).

StockTwits is a social network for investors and focuses on trading-related opinions, in contrast to other social media websites, like Twitter and Facebook, which cater to sharing broader discussions and opinions. Only around 11 percent of the users on StockTwits share their ideas simultaneously on Twitter and StockTwits, despite Twitter using cashtags in place of hashtags for stock-related posts since 2012. Similarly, StockTwits also allows linking to Facebook, but users seldom link their StockTwits posts to Facebook. Thus, unlike other social media platforms, the posts shared on StockTwits represent a sample of posts focused on individual investors' ideas and opinions about trading stocks. In the timeframe of data used in this study, StockTwits used the tagline, "A place for investors to grow and learn." We conjecture that the popularity of posting opinions on social media by smaller investors stems from both a sense of belonging to a social network focused on investing and the ability to potentially highlight their own ability to predict stock market outcomes.

There are two additional aspects of StockTwits that make this setting ideal for examining individual investors' differences in beliefs about stocks. First, StockTwits users can unambiguously indicate their opinion about a stock on a given post by attaching a "bullish" or "bearish" flag that appears as a text icon immediately after their post (see Figure 1 for examples). In addition to allowing for the identification of sentiment for these posts, the existence of these flags following the text of the post makes it possible to use supervised machine learning to infer the sentiment of unflagged posts even without the availability of a StockTwits or social media-specific dictionary. This is important, as the use of language specific to StockTwits, and more generally on social media, differs from traditional language use.

Second, StockTwits caters to a broad group of individuals with differing experiences and trading objectives, including their expected holding period (or horizon) and the information they use to trade (or trading approach). As part of setting up a profile on StockTwits, users can select summary information about their trading objectives, including the selection of Swing Trader, Position Trader, Day Trader, and/or Long Term Investor to appear as their "Horizon," the selection of Technical, Momentum, Fundamental, Growth, Value, or Global Macro to appear as their "Trading Approach," and finally Novice, Intermediate, or Professional to appear as their "Experience." We plot the distribution of answers to Horizon, Trading Approach, and Experience in Figure 2. Although many users choose not to provide an answer, those that do offer a wide number of different answers, consistent with StockTwits users reflecting the broad, heterogenous set of market participants that discuss a broad set of signals and news events on social media.

Review of Theory and Hypothesis

A large analytical literature considers how the disclosure of accounting or other information is expected to increase trading volume. Bamber et al. (2011) summarize this literature and suggest that analytical models commonly predict that trading around news events is due to differential belief revision from differences in (1) investors' beliefs before the news event or (2) investors' interpretation of the news.

Kim and Verrecchia (1991) suggest that dispersion in investors' beliefs before the news event increases trading volume. Specifically, when investors receive noisy information signals before an earnings event, this creates dispersion in beliefs before an information event. When new information arrives that is more precise, investors will trade toward their optimal share exposure based on this new information. These models assume a common interpretation of new information, and in rational models such as Kim and Verrecchia (1991), trading volume is predicted to be proportional to the absolute value of returns around the news event. The evidence in Kandel and Pearson (1995), however, of significant trading around earnings events even when returns are near zero suggests that the revision process around news events is

[7] https://about.stocktwits.com/

FIGURE 1

Example Discussion on StockTwits for Fitbit with Bullish and Bearish Tags ($FIT)

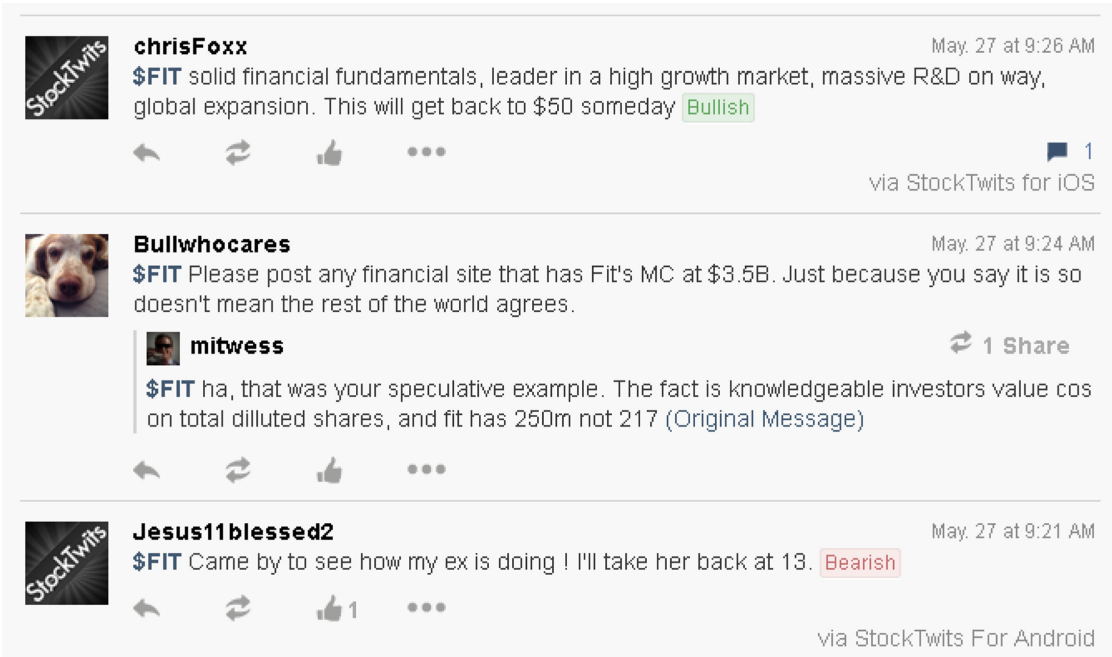

In the first and last posts above, the green "Bullish" and red "Bearish" text icons reflect user-tagged posts. These posts have an explicit opinion tied to them. The post in the middle is not user-flagged and cites another earlier post (shown as indented on this screen). The text of the new unflagged post (i.e., excluding the cited post) is used to calculate sentiment using machine learning models.
(The full-color version is available online.)

not solely due to the dispersion of beliefs prior to the announcement. They model a trading process that includes differential interpretation of the news. Kim and Verrecchia (1997) incorporate differential interpretation in their model of trading volume to incorporate an additional private signal around the news release. In their model, trading volume from the convergence of prior dispersion is predicted to be proportional to the absolute value of returns around the news event, and any excess trading volume is predicted to be due to a differential interpretation of the news event.

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

H1: The prior dispersion of user opinions and the change in the dispersion of user opinions on social media are positively associated with trading volume.

There are several reasons why we may fail to find support for our hypothesis. First, if StockTwits users hold opinions not reflective of the wider market opinions of those investors' trading following an earnings announcement, we may not find support for our hypothesis.[8] Second, we may only find weak evidence of changes in disagreement around the announcement if earnings announcements do not elicit a significant number of additional social media users willing to express their opinions or if all users share very similar opinions in a short time period following the earnings announcement. Third, Cookson and Niessner (2020, 215) provide evidence that the spike in trading volume in the weeks around earnings announcements is largely unexplained by the daily level of disagreement. As we are examining preannouncement disagreement and the change in disagreement around the announcement, it is unclear what the implications of the Cookson and Niessner (2020) result means for our hypothesis. It is possible that preannouncement disagreement may not be resolved by the earnings announcement and possible that the earnings announcement does not elicit a change in disagreement for investors on social media. In either case, we may not find support for our hypothesis. Fourth, we

[8] As we are examining trading volume and not returns to the announcements, our hypothesis is not dependent on StockTwits users being the marginal investor. Part of the effect we document could be due to the observation that many StockTwits users actively trade in relation to their posts. For example, after our sample period, StockTwits users can now post evidence of their trades by linking their Robinhood brokerage account. See https://blog.stocktwits.com/we-teamed-up-with-robinhood-to-bring-true-social-trading-to-the-stocktwits-community. Future research may investigate return effects in a social media setting by potentially leveraging this information when the data are made available for academic use.

FIGURE 2
Postsentiment Distributions by Classification Method

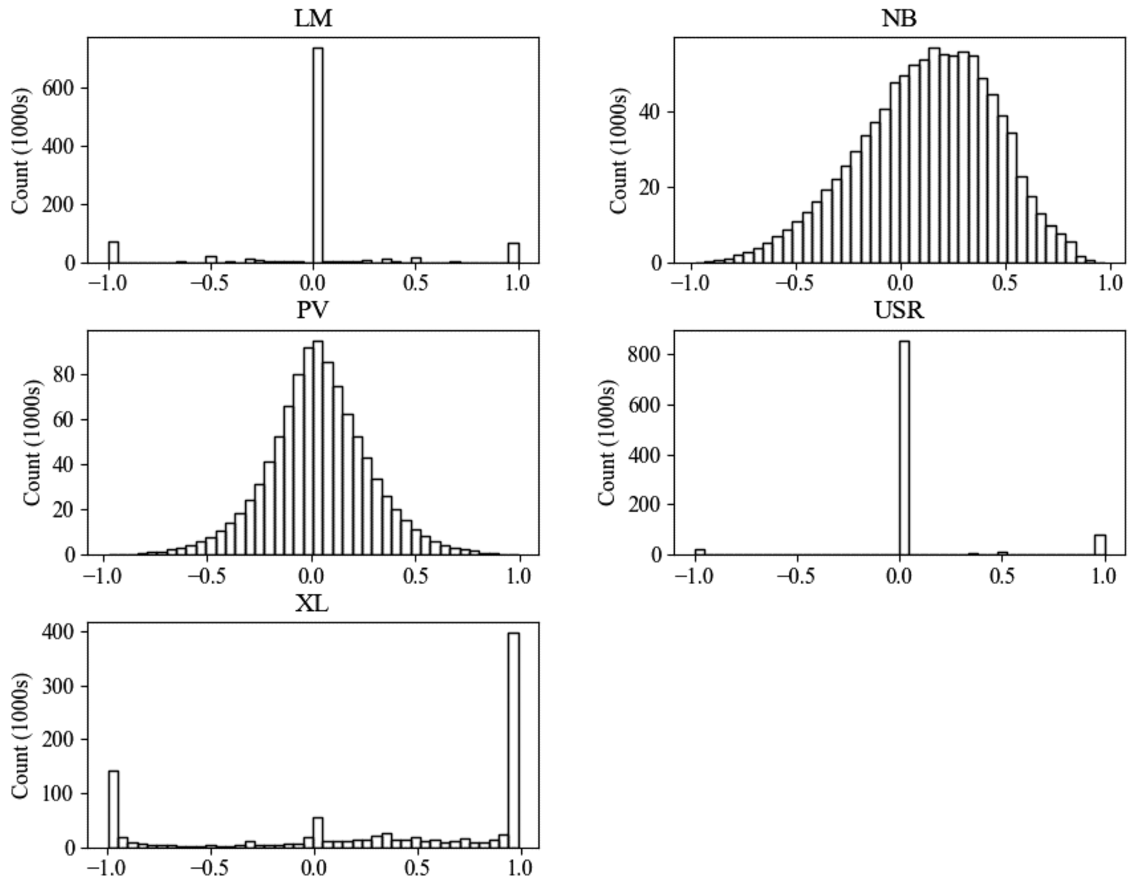

All models are normalized to classify sentiment between −1 and 1, and the bin width for the histograms is set to 0.01. LM, Loughran-McDonald dictionary-based classifier; NB, Naïve Bayes classifier; PV, paragraph vector classifier; USR, self-reported user sentiment classification with unflagged posts set to 0; XL, XLNet classifier.

require textual analysis to calculate the opinions of individual investors, using machine learning to estimate sentiment. Whereas the content of posts has lexical and syntactic elements (titles, important sentences, etc.) that provide insight into the opinions of individuals, sentiment derived from textual analysis of user posts is not a perfect indicator of users' opinions. Like the difficulties with interpreting written text, our measures of disagreement rely on the ability of the textual analysis calculations, and if these models fail to discriminate between positive and negative posts, then we may fail to find evidence to support our hypothesis (Loughran and McDonald 2011; Guay, Samuels, and Taylor 2016).

Additional Empirical Predictions

Technological advances in internet and mobile technologies over the past few decades have made it easier for investors to connect and discuss earnings information. Social media sites like StockTwits provide a network that is open to the public, allowing a diverse group of investors the chance to share their opinions with others. Social media networks that discuss trading ideas and opinions provide an interactive and information-rich environment with real-time updates.[9] These features allow us to make predictions about how user characteristics may interact with disagreement around earnings announcements.

[9] These websites appear to both reflect and generate information about stocks. 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.

We base our predictions on recent theory, which assumes individual investors face processing costs when new information arrives (see Veldkamp 2011 for a review). An important implication of this literature is that it leads to rational inattention—investors choose what information to acquire and integrate into their trading strategies. We follow the taxonomy in Blankespoor et al. (2020) to consider investors facing frictions relating to awareness, acquisition, and integration costs for accounting information. In addition to attention having the main effect on trading volume around news events, we expect that rational inattention will interact with disagreement. Odean (1998) provides a model in which trading volume increases with investor confidence and attention. We expect that processing costs for bulls and bears decrease as attention (the amount of processed information) and investor diversity (the differences in information shared by investors with varied backgrounds) increase.

Our prediction for frictions in interpreting an earnings announcement (a form of inattention) suggests that all else equal, lower levels of attention to an announcement will decrease disagreement-based trading volume. Lower attention could be due to either a lack of awareness of the earnings announcement or a lack of acquisition of the information in the earnings announcement. In contrast, a lack of awareness about an upcoming earnings announcement could trigger an increase in attention to the information in the earnings news, a surprise effect, resulting in increased attention to the stock from additional investors. The arrival of new investors, or increased attention around an earnings announcement, could lead to either an increase or decrease in disagreement about the new information in the earnings report. Thus, we examine whether an increase in user attention strengthens the association between prior dispersion and trading volume and the association between disagreement about new information and trading volume.

III. VARIABLE MEASUREMENT

Measurement of Trading Volume Reactions

We follow prior literature and calculate the median-adjusted trading volume to measure trading volume reactions to earnings announcements. We use this measure as it is intended to control for cross-sectional differences in liquidity trading coincident with the earnings announcement (Bamber et al. 1997, 2011). Specifically, we calculate turnover as trading volume divided by shares outstanding using a three-day window around the earnings announcement:

$$TURN_t = \frac{\sum_{j=-1}^{j=1} VOL_{t+j}}{\sum_{j=-1}^{j=1} SHROUT_{t+j}}, \tag{1}$$

where VOL_t is the trading volume on day t and $SHROUT_t$ is the number of shares outstanding on day t both from the daily CRSP database. We select the date $t = 0$ as the trading day of the news event if the event occurs before 4 PM EST or the next trading day if the announcement is after 4 PM EST on a trading day or occurs on a nontrading day.

We calculate $ADJTURN_t$ as the median-adjusted shares traded in the three days around an earnings announcement, divided by the sum of the number of shares outstanding less the same for a moving three day window in the previous 249 days as follows:

$$ADJTURN_t = TURN_t - median(TURN_{t-259,\,t-2}) \tag{2}$$

Measurement of Independent Variables

For each stock, we measure investor disagreement as the standard deviation of individual user sentiment calculated from that user's posts. In the cases where a bullish icon is attached to the post, we set sentiment to 1 (the maximum), and in the cases where a bearish icon is attached, we set sentiment to -1 (the minimum). In the cases where a bullish or bearish icon is not attached to the post, we calculate the sentiment of each post using textual analysis. Our primary method to calculate sentiment in these instances uses the Naïve Bayes machine learning classifier commonly used in the accounting literature (Li 2010). To supplement this technique, we also validate and test alternative techniques to calculate sentiment from an individual post. Specifically, we provide supplementary analysis of measures of sentiment based on a dictionary approach following Loughran and McDonald (2011) which counts the number of positive and negative words in each post and two supervised machine learning approaches that are based on recent advances in computer science: (1) the paragraph vector model of Le and Mikolov (2014), and (2) a neural networks approach using the XLNet model of Yang et al. (2019).

We measure preannouncement disagreement using the sentiment of users in the 24-hour period prior to the time of the earnings announcement:

$$PDIS_{t\in[-1,0]} = \sqrt{\frac{\Sigma_i \left(Bull_{it} - \overline{Bull_t}\right)^2}{n-1}}, \tag{3}$$

where $Bull_{it}$ is a continuous measure of sentiment that takes values between -1 and 1 for post i in timespan t, based on either a bearish or bullish icon attached to the post, or a textual analysis-based measure of sentiment derived from the text of the post, n is the number of unique users posting during this window. We require at least three individuals to post within the window to calculate this measure.[10]

We measure the change in disagreement around the earnings announcement as the standard deviation of individual user sentiment in the 24-hour period following the earnings announcement using the standard deviation formula in Equation (3) for disagreement in the time $t \in [0,1]$ ($POSTDIS$) less $PDIS$. We measure *post disagree* as in Equation (3), except that we use posts from the 24-hour period following the earnings announcement event. Specifically, we measure the change in disagreement as follows:

$$\Delta DIS_{t\in[-1,1]} = POSTDIS_{t\in[0,1]} - PDIS_{t\in[-1,0]}. \tag{4}$$

To examine our research questions based on disclosure processing costs, we measure investor attention as the ranked number of unique users posting about the stock during the 48 hours around the news event. Our measure differs slightly from that of Antweiler and Frank (2004), who use the log number of posts on message boards as a measure of attention, Curtis et al. and Curtis, Richardson, and Schmardebeck (2016), who use a proprietary measure of attention, "Smart Velocity," provided by MarketIQ, which can be considered as a measure of social media buzz.[11] We use ranked attention to eliminate ambiguity in interpreting the interaction of two change of variable changes around the announcement and to eliminate possible concerns due to nonlinearity when using unlogged measures of attention. Our measure focuses on the number of users rather than the number of posts because the number of posts can be skewed by one excited individual and is more consistent with the theory of disagreement as being between investors. We use a ranked measure of attention to account for ambiguities in change models.[12] We subtract the mean daily attention in the -10 to -30 days relative to the announcement before ranking to account for firm and time differences in base-level attention. Our measure of attention is the percentile rank (from 0 to 1) of:

$$PATTN_{t\in[-1,0]} = \log\left(Attention_{t\in[-1,0]}\right) - \log\left(median\left(Attention_{t\in[-30,-10]}\right)\right),$$

and our measure of change in attention is the percentile rank of:

$$\Delta ATTN_{t\in[-1,1]} = \log\left(Attention_{t\in[0,1]}\right) - \log\left(Attention_{t\in[-1,0]}\right). \tag{5}$$

We measure the heterogeneity of individuals posting during the event window using a measure of entropy that incorporates key information on users' self-disclosed profiles. Our measure differs from Cookson and Niessner's (2020) measure of investor heterogeneity as they focus solely on heterogeneity from trading strategies. We use an entropy-based measure to calculate a composite measure of heterogeneity.[13] Specifically, our entropy measure is increasing in the uniformity of the distribution of the set of users' backgrounds and is calculated as follows:

$$H = \sum\nolimits_{i=0}^{n} p_i * \log_2(p_i), \tag{6}$$

[10] When individuals post multiple times, we take the average sentiment of their posts as this aligns with the construct of disagreement between individuals rather than between posts.

[11] Social media platforms also disclose user engagement in terms of the number of active users rather than the number of posts (for example, https://www.sec.gov/edgar/search/#twtr-ex992_6.htm).

[12] As a robustness analysis, we also considered binary indicators for high attention and find similar but weaker results to our ranked measure. Specifically, we measure high attention (and high change in attention) as either (1) those in the top quintile or (2) those above the median.

[13] This measure is also known as Shannon entropy, or information entropy. For simplicity, we will refer to this as entropy. 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' experience levels could be encoded as novice, intermediate, professional, or undeclared).

where H is the entropy and p_i is the prior probability of a user having the background classification i in the event period. In this study, we use three different categories of StockTwits users' backgrounds 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 self-report information about their background in these categories. For example, to construct the measure of entropy for users' experience levels, 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:

$$H_{total} = H_{trading\ strategy} + H_{experience} + H_{holding\ period}. \tag{7}$$

Measurement of Control Variables

We base our selection of control variables on Barron et al. (2018). As our dependent variable is median-adjusted trading volume, which controls for liquidity trading under the assumption that liquidity trading is constant (Kim and Verrecchia 1997), we follow Bamber et al. (1997) and include market turnover to control for macroeconomic-related liquidity trading shocks unrelated to information in the news event. Theory suggests that variation in the precision of pre-announcement information will result in trade around the release of new information (Kim and Verrecchia 1991). Following Barron et al. (2018), we include the absolute value of returns to control for the magnitude of belief revision that occurs at an earnings announcement. Finally, we include the closing price two days before earnings are announced to control transaction costs that may dampen information-based trading (Ahmed et al. 2003; Barron and Karpoff 2004).[14]

IV. EMPIRICAL ANALYSIS

Data and Sample

We collect daily market data (returns, volume, and prices) from the CRSP (The Center for Research on Securities) daily securities file, earnings data from I/B/E/S (Thomson Reuters Institutional Brokers Estimate System), industry data from Compustat, and social media data from StockTwits.[15] We collect financial information and StockTwits user posts for a sample of public firms with quarterly earnings announcements between July 2010 and December 2015. We began in July 2010 when StockTwits first implemented the ability for users to add a bullish or bearish icon to their posts, which we require to calculate user sentiment. We end on December 31, 2015, as it is the final year that we have data available from StockTwits.[16]

We provide our sample selection criteria in Table 1. In Table 1, Panel A, we outline the underlying number of posts shared on StockTwits (labeled as "ideas" on StockTwits and often labeled "tweets"). The StockTwits dataset has 37,037,180 total user posts between July 1, 2010, and December 31, 2015, which is approximately 15,654 posts per day, the majority of which occur during trading hours on weekdays. Of the 37,037,180 posts, 5,160,405 are user flagged as bullish, and 1,281,537 are user flagged as bearish, with user-flagged posts representing 17.4 percent of posts. Of the 82.6 percent of posts that are not user flagged, many express an opinion that readers would not consider neutral but instead express an opinion with some degree of positivity or negativity. Thus, we use machine learning methods to calculate a continuous measure of sentiment for these unflagged posts.

In Table 1, Panel B, we provide information on the number of posts around earnings announcements. We require that a firm has at least three posts in the 24-hour window before the earnings announcement and at least three posts in the 24-hour window after the earnings announcement. This results in a sample of 306,934 user posts in the 24-hour window prior to the earnings announcement and 698,688 user posts in the 24 hours following the earnings announcement. This suggests that the 48 hours surrounding earnings announcements account for approximately 2.71 percent of all posts on StockTwits, with a substantial increase in user posts following the earnings announcement.

---

[14] Prior research has also included firm size and the correlation between prior returns and trading volume. We exclude firm size following the recommendation in (Barron et al. 2018, 1671) that "controlling for firm size in a regression model of abnormal trading volume may filter out some of the effect of interest in studies of investor diversity or differential prior precision around earnings announcements" and exclude the correlation between prior returns and trading volume as it is captured in the absolute return around the earnings announcement. As robustness checks we included both variables and find that our inferences regarding prior disagreement and changes in disagreement are unaffected (not tabulated).

[15] See http://stocktwits.com/developers/docs/start for a description of the raw data provided by StockTwits.

[16] Data prior to the incorporation of user flags is available from the inception of StockTwits on July 9, 2009. Our results are not sensitive to including this sample period (not tabulated).

TABLE 1

Sample Selection

Panel A: StockTwits Data	
Full dataset from July 1, 2010, to December 31, 2015	
Social media observations, n	37,037,180
Unique users, n	204,143
User-identified bullish posts, n	5,160,405
User-identified bearish posts, n	1,281,537
Panel B: Earnings Announcements StockTwits Data	
Posts and users in 48 hours centered on earnings	
Posts in 24 h before	306,934
Posts in 24 h after	698,688
Observations with at least 3 unique users before	22,349
Observations with at least 3 unique users after	27,559
Panel C: Merged Sample Data	
Unique I/B/E/S earnings announcements in the summary database	86,607
Less: Missing Compustat data	(3,404)
Missing CRSP data (249 prior trading days not available)	(6,495)
Missing StockTwits data (3 unique users in 12 hours pre and post)	(56,131)
Total	20,577

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.

In Table 1, Panel C, we outline the sample selection criteria to merge StockTwits data with other data sources. We identify 86,607 unique firm quarters with earnings announcements in the I/B/E/S Detail History database between July 1, 2010 and December 31, 2015. We used the Compustat Securities database for the merge and lost 3,404 observations that either did not have an I/B/E/S ticker or did not have a corresponding quarterly announcement (the Compustat fiscal year is adjusted for the month of year when merging with I/B/E/S). We lose 6,495 observations when merging I/B/E/S with CRSP using the WRDS iclink.sas macro and requiring that data are available on CRSP for the 249 days before the announcement. Finally, we use cashtags and the timestamps for each post to merge individual posts to firm-quarter earnings announcements. We lose 56,131 observations when merging with our StockTwits firm-quarter observations, yielding a sample of 20,577 firm-quarter observations with 1,532,818 posts in the 48 hours around these firm-quarter observations.

Sentiment Classifier Validation

Before proceeding to our empirical analysis, we present model validation statistics for the different measures of post sentiment in Table 2. Our goal for undertaking a sentiment analysis is to calculate the opinions expressed by users who did not flag their posts as either bullish or bearish. Since some users do flag their posts as bullish or bearish;, we can validate the output from machine learning models on a holdout sample of flagged posts. This validation process allows us to provide some assurance about the sentiment measures being calculated for the unflagged posts.

For each machine learning model (Naïve Bayes, Paragraph Vector, and XLNet), we randomly select a sample of 80,000 posts from the 353,662 user-flagged posts, 40,000 of which are selected from user-flagged bullish posts and 40,000 from user-flagged bearish posts. We use these 80,000 posts as our training data to calculate the parameters of each model. We use the parameters from each model to calculate a sentiment measure for the remaining 273,662 user-flagged posts (217,702 bullish and 55,960 bearish) made in the 48 hours around earnings announcements. For each of the 273,662 holdout observations, we compare the calculated sentiment from the model (scores are normalized to 1 for positive sentiment and -1 for negative sentiment in our validation tests) with the user-flagged sentiment (coded as 1 for bullish and -1 for bearish). We calculate three measures of accuracy for each model: precision, recall, and f1-score.

TABLE 2

Sentiment Classifier Validation Statistics

Panel A: Machine Learning Methods

	Precision (Percent)	Recall (Percent)	F1-Score (Percent)	Support
(a) Naïve Bayes				
Bear	45	75	56	55,960
Bull	92	77	84	217,702
Weighted average	83	76	78	273,662
(b) Paragraph vector				
Bear	29	62	4	55,960
Bull	86	62	72	217,702
Weighted average	75	62	65	273,662
(c) XLNet				
Bear	54	79	64	55,960
Bull	94	83	88	217,702
Weighted average	86	82	83	273,662
Panel B: Dictionary Methods				
(d) Full sample Loughran-McDonald dictionary				
Bear	6	2	9	55,960
Bull	45	17	24	217,702
Weighted average	37	17	21	273,662
(e) Restricted sample Loughran-McDonald dictionary				
Bear	35	71	47	16,404
Bull	88	63	74	57,415
Weighted average	77	65	68	73,479
Panel C: Human Classifiers				
(f) Mechanical Turk workers				
Bear	18	8	29	50,000
Bull	47	72	57	50,000
Weighted average	32.5	76	43	100,000

The machine learning models in Panel A are all trained using the same 80,000 posts (40,000 bullish and 40,000 bearish) that were randomly selected from the 353,662 user-flagged posts around earnings announcements. Panel B reports dictionary-based approaches. Classifier (d) is based on the same sample of posts used to validate the machine learning models; note that this results in a score of zero for posts that do not include a word from the Loughran-McDonald dictionary, classifier (e) restricts the sample to only posts that have a positive or negative word in the Loughran-McDonald dictionary. Panel C reports human classifications based on a subsample of 20,000 randomly selected posts (10,000 bullish and 10,000 bearish) that have been classified by five different Mechanical Turk workers. In this case, the Mechanical Turk workers selected "Neutral" for approximately 40 percent of posts.

Variable Definitions:
Precision = the number of true positives divided by the number of true positives and false positives;
Recall = the number of true positives divided by the number of true positives and false negatives; and
F1-score = the harmonic mean of precision and recall.

Precision is defined as the number of true positives divided by the number of true positives and false positives, which is a measure of type-1 classification error. Recall is defined as the number of true positives divided by the number of true positives and false negatives, which is a measure of type-2 error. The f1-score is the harmonic mean of precision and recall, which is a measure of the overall level of correctness of classification.

We report the validation statistics for the machine learning models in Table 2, Panel A. We report the methods from oldest to newest: Naïve Bayes, Paragraph Vector, and XLNet. We find that all methods have a higher precision for classifying bullish over bearish posts. In addition, the methods all have a high recall for classifying bullish posts. We believe this to be related to users being less willing to be overtly bearish, as seen in the total number of user-flagged bullish and bearish posts. The relative accuracy of the models for bullish and bearish posts is much closer, based on recall which measures type 2 errors. Overall, XLNet is the most accurate classification method by all measures, with an f1-score of 83 percent, with the Naïve Bayes at 78 percent. The high level of accuracy of these models suggests that they are all good candidates for distinguishing between bullish and bearish posts.[17]

In Table 2, Panel B, we report similar measures for dictionary-based approaches based on the Loughran-McDonald (LM) financial dictionary. We present these statistics for two samples, the first is the same holdout sample used for the machine learning models, and the second is a sample restricted to only the posts that contain one or more of the words in the Loughran-McDonald dictionary. As expected, most likely due to differences in language use, the dictionary-based models are generally not as good at classifying social media posts. The accuracy of the dictionary models increases from a weighted average f1-score of 21 percent to 68 percent when using the restricted sample, but at a loss of 73 percent of sample observations.

To further assess these models, we include accuracy statistics based on human classification of a randomly selected subsample of 20,000 posts from the StockTwits dataset (split evenly between bullish and bearish flagged posts). For this sample, we had five different Mechanical Turk workers classify each message to further investigate the sources of error. We remove the bullish and bearish icons leaving only the text of the post available to the Mechanical Turk worker. The Mechanical Turk workers were asked to categorize each post as either bullish, bearish, or neutral. We again find that bearish posts are relatively harder to classify. In addition, it is notable that the classification errors are high in part because approximately 40 percent of the posts are classified by the Mechanical Turk workers as neutral. This result is consistent with readers interpreting the posts in a more continuous manner than bullish and bearish.

In Figure 2, we plot the distribution of sentiment using the 1,532,818 posts in the 48 hours around the earnings announcement in our sample. If individuals post more than once in either of the 24-hour windows before or after the same firm's earnings announcement, we average their sentiment scores. The difference in the continuity of the distributions between the dictionary-based measure (LM) and the machine-learning measure distributions (NB, PV, and XL) highlights the advantage of using one of the machine-learning approaches. Specifically, they can create a distribution of sentiment based on the probability that a post is bullish or bearish. On this dimension, the XLNet model appears to provide a more extreme distribution, classifying a significant amount of unflagged posts as 100 percent likely to be bullish, which will bias measures of disagreement.

Descriptive Statistics

We present descriptive statistics in Table 3, Panel A for firms that had at least three different social media users post in the 24 hours before and three users post in the 24 hours after earnings are announced. We present descriptive statistics for firms not discussed on social media in either or both periods in Table 3, Panel B. As expected, in Table 3, Panel A, the average median-adjusted volume is significantly positive around earnings announcements and displays considerable variation, consistent with prior literature. The subset of firms not discussed on social media has a lower average trading volume (mean 0.014 for social media firms versus mean 0.006 for nonsocial media-discussed firms), smaller, and lower per-share prices than the subset of firms discussed on social media. We find that preannouncement disagreement has a range from 0.126 to 0.466 and a median of 0.284, and the change in disagreement has a range from −0.144 to 0.200 and a median of 0.026, consistent with a differential response to disclosures. In Table 3, Panel C, we report correlations. Both the base logged measures of attention and the ranked measures of attention have a similar correlation with $ADJTURN$, this similarity arises from the fact that attention generally increases after the announcement.

We provide the distributions of the self-described trading horizon, strategy, and experience, which we use to calculate investor heterogeneity in Figure 3. StockTwits users have the option to choose to display these self-descriptions on their profile publicly. Whereas Figure 3 suggests that many users chose not to volunteer their attributes, those that do have heterogeneous trading strategies, horizons, and experience. We combine these three different characteristics in our entropy-based measure of heterogeneity, treating no answer as a separate category for each characteristic.

We provide visual evidence of user posting activity, which we use to measure attention, in Figure 4. In Figure 4, Panel A, we report aggregate daily user posts and find a prominent spike on the day of the earnings announcement consistent with prior research (Curtis et al. 2016). In Figure 4, Panel B, we disaggregate posts by user characteristics and plot the results minute by minute. These figures are stacked bar plots in which the total number of posts and the contribution by user characteristics are represented by different shading. We find a prominent spike right around the minute of the earnings announcement, followed by an exponential decay in attention in the hour after earnings are announced. We also find that users with different characteristics are represented both before and after the earnings announcement.

[17] We report validation for a holdout sample that is not used in the training data, which uses a randomly sampled equal proportion of bullish and bearish posts. Another popular machine learning validation method is cross validation, which examines the consistency of a model's parameters by comparing the training results across different subsets of the data. As our sample is unbalanced toward user-flagged posts that are bullish, a standard cross validation exercise could lead to overfitting on the bullish category, which could bias the validation in favor of models that are better able to fit bullish sentiment.

TABLE 3

Descriptive Statistics

Panel A: Firms with Social Media Activity

Variable	n	Mean	Variance	p10	p50	p90
ADJTURN	20,577	0.014	0.000	0.001	0.008	0.037
PDIS	20,577	0.292	0.017	0.126	0.284	0.466
Δ*DIS*	20,577	0.027	0.019	−0.144	0.026	0.200
PATTN (log)	20,577	1.735	0.747	0.849	1.609	2.773
Δ*ATTN (log)*	20,577	0.538	0.314	−0.154	0.511	1.253
PATTN (rank)	20,577	0.500	0.083	0.100	0.500	0.900
Δ*ATTN (rank)*	20,577	0.500	0.083	0.100	0.500	0.900
PRC	20,577	3.295	1.348	1.703	3.506	4.532
ARET	20,577	0.056	0.006	0.006	0.036	0.121
MKTVOL	20,577	0.009	0.000	0.008	0.009	0.011

Panel B: Firms without Social Media Activity

Variable	n	Mean	Variance	p10	p50	p90
ADJTURN	56,131	0.006	0.000	−0.001	0.002	0.017
PRC	56,131	2.700	1.319	1.131	2.834	4.011
ARET	56,131	0.056	0.005	0.006	0.037	0.127
MKTVOL	56,131	0.011	0.000	0.008	0.010	0.013

Panel C: Correlation Matrix for Firms with Social Media Activity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
(1) *ADJTURN*	1	0.14*	0.06*	0.14*	0.06*	0.26*	0.37*	0.26*	0.37*	−0.06*	0.25*	0.08*
(2) *PDIS*	0.15*	1	−0.61*	1.00*	−0.61*	0.23*	0.09*	0.23*	0.09*	−0.01	0.10*	−0.08*
(3) Δ*DIS*	0.05*	−0.63*	1	−0.61*	1.00*	−0.06*	0.18*	−0.06*	0.18*	−0.01*	−0.01	−0.02*
(4) *PDIS (rank)*	0.15*	0.98*	−0.62*	1	−0.61*	0.23*	0.09*	0.23*	0.09*	−0.01	0.10*	−0.08*
(5) Δ*DIS (rank)*	0.04*	−0.61*	0.96*	−0.61*	1	−0.06*	0.18*	−0.06*	0.18*	−0.01*	−0.01	−0.02*
(6) *PATTN (log)*	0.27*	0.25*	−0.05*	0.26*	−0.06*	1	−0.11*	1.00*	−0.11*	0.06*	0.08*	−0.03*
(7) Δ*ATTN (log)*	0.36*	0.08*	0.18*	0.078*	0.18*	−0.08*	1	−0.11*	1.00*	0.04*	0.02*	−0.02*
(8) *PATTN (rank)*	0.24*	0.21*	−0.05*	0.23*	−0.06*	0.91*	−0.13*	1	−0.11*	0.06*	0.08*	−0.03*
(9) Δ*ATTN (rank)*	0.34*	0.10*	0.18*	0.09*	0.18*	−0.06*	0.96*	−0.11*	1	0.04*	0.02*	−0.02*
(10) *PRC*	−0.06*	−0.04*	−0.01*	−0.03*	−0.01*	0.10*	0.04*	0.06*	0.05*	1	−0.25*	0.06*
(11) *ARET*	0.26*	0.10*	−0.01*	0.1*	−0.01*	0.09*	−0.06*	0.08*	−0.05*	−0.29*	1	0.06*
(12) *MKTVOL*	0.08*	−0.08*	−0.02*	−0.08*	−0.02*	−0.01	−0.03*	−0.02*	−0.02*	0.06*	0.06*	1

* Denotes significance at the < 5 percent level.

In Panel C we report Spearman rank correlations above the diagonal and Pearson parametric correlations below the diagonal; the number of firm-quarter observations is 20,577.

Variable Definitions:

ADJTURN = median-adjusted turnover calculated as the three-day earnings turnover less the median three-day turnover in the previous 249 trading days;

PDIS = the standard deviation of user sentiment in the 24 hours before earnings are announced;

Δ*DIS* = the standard deviation of sentiment in the 24 hours after earnings less the standard deviation of sentiment in the 24 hours before earnings;

PATTN (*log*) = the natural log of the number of active users in the 24 hours before earnings are announced less the average daily attention in days −10 to −30 relative to the announcement;

Δ*ATTN* (*log*) = the natural log of the number of active users in the 24 hours after earnings less the number of active users in the 24 hours after earnings are announced;

PATTN (*rank*) = the within-sample percentage rank of *PATTN* (*log*);

Δ*ATTN* (*rank*) = the within-sample percentage rank of Δ*ATTN* (*log*);

PRC = the natural logarithm of the market price of a share of common stock two days before the earnings announcement;

ARET = the absolute value of returns in the three trading days centered on the earnings announcement; and

MKTVOL = the turnover for all firms in the CRSP dataset in the three days centered on earnings.

FIGURE 3
Shows Self-Reported Investor Heterogeneity

Panel A: Count of Distinct Users by Attribute

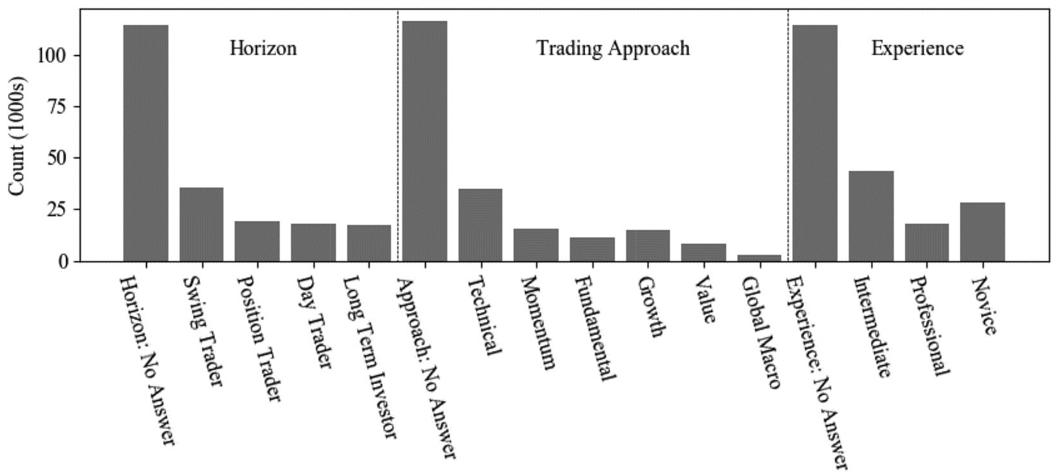

Panel B: Count of Posts by Attribute

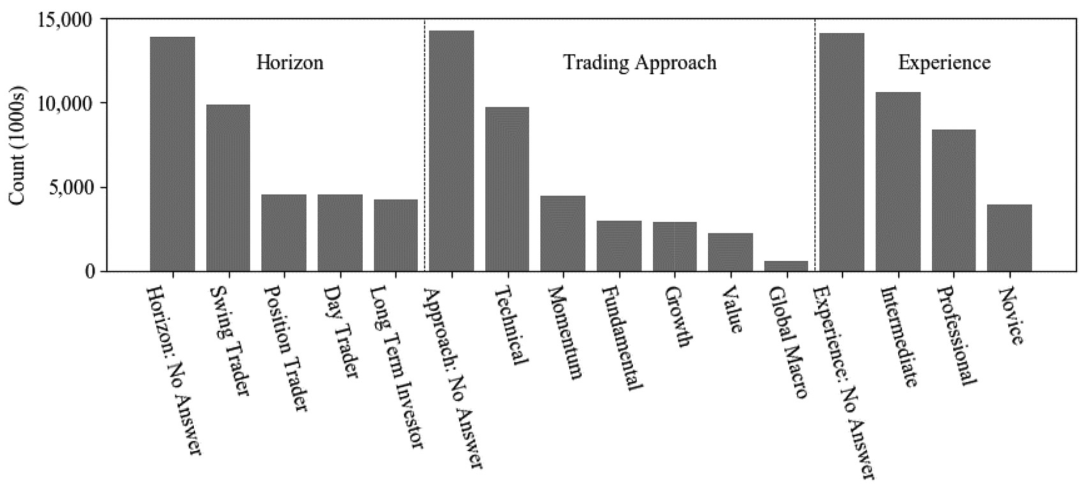

(Panel A) Count of distinct users by attribute. (Panel B) Count of posts by attribute. For each of these three categories, individuals can select from the choices above or leave them blank (in which case they are included in the "no answer" category). Panel A plots the distribution of these responses by the distinct user, and Panel B plots the distribution of these responses by post. The plots represent 1,532,818 posts from the 48 hours around the earnings announcements in our sample.

Tests of Hypothesis

Following the prior literature, we anticipate that both preannouncement disagreement and changes in disagreement around an earnings announcement are positively associated with trading volume. To test this hypothesis, we consider the following regression model:

$$ADJTURN_{j,t} = \alpha + \beta_1 PDIS_{j,t} + \beta_2 \Delta DIS_{j,t} + \sum \beta_i Controls_{j,t} + \gamma_1 YearFE_t + \gamma_2 MonthFE_t + \gamma_3 IndustryFE_j + e_{j,t}, \quad (8)$$

where $ADJTURN$ for firm j is measured over the three days around an earnings announcement at time $t = 0$. We are most interested in the investor disagreement variables based on the explicit opinions of users. Following H_1, we expect that $\beta_1 > 0$ and $\beta_2 > 0$.

FIGURE 4

Shows Time Series Plots of Post Volume Around Earnings Announcements

Panel A: Daily Post Volume Around Earnings Announcements

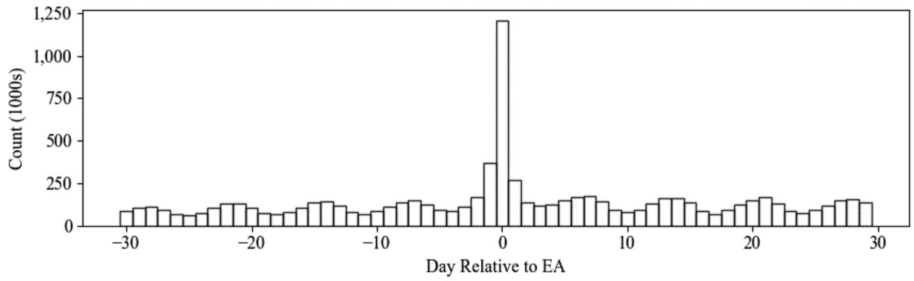

Panel B: Minute Post Volume By Investor Type Around Earnings Announcements

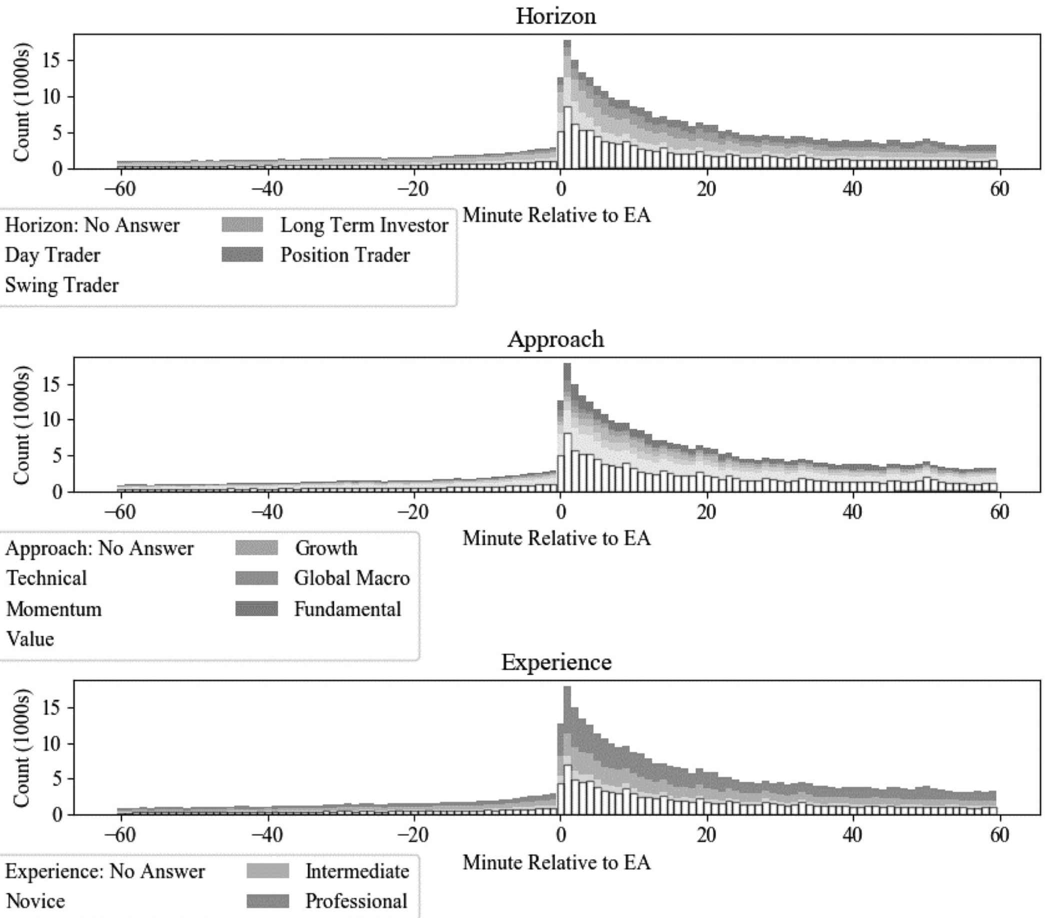

(Panel A) Daily post volume around earnings announcements. (Panel B) Minute post volume by investor type around earnings announcements. Figures are constructed using our sample of 20,577 quarterly earnings announcements. Panel B plots minute-by-minute post volume by StockTwits users' self-described trading horizon, approach, and experience. For each of these three categories, individuals can select from the choices above or leave them blank (in which case they are included in the "no answer" category).

We report estimates of Equation (8) in Table 4 based on four different sentiment analysis methods, disagreement based on only user-flagged posts (USR) and a model that combines the information in the Naïve Bayes classifier combined with the information in user-flagged posts (NB + USR). The NB + USR model replaces the Naïve Bayes calculated sentiment with −1 for bearish-flagged posts and 1 for bullish-flagged posts. All regressions include the full sample

TABLE 4

Different Sentiment Measurement Methods

$$ADJTURN_j = \alpha + \beta_1 PDIS_j + \beta_2 \Delta DIS_j + \sum Controls_j + \sum FE + e_j$$

	(1) NB + USR	(2) LM	(3) PV	(4) NB	(5) XL	(6) USR
PDIS	0.046***	0.023***	0.035***	0.055***	0.011***	0.028***
	(18.79)	(17.81)	(12.51)	(17.77)	(11.08)	(17.35)
ΔDIS	0.034***	0.015***	0.026***	0.036***	0.004***	0.021***
	(21.12)	(18.07)	(12.79)	(18.14)	(6.06)	(19.54)
PRC	0.000	0.001*	0.001*	0.000	0.001**	0.001***
	(1.56)	(1.80)	(1.79)	(0.78)	(2.02)	(2.70)
ARET	0.098***	0.102***	0.106***	0.102***	0.105***	0.100***
	(19.81)	(20.13)	(20.10)	(19.97)	(20.24)	(19.96)
MKTVOL	0.528***	0.467***	0.493***	0.435***	0.486***	0.763***
	(4.29)	(3.78)	(3.99)	(3.53)	(3.94)	(6.19)
Observations	20,577	20,577	20,577	20,577	20,577	20,577
Adj. R^2	0.53	0.52	0.51	0.52	0.51	0.53
Sentiment distribution		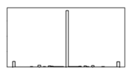		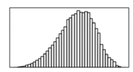	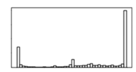	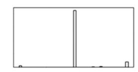

***, **, * Denote significance at the less than 10 percent, 5 percent, and 1 percent levels, respectively.
Column (1) combines NB and USR by replacing user-flagged posts with −1 for bearish and 1 for bullish flagged posts. All models include industry, year, and month fixed effects. Sentiment distribution is the underlying sentiment of posts in the 48 hours centered on earnings. The t-statistics reported in parentheses are based on standard errors clustered by firm: industry, year, and month fixed effects are included.

Variable Definitions:
LM = the Loughran-McDonald dictionary-based classifier;
NB = the Naïve Bayes classifier;
PV = the paragraph vector classifier;
USR = the self-reported user sentiment classification with unflagged posts set to 0;
XL = the XLNet classifier;
ADJTURN = median-adjusted turnover calculated as the three-day earnings turnover less the median three-day turnover in the previous 249 trading days;
PDIS = the standardized standard deviation of user sentiment in the 24 hours before earnings are announced;
ΔDIS = the standardized standard deviation of sentiment in the 24 hours after earnings less the standard deviation of sentiment in the 24 hours before earnings;
PRC = the natural logarithm of the market price of a share of common stock two days before the earnings announcement;
ARET = the absolute value of returns in the three trading days centered on the earnings announcement; and
MKTVOL = the turnover for all firms in the CRSP dataset in the three days centered on earnings.

of 20,577 firm-quarter observations and have median adjusted volume (*ADJTURN*) as the dependent variable. We report the results for opinions derived using the Naïve Bayes plus user-flagged posts (NB + USR) in column (1), where the coefficient on prior user disagreement is positive and significant, $\beta_1 = 0.046\,(\mathrm{t} = 18.79)$ as is the change in user disagreement, $\beta_2 = 0.034\,(\mathrm{t} = 21.12)$. In column (2), we report the Loughran-McDonald dictionary-based approach. The coefficients on the disagreement measures remain positive and significant, with $\beta_1 = 0.023\,(\mathrm{t} = 17.81)$, $\beta_2 = 0.015\ (\mathrm{t} = 18.07)$. The dictionary-derived results should be interpreted with caution as 10,094 of the 22,577 observations do not have a post that used a word from the dictionary corpus, meaning that if disagreement between investors does exist, it is arbitrarily set to zero for that earnings announcement. In column (3), the paragraph vector-derived sentiment is positive and significant, with $\beta_1 = 0.035\,(\mathrm{t} = 12.51)$ and $\beta_2 = 0.026\,(\mathrm{t} = 12.79)$. In column (4), the Naïve Bayes-derived user opinion, $\beta_1 = 0.055(\mathrm{t} = 17.77)$ and $\beta_2 = 0.036\,(\mathrm{t} = 18.14)$. In column (5), the XLNet-derived user opinion, $\beta_1 = 0.011\,(\mathrm{t} = 11.08)$ and $\beta_2 = 0.004\,(\mathrm{t} = 6.06)$, the lower coefficient likely due to a downward bias in disagreement expected from the underlying sentiment distribution not capturing the null posts well. In column (6), we provide the results of user-given sentiment-derived disagreement and find that both user disagreement and change in disagreement are positively associated with trading volume. Collectively, these results provide strong support for H1 and evidence that our measure of disagreement is robust to the use of different sentiment analysis techniques.

The control variables included in the model are generally consistent with predictions based on prior literature.[18] The level of marketwide trading volume is positively associated with trading volume around earnings announcements. The coefficients on the absolute value of the returns are positive and significant. Share price two days before earnings are announced (PRC) is positively and often significantly associated with trading volume, consistent with lower transaction costs at higher prices facilitating information-based trading.

Examination of Additional Predictions

Our research questions relate to the possibility of rational inattention theories having an interactive effect on the relation between disagreement and trading volume. To test our research questions, we consider regression models that interact proxies for attention with disagreement, both before and with the change after the earnings announcement. Specifically, we consider regression models of the type:

$$\begin{aligned} ADJTURN_j = \alpha + \beta_1 PDIS_j + \beta_2 \Delta DIS_j + \beta_3 PATTN_j + \beta_4 \Delta ATTN_j + \beta_5 (PATTN_j \times PDIS_j) \\ + \beta_6 (PATTN_j \times \Delta DIS_j) + \beta_7 (\Delta ATTN_j \times \Delta DIS_j) + \sum \beta_i Controls_{j,t} + \gamma_1 Year\,FE_t \\ + \gamma_2 Month\,FE_t + \gamma_3 Industry\,FE_j + e_{j,t}. \end{aligned} \quad (9)$$

We begin our analysis by examining the incremental effects of including attention and the change in attention around the earnings announcement. In column (1) of Table 5, we report the main effects of our measures of investor disagreement without controls or fixed effects. We find that both measures of disagreement are positive and significant ($\beta_1 = 0.051$, t = 18.28; $\beta_2 = 0.038$, t = 20.71) and that these measures of disagreement explain five percent of the variation in trading volume. These results are like those reported by Cookson and Niessner (2020, 213), who conclude that social media-based measures of disagreement have only weak explanatory power over the spike in trading volume around earnings announcements using a daily model of disagreement. As earnings announcements require investors to process the information in accounting disclosures, we expect attention to have a significant effect on the link between disagreement and trading around earnings announcements. In column (2), we add attention levels and find that these increase the explained variation fourfold to an adjusted R^2 of 20 percent. The coefficients on our measures of disagreement decrease to $\beta_1 = 0.017$ and $\beta_2 = 0.009$, consistent with disclosure processing costs playing a role in the link between investor disagreement and trading volume.

Social media posts could also lower processing costs for other investors as they often relay the highlights of investors' information searches, thereby reducing acquisition and integration processing costs for other investors reading these posts around disclosure events (Curtis et al. 2016; Blankespoor et al. 2020). Like any information source, there is noise in social media posts. However, prior literature has found evidence that individual investors' online posts contain value-relevant information (Bagnoli, Beneish, and Watts 1999; Antweiler and Frank 2004; Bartov, Faurel, and Mohanram 2018).[19] Elevated activity on social media also potentially reduces awareness costs about the earnings announcement, as increased social media posts may be attention-grabbing news to other investors (Da, Engelberg, and Gao 2011; Barber and Odean 2008). As such, increased attention to social media can impact awareness, acquisition, and integration costs. With lower disclosure processing costs, earnings news becomes available to a broader investor base with more diverse priors and can change more investors' priors.

The interaction between disagreement and attention describes how lower disclosure processing costs impact disagreement. In columns (3) to (9), we include the three attention interactions $PATTN \times PDIS$, $PATTN \times \Delta DIS$, and $\Delta ATTN \times \Delta DIS$. In each column, we report our model for specifications that differ based on the inclusion of different time-varying controls and fixed effects. Across all specifications, we find that all three attention interactions are reliably positive and statistically significant. The main effects for $PATTN$ and $\Delta ATTN$ are also reliably positive and statistically significant. In contrast, the main effect of $PDIS$ and ΔDIS measures are inconsistent across the specifications. In column (5), we find that the main effect of change in disagreement becomes negative and significant when $ARET$ is included in the model. Since $ARET$ is a proxy for the change in priors, it is likely the correlation between $ARET$ and our

[18] We also considered specifications that include firm size, measured as the logarithm of market capitalization (not tabulated). The relation between trading volume and size is negative when included in the equation in Equation (8), which suggests that conditional on disagreement larger firms have incrementally lower trading volume. This is consistent with the increasing diversity of institutional firms following large firms found in Barron et al. (2018). We find the relation between size and trading volume is positive for a sample of firms with available data in I/B/E/S, Compustat, and CRSP, but without a requirement for social media activity (not tabulated).

[19] We examined a random sample of StockTwits posts before and after earnings, finding both posts which discuss earnings information and other posts discussing technical trading strategies. For example, some posts contain discussion of important components of earnings and additional information that highlight the context of the information in the earnings disclosures, which could aid other users process earnings disclosures.

TABLE 5

Attention and Disagreement

$$ADJTURN_j = \alpha + \beta_1 PDIS_j + \beta_2 \Delta DIS_j + \beta_3 PATTN_j + \beta_4 \Delta ATTN_j + \beta_5 (PATTN_j \times PDIS_j) + \beta_6 (PATTN_j \times \Delta DIS_j) + \beta_7 (\Delta ATTN_j \times \Delta DIS_j) + \sum Controls_j + \sum FE + e_j$$

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PDIS	0.051***	0.017***	0.007*	0.004	−0.000	0.009**	0.020***	0.002	0.012***
	(18.28)	(7.29)	(1.67)	(1.09)	(−0.04)	(2.27)	(5.77)	(0.55)	(3.53)
ΔDIS	0.038***	0.009***	−0.005	−0.006*	−0.008**	−0.005	−0.001	−0.007**	−0.005
	(20.71)	(5.84)	(−1.49)	(−1.65)	(−2.40)	(−1.46)	(−0.45)	(−2.09)	(−1.45)
PATTN		0.020***	0.012***	0.012***	0.011***	0.011***	0.010***	0.007***	0.007***
		(15.72)	(4.39)	(4.59)	(4.54)	(4.25)	(4.28)	(3.53)	(3.43)
$\Delta ATTN$		0.027***	0.026***	0.027***	0.028***	0.026***	0.024***	0.022***	0.023***
		(26.28)	(24.87)	(26.04)	(27.26)	(24.93)	(25.58)	(23.56)	(27.00)
$PATTN \times PDIS$			0.025***	0.025***	0.023***	0.026***	0.026***	0.024***	0.024***
			(2.94)	(2.99)	(2.84)	(3.18)	(3.56)	(3.56)	(3.53)
$PATTN \times \Delta DIS$			0.013**	0.013**	0.012*	0.016**	0.016***	0.014**	0.014***
			(2.13)	(2.09)	(1.91)	(2.51)	(2.83)	(2.45)	(2.62)
$\Delta ATTN \times \Delta DIS$			0.019***	0.017***	0.017***	0.021***	0.027***	0.017***	0.024***
			(4.57)	(4.13)	(4.31)	(5.15)	(7.02)	(4.28)	(6.48)
PRC				−0.002***					−0.000
				(−6.17)					(−1.37)
ARET					0.113***				0.101***
					(21.59)				(20.26)
MKTVOL						1.634***			0.491***
						(10.29)			(4.16)
Constant	0.001	−0.012***	−0.008***	−0.003*	−0.011***	−0.024***			
	(1.43)	(−12.87)	(−7.03)	(−1.79)	(−9.47)	(−11.96)			
Observations	20,577	20,577	20,577	20,577	20,577	20,577	20,577	20,577	20,577
Adj. R^2	0.05	0.20	0.20	0.21	0.26	0.21	0.53	0.53	0.59
Year FE	no	no	no	no	no	no	yes	no	yes
Month FE	no	no	no	no	no	no	yes	no	yes
Industry FE	no	no	no	no	no	no	no	yes	yes

***, **, * Denote significance at the less than 10 percent, 5 percent, and 1 percent levels, respectively.
The t-statistics reported in parentheses are based on standard errors clustered by firm.

Variable Definitions:
ADJTURN = median-adjusted turnover calculated as the three-day earnings turnover less the median three-day turnover in the previous 249 trading days;
PDIS = the standard deviation of user sentiment in the 24 hours before earnings are announced;
ΔDIS = the standard deviation of sentiment in the 24 hours after earnings less the standard deviation of sentiment in the 24 hours before earnings;
PATTN = the ranked percentage of the natural logarithm of the number of users in the 24 hours before earnings are announced less the average daily number of users in the −10 to −30 days before earnings; and
$\Delta ATTN$ = the ranked percentage of the natural logarithm of the number of users in the 24 hours after earnings are announced less the natural logarithm of the number of users in the 24 hours after earnings are announced.

We also report control variables:
PRC = the natural logarithm of the market price of a share of common stock two days before the earnings announcement;
ARET = the absolute value of returns in the three trading days centered on the earnings announcement; and
MKTVOL = the turnover for all firms in the CRSP dataset in the three days centered on earnings.

disagreement measures that is driving the coefficient on ΔDIS negative. We examine the impact of time and industry fixed effects in columns (7) and (8) and find little impact on the interactions of interest. Accounting for the impact of omitting the intercept on R^2, the fixed effects explain approximately 4 percent of the variation.[20] Our results indicate that disagreement-based trading activity around earnings announcements is significantly higher when processing costs are lower.[21]

Using the results in column (9) as an example, we interpret the interactive effects as how disclosure processing costs impact different aspects of disagreement. First, the coefficient on the interactive effect $PATTN \times PDIS$ ($\beta_5 = 0.024$, t = 3.53) provides evidence that lower disclosure processing costs amplify trading related to existing differences in priors. As the variable ΔDIS is positive in our sample on average, the positive coefficient on $PATTN \times \Delta DIS$ ($\beta_6 = 0.014$, t = 2.62) suggests lower processing costs before the earnings announcement are associated with increasing trading associated with increasing divergence of beliefs around earnings announcements. Similarly, as the variable $\Delta ATTN$ is ranked, the positive coefficient on $\Delta PATTN \times \Delta DIS$ ($\beta_7 = 0.024$, t = 6.48) is consistent with lower disclosure processing costs following the disclosure increasing trading associated with increasing divergence of beliefs around earnings announcements. We conjecture that the effects of preannouncement attention are most likely consistent with awareness, and postannouncement attention is most likely consistent with acquisition and integration costs.

We examine alternate proxies for processing costs in Table 6. We include our attention result from Table 5 in column (1). We include our percentage ranked measure of heterogeneity in column (2) and find similar results, consistent with a varied set of opinions helping to decrease processing costs. In columns (3) and (4), we examine whether user influence appears to moderate the association between disagreement and trading volume by including indicator variables for a post from one of the top 1,000 most interactive users by Eigen centrality (column (3)), and a post from one or more of the top 1,000 most-followed users (column (4)).[22] To examine the possibility that user influence could impact trading, we include indicator variables for posts from high-visibility users. In both columns (3) and (4), the main effects remain positive and significant, and the interaction [high vis char] $\times$ predisagreement is the only significant interaction ($\beta_5 = 0.008$, t = 2.43). This provides evidence that high-visibility users may draw additional attention to the announcement, but they do not necessarily help investors process the announcement.

V. FURTHER ANALYSIS

Comparison to Dispersion in Analysts' Forecasts

For a subset of our analysis, to highlight the additional information in social media posts regarding current market events, we compare disagreement between analysts regarding annual earnings estimates and investor disagreement on StockTwits. A major difference between analyst forecast-derived measures and social media investor-derived measures of investor disagreement is that opinions about firm value on social media are not explicitly tied to an earnings forecast. In addition, the population of analysts is relatively sophisticated in terms of their financial knowledge, at least relative to the full population of investors. These reasons imply that disagreement based on social media opinions is likely distinct from analyst forecast measures of disagreement. We formally test this assertion in this section.

To test this assertion, we follow Bamber et al. (1997) and calculate the natural log of the standard deviation of analysts' annual earnings forecasts ($ANDIS$) t around quarterly earnings announcements. We calculate the jumbling of forecasts (*jumbling*) as 1.1 less the Pearson correlation coefficient for analysts' estimates in the 45 days prior and the 30 days after the announcement, where 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 logarithm of the standard deviation in the period after earnings less the logarithm of the standard deviation of analyst estimates in the period before ($\Delta ANDIS$). We require each firm-quarter observation to have three or more

[20] The total sum of squares (SS_{tot}) in models without an intercept increases with the variance of the predictor variables, which mechanically increases the R^2. The R^2 for the specification in columns (7) and (8) without fixed effects is 0.49. This specification (without an intercept) is consistent with the calculation of the R^2 in Cookson and Niessner (2020).

[21] We examine cross-sectional variation in abnormal trading volume around the earnings announcement. In contrast, Cookson and Niessner (2020) examine the increased trading on earnings announcement days relative to one week prior and three weeks after the earnings announcement. Their approach is based on observing a lower coefficient on the indicator for earnings announcement days (EA) when disagreement is added to their regression model. They find that the coefficient EA only declines by one-eighth when including disagreement and do not consider attention in their test. As a specification check, when we replicate their model and included user attention, we find that the coefficient on EA is insignificant, consistent with our findings that disagreement and user attention combined explain cross-sectional variation in abnormal trading on the earnings announcement day (not tabulated).

[22] We verify that the distribution of importance across our dataset follows a Pareto distribution in which the top 1,000 most visible users account for a significant majority of the mentions and comments from other users in the network (not tabulated).

TABLE 6

Disagreement and Trading Volume Interactions with User Characteristics

$$ADJTURN_j = \alpha + \beta_1 PDIS_j + \beta_2 \Delta DIS_j + \beta_3 PATTN_j + \beta_4 [CHAR]_j + \beta_5 \left([CHAR]_j \times PDIS_j\right) + \beta_6 \left([CHAR]_j \times \Delta DIS_j\right) + \beta_7 \left([\Delta CHAR] \times \Delta DIS_j\right) + \sum Controls_j + \sum FE + e_j$$

[CHAR] =	(1) Attention	(2) Heterogeneity	(3) Engagement	(4) Followers
PDIS	0.012***	0.008**	0.040***	0.041***
	(3.53)	(2.55)	(15.62)	(14.72)
ΔDIS	−0.005	0.001	0.034***	0.035***
	(−1.45)	(0.27)	(17.82)	(16.90)
[CHAR]	0.007***	0.007***	−0.002	0.002
	(3.43)	(3.33)	(−1.20)	(0.46)
$[\Delta CHAR]$	0.023***	0.004***	0.002*	0.001
	(27.00)	(19.76)	(1.74)	(0.38)
$[CHAR] \times PDIS$	0.024***	0.039***	0.009***	0.008**
	(3.53)	(5.46)	(2.74)	(2.43)
$[CHAR] \times \Delta DIS$	0.014***	0.039***	−0.009	−0.009
	(2.62)	(6.27)	(−1.41)	(−0.48)
$[\Delta CHAR] \times \Delta DIS$	0.024***	0.004***	0.008	0.007
	(6.48)	(4.32)	(1.30)	(0.39)
PRC	−0.000	−0.000	0.000	0.000
	(−1.37)	(−0.66)	(1.34)	(1.19)
ARET	0.101***	0.096***	0.098***	0.098***
	(20.26)	(19.62)	(19.76)	(19.75)
MKTVOL	0.491***	0.660***	0.495***	0.435***
	(4.16)	(5.38)	(4.02)	(3.55)
Observations	20,577	20,577	20,577	20,577
Adj. R^2	0.59	0.55	0.53	0.53

***, **, * Denote significance at the less than 10 percent, 5 percent, and 1 percent levels, respectively.
All models include industry, year, and month fixed effects. The t-statistics reported in parentheses are based on standard errors clustered by firm: industry, year, and month fixed effects are included.

Variable Definitions:
ADJTURN = median-adjusted turnover calculated as the three-day earnings turnover less the median three-day turnover in the previous 249 trading days;
PDIS = the standard deviation of user sentiment in the 24 hours before earnings are announced;
ΔDIS = the standard deviation of sentiment in the 24 hours after earnings less the standard deviation of sentiment in the 24 hours before earnings;
char = a measure of user characteristics which we vary across columns;
Attention = the ranked percentage of *Attention* (from 0 to 1), and change in *Attention* is the ranked percentage change in *Attention*;
Heterogeneity = the ranked percentage *Heterogeneity* (from 0 to 1), and change in *Heterogeneity* is the ranked percentage change in *Heterogeneity*;
Engagement = an indicator variable that takes a value of one if a user that is in the top 1,000 most central users posts in the time before earnings, and change in *Engagement* takes the value of 1 if one of the top 1,000 most central users posted in the 24 hours after earnings and none of these users posted in the 24 hours before earnings; and
Followers = an indicator variable that takes a value of 1 if a user that is in the top 1,000 most followed users posts in the time before earnings, and change in *Followers* takes the value of 1 if one of the top 1,000 most followed users posted in the 24 hours after earnings and none of these users posted in the 24 hours before earnings.

We also report control variables:
PRC = the natural logarithm of the market price of a share of common stock two days before the earnings announcement;
ARET = the absolute value of returns in the three trading days centered on the earnings announcement; and
MKTVOL = the turnover for all firms in the CRSP dataset in the three days centered on earnings.

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. We add these proxies for analyst disagreement into our regression model in Equation (8), resulting in the following:

$$\begin{aligned} ADJTURN_j = \alpha &+ \beta_1 PDIS_{j,t} + \beta_2 \Delta DIS_{j,t} + \beta_3 ANDIS_j + \beta_4 \Delta ANDIS_j + \beta_5 ANJUMB_j \\ &+ \beta_6 PDIS_{j,t}\, x\, [CHAR_{j,t}] + \beta_7 \Delta DIS_{j,t}\, x\, [CHAR_{j,t}] + \beta_8 ANDIS_j\, x\, [CHAR_{j,t}] \\ &+ \beta_9 \Delta ANDIS_j x\, [CHAR_{j,t}] + \beta_{10} ANJUMB_j x\, [CHAR_{j,t}] + \sum \beta_i Controls_{j,t} + \gamma_1 Year\, FE_t \\ &+ \gamma_2 Month\, FE_t + \gamma_3 Industry\, FE_j + e_{j,t}, \end{aligned} \tag{10}$$

where the subscript j represents an earnings announcement from firm j, and [$CHAR$] is an indicator variable for high news coverage before (2), high news coverage after (3), high net income dispersion (4), high $PATTN$ (5), and high attention post (6).

We report the results in Table 7. We limit this analysis to the subsample of our firms with at least three analyst earnings revisions before and after the earnings announcement. In column (1), we provide the results for the main effects in Equation (10). All measures of disagreement, both social media and analyst based, are significant.[23] This result suggests that sentiment provides incremental information regarding the level of investor disagreement.

We next examine whether alternate measures of attention interact with our social media measures of disagreement and analyst-based measures of disagreement. In general, attention interacts with social media measures of disagreement more consistently than analyst-based measures of disagreement. In column (2), we examine the impact of traditional news on our measures of disagreement using an indicator variable that is equal to 1 if the number of news articles in the 24 hours before earnings is in the top quintile in our sample and find evidence that press coverage in the day before the earnings announcement does not significantly impact the main results. Firms with high press coverage the day before the earnings announcement have a lower coefficient on the interaction $\Delta DIS \times$ *High News Coverage Pre* ($\beta = -0.015$, t $= -2.03$). In the context of retail investors searching for the highest gains, we interpret this as social media disagreement having a weaker association with trading volume when more traditional information intermediaries are available. We examine the interactive effects of news coverage on the day after the earnings announcement in column (3) and find little impact on the main effects and a marginally significant association between $ANJUMB \times$ *High News Coverage Post* ($\beta = -0.003$, t $= -1.79$). In column (4), we report interactions between disagreement and firms with a high standard deviation of GAAP earnings. We find a negative incremental effect for both preannouncement social media disagreement ($\beta = -0.031$, t $= -3.77$) and change in disagreement ($\beta = -0.017$, t $= -2.45$), with none of the other interactions significant at conventional levels. These results are consistent with a smaller revision of prior disagreements, potentially consistent with earnings being a noisier signal for those firms. In columns (5) and (6), we report interactions with an indicator for firms with high $PATTN$ and post the earnings announcement using our social media measure of attention. The interactions with social media attention after the earnings announcement indicate that attention on social media helps investors quickly process the implications for future earnings as well as other value-relevant components of earnings.

Additional Cross-Sectional Evidence

We next examine potential cross-sectional differences in the association between disagreement and trading volume. We examine several proxies that together provide insight into the potential cross-sectional variation of differences in earnings news, existing information, and attention to the earnings announcement. For each proxy, we sort firms into quintile portfolios, with portfolio 1 being the lowest and quintile 5 being the highest. We report the results of this analysis in Table 8. Due to the large number of tests reported in Table 8, we only report the coefficients on $PDIS$ and ΔDIS to preserve space.

We consider variation in earnings news by examining the variation in absolute standardized earnings surprises and signed earnings surprises. In Table 8, Panel A, we report an increased role for prior disagreement and the change in disagreement around earnings announcements with the largest absolute earnings surprises. Specifically, in column (1), for the smallest absolute earnings surprises, the coefficients on $PDIS$ ($\beta = 0.034$, t $= 9.28$) and ΔDIS ($\beta = 0.026$, t $= 9.67$) are smaller than those in column (5), for the largest absolute earnings surprises, with coefficients $PDIS$ ($\beta = 0.061$, t $= 14.08$) and ΔDIS ($\beta = 0.046$, t $= 13.64$). In Table 8, Panel B, we find that this effect is asymmetric, with extreme positive surprises (column (5)) playing a greater role than extreme negative earnings surprises (column (1)).[24]

[23] Analyst dispersion and change in dispersion are only significant when logged, indicating dispersion is somewhat heteroskedastic.
[24] We also examine the difference between positive and negative news by separating into firms that meet/beat versus those that miss and find no statistical evidence of a differential role for positive versus negative earnings surprises (not tabulated).

TABLE 7

Comparison of Social Media and Analyst Forecast Measures of Disagreement

$$ADJTURN_j = \alpha + \beta_1 PDIS_j + \beta_2 \Delta DIS_j + \beta_3 PATTN_j + \beta_4 [CHAR]_j + \beta_5 \left([CHAR]_j \times PDIS_j\right) + \beta_6 \left([CHAR]_j \times \Delta DIS_j\right) + \beta_7 \left([\Delta CHAR] \times \Delta DIS_j\right) + \sum Controls_j + \sum FE + e_j$$

	(1) Base Model	*[CHAR] =*	(2) *[High News Pre]*	(3) *[High News Post]*	(4) *[High NI Dispersion]*	(5) *[High PATTN]*	(6) *[High Attention Post]*
PDIS	0.052***		0.052***	0.050***	0.054***	0.038***	0.031***
	(14.27)		(13.12)	(13.59)	(13.51)	(12.29)	(11.29)
ΔDIS	0.036***		0.036***	0.034***	0.037***	0.028***	0.022***
	(14.28)		(13.34)	(13.23)	(13.41)	(12.17)	(10.67)
ANDIS	0.001***		0.001*	0.001*	0.001**	0.000	0.000
	(2.60)		(1.78)	(1.75)	(2.47)	(1.43)	(0.46)
ΔANDIS	0.001***		0.001***	0.001***	0.001***	0.001***	0.001*
	(3.58)		(2.82)	(2.81)	(3.30)	(2.80)	(1.95)
ANJUMB	0.002***		0.002***	0.002***	0.003***	0.002***	0.001***
	(5.01)		(4.67)	(4.04)	(5.31)	(4.67)	(3.46)
[CHAR]			0.002	0.002	0.012***	−0.003	0.006
			(0.38)	(0.31)	(4.05)	(−0.66)	(1.37)
PDIS × [CHAR]			−0.006	0.005	−0.031***	0.032***	0.026***
			(−0.63)	(0.51)	(−3.77)	(3.40)	(2.66)
ΔDIS × [CHAR]			−0.015**	0.004	−0.017**	0.018*	0.033***
			(−2.03)	(0.63)	(−2.45)	(1.87)	(3.32)
ANDIS × [CHAR]			−0.000	0.000	−0.000	0.000	0.001*
			(−0.05)	(0.11)	(−0.18)	(0.79)	(1.73)
ΔANDIS × [CHAR]			0.000	0.001	−0.000	0.000	0.001
			(0.35)	(0.64)	(−0.27)	(0.16)	(0.93)
ANJUMB × [CHAR]			−0.001	0.003*	−0.001	0.002	0.005***
			(−0.46)	(1.79)	(−0.91)	(1.13)	(3.18)
Observations	7,229		6,424	6,605	6,702	7,229	7,229
Adj. R^2	0.60		0.61	0.61	0.59	0.61	0.64
Controls	Included		Included	Included	Included	Included	Included

***, **, * Denote significance at the less than 10 percent, 5 percent, and 1 percent levels, respectively.
The t-statistics reported in parentheses are based on standard errors clustered by firm: industry, year, and month fixed effects are included.

Variable Definitions:
ADJTURN = median-adjusted turnover calculated as the three-day earnings turnover less the median three-day turnover in the previous 249 trading days.

Our social media-based disagreement measures:
PDIS = the standard deviation of user sentiment in the 24 hours before earnings are announced; and
ΔDIS = the standard deviation of sentiment in the 24 hours after earnings less the standard deviation of sentiment in the 24 hours before earnings.

Our analysts measure of disagreement:
ANDIS = the standard deviation of analyst Earnings Per Share (EPS) forecasts before the earnings announcement;
ΔANDIS = the change in analyst EPS forecasts after the earnings announcement;
ANJUMB = 1.1 minus the person correlation coefficient between annual earnings estimates around the quarterly announcement;
High News Pre = a dummy variable equal to 1 for observations for which the number of news articles in the 24 hours before earnings was in the top quintile;
High News Post = a dummy variable equal to 1 for observations for which the number of news articles in the 24 hours after earnings were in the top quintile;
High NI Dispersion = a dummy variable equal to 1 for observations in which the standard deviation of net income divided by book value was in the top quintile;
High PATTN = a dummy variable equal to 1 for observations in which the level of attention in the 24 hours before the announcement was in the top quintile;
High Attention Post = a dummy variable equal to 1 for observations in which the level of attention in the 24 hours after the announcement was in the top quintile;
PRC = the natural logarithm of the market price of a share of common stock two days before the earnings announcement;
ARET = the absolute value of returns in the three trading days centered on the earnings announcement; and
MKTVOL = the turnover for all firms in the CRSP dataset in the three days centered on earnings.

TABLE 8

Cross-Sectional Analyses of Disagreement and Trading Volume around Earnings Announcements

	(1) Q1 (lowest)	(2) Q2	(3) Q3	(4) Q4	(5) Q5 (highest)
Panel A: Absolute Earnings Surprise					
PDIS	0.034***	0.046***	0.042***	0.042***	0.061***
	(9.28)	(10.70)	(11.12)	(10.93)	(14.08)
Δ*DIS*	0.026***	0.033***	0.031***	0.032***	0.046***
	(9.67)	(11.52)	(10.45)	(11.08)	(13.64)
Observations	4,116	4,115	4,115	4,115	4,116
Adj. R^2	0.55	0.56	0.54	0.55	0.56
Panel B: Signed Earnings Surprise					
PDIS	0.044***	0.040***	0.041***	0.041***	0.060***
	(11.45)	(11.43)	(9.47)	(10.78)	(13.92)
Δ*DIS*	0.037***	0.028***	0.032***	0.031***	0.043***
	(12.26)	(11.40)	(10.52)	(10.64)	(13.09)
Observations	4,116	4,115	4,115	4,115	4,116
Adj. R^2	0.53	0.54	0.55	0.57	0.55
Panel C: Market Value of Equity					
PDIS	0.067***	0.072***	0.066***	0.046***	0.016***
	(15.35)	(14.03)	(13.95)	(11.26)	(5.44)
Δ*DIS*	0.048***	0.052***	0.046***	0.031***	0.012***
	(15.01)	(14.20)	(13.24)	(10.98)	(6.37)
Observations	4,116	4,115	4,115	4,115	4,116
Adj. R^2	0.53	0.60	0.63	0.60	0.52
Panel D: Analyst Following					
PDIS	0.048***	0.052***	0.055***	0.038***	0.036***
	(11.29)	(12.49)	(11.90)	(8.47)	(6.63)
Δ*DIS*	0.031***	0.040***	0.042***	0.030***	0.026***
	(11.15)	(12.98)	(12.54)	(9.35)	(6.73)
Observations	4,116	4,115	4,115	4,115	4,116
Adj. R^2	0.49	0.55	0.58	0.60	0.59
Panel E: Institutional Ownership					
PDIS	0.046***	0.046***	0.038***	0.056***	0.063***
	(9.76)	(10.40)	(8.85)	(12.74)	(13.60)
Δ*DIS*	0.031***	0.033***	0.027***	0.039***	0.052***
	(10.63)	(10.20)	(9.22)	(11.98)	(14.37)
Observations	4,116	4,115	4,115	4,115	4,116
Adj. R^2	0.52	0.49	0.56	0.60	0.62
Panel F: Prior Earnings Volatility					
PDIS	0.049***	0.030***	0.041***	0.039***	0.064***
	(12.34)	(8.98)	(11.97)	(8.37)	(11.96)
Δ*DIS*	0.037***	0.025***	0.031***	0.029***	0.047***
	(12.41)	(10.89)	(11.72)	(9.20)	(11.66)
Observations	4,116	4,115	4,115	4,115	4,116
Adj. R^2	0.52	0.55	0.55	0.55	0.57

(continued on next page)

TABLE 8 (continued)

Panel G: RavenPack Event Volume (−10, −2)					
PDIS	0.040***	0.053***	0.044***	0.045***	0.024***
	(11.98)	(14.85)	(12.23)	(10.79)	(4.81)
Δ*DIS*	0.033***	0.038***	0.032***	0.028***	0.020***
	(13.11)	(13.96)	(11.76)	(9.75)	(5.84)
Observations	3,679	3,678	3,680	3,681	3,682
Adj. R^2	0.51	0.53	0.57	0.58	0.52
Panel H: Number of Earnings Announcements on the Same Day					
PDIS	0.050***	0.050***	0.049***	0.045***	0.038***
	(9.64)	(11.98)	(11.40)	(11.83)	(10.75)
Δ*DIS*	0.037***	0.037***	0.035***	0.036***	0.027***
	(10.45)	(11.79)	(10.98)	(12.82)	(10.44)
Observations	4,116	4,115	4,115	4,115	4,116
Adj. R^2	0.55	0.55	0.56	0.52	0.50
Panel I: Number of Active Users in 24 Hours before Earnings					
PDIS	0.030***	0.029***	0.033***	0.032***	0.068***
	(8.90)	(9.54)	(9.51)	(8.70)	(9.27)
Δ*DIS*	0.029***	0.024***	0.027***	0.028***	0.047***
	(10.29)	(11.08)	(11.95)	(9.60)	(8.58)
Observations	4,116	4,115	4,115	4,115	4,116
Adj. R^2	0.48	0.48	0.52	0.54	0.61
Panel J: Change in Number of Users in the 48 Hours around Earnings					
PDIS	0.025***	0.026***	0.028***	0.039***	0.066***
	(8.87)	(9.04)	(8.27)	(9.74)	(10.21)
Δ*DIS*	0.012***	0.016***	0.017***	0.026***	0.053***
	(6.35)	(7.59)	(7.18)	(8.26)	(9.60)
Observations	4,116	4,115	4,115	4,115	4,116
Adj. R^2	0.54	0.49	0.50	0.56	0.63

***, **, * Denote significance at the less than 10 percent, 5 percent, and 1 percent levels, respectively.
The t-statistics reported in parentheses are based on standard errors clustered by firm: industry, year, and month fixed effects are included.

Variable Definitions:
ADJTURN = median-adjusted turnover calculated as the three-day earnings turnover less the median three-day turnover in the previous 249 trading days;
PDIS = the standard deviation of user sentiment in the 24 hours before earnings are announced;
ΔDIS = the standard deviation of sentiment in the 24 hours after earnings less the standard deviation of sentiment in the 24 hours before earnings;
PRC = the natural logarithm of the market price of a share of common stock two days before the earnings announcement;
ARET = the absolute value of returns in the three trading days centered on the earnings announcement; and
MKTVOL is the turnover for all firms in the CRSP dataset in the three days centered on earnings.

We consider variation in existing information, or information environment, by examining variation in firm size (Table 8, Panel C), analyst following (Table 8, Panel D), institutional holdings (Table 8, Panel E), prior earnings volatility (Table 8, Panel F), and media coverage (Table 8, Panel G). Based on prior studies, we expect that earnings announcements are more informative when the level of existing information, or the information environment, is lower. Our results generally support our expectation, with the coefficients on *PDIS* and Δ*DIS* generally being significantly lower for higher levels of existing information.

Finally, we consider variation in attention to the earnings announcement by examining variation in the number of firms reporting earnings on the same day (Table 8, Panel H) and user attention (Table 8, Panels I and J). Based on the disclosure processing costs literature, we expect that the association between disagreement and trading volume around earnings announcements will be lower when attention is lower. We find evidence consistent with our expectations. For example, in Table 8, Panel H, on the days with the lowest number of earnings announcements, or least busy earnings days (column (1)), the coefficients on *PDIS* ($\beta = 0.050$, t = 9.64) and Δ*DIS* ($\beta = 0.037$, t = 10.45), whereas in column

TABLE 9

Random Sample of Nonearnings Announcement Days

$$ADJTURN_j = \alpha + \beta_1 PDIS_j + \beta_2 \Delta DIS_j + \beta_3 PATTN_j + \beta_4 \Delta ATTN_j + \beta_5 (PATTN_j \times PDIS_j) + \beta_6 (PATTN_j \times \Delta DIS_j) + \beta_7 (\Delta ATTN_j \times \Delta DIS_j) + \sum Controls_j + \sum FE + e_j$$

	(1)	(2)	(3)
PDIS	0.017***	0.002	−0.005
	(8.69)	(1.49)	(−1.25)
ΔDIS	0.008***	0.002**	0.004***
	(7.70)	(2.30)	(2.86)
PATTN		0.022***	0.014***
		(15.15)	(5.11)
$\Delta ATTN$		0.006***	0.007***
		(8.96)	(4.43)
$PDIS \times PATTN$			0.022***
			(2.74)
$PDIS \times \Delta ATTN$			−0.003
			(−0.85)
$\Delta DIS \times \Delta ATTN$			−0.005**
			(−2.07)
PRC	−0.002***	−0.002***	−0.002***
	(−6.24)	(−6.76)	(−6.88)
ARET	0.494***	0.462***	0.462***
	(17.71)	(17.40)	(17.39)
MKTVOL	0.499***	0.391***	0.399***
	(6.65)	(5.36)	(5.42)
Observations	50,000	50,000	50,000
Adj. R^2	0.42	0.46	0.46

***, **, * Denote significance at the less than 10 percent, 5 percent, and 1 percent levels, respectively.
The t-statistics reported in parentheses are based on standard errors clustered by firm.

Variable Definitions:
ADJTURN = median-adjusted turnover calculated as the three-day earnings turnover less the median three-day turnover in the previous 249 trading days;
PDIS = the standard deviation of user sentiment in the 24 hours before a randomly selected nonearnings day;
ΔDIS = the standard deviation of sentiment in the 24 hours after a random nonearnings day less the standard deviation of sentiment in the 24 hours before the same randomly selected nonearnings day;
PATTN = the ranked percentage of the number of users in the 24 hours before a randomly selected nonearnings day less the average daily number of users in the −10 to −30 days before the randomly selected nonearnings day; and
$\Delta ATTN$ = the ranked percentage of the number of users in the 24 hours after the randomly selected nonearnings day less the number of users in the 24 hours after the randomly selected nonearnings day.

We also report control variables:
PRC = the natural logarithm of the market price of a share of common stock two days before the earnings announcement;
ARET = the absolute value of returns in the three trading days centered on the randomly selected nonearnings day; and
MKTVOL = the turnover for all firms in the CRSP dataset in the three days centered on the firm's nonearnings-announcement day.

(5), the busiest earnings days, the coefficients on *PDIS* ($\beta = 0.038$, t = 10.75) and ΔDIS ($\beta = 0.027$, t = 10.44) are significantly smaller.

Placebo Tests

Cookson and Niessner (2020) find that attention has a minimal impact on disagreement using a sample of largely nonearnings announcement days. To reconcile with their findings, we examine a random sample of 50,000 nonearnings announcement days. As these days do not include significant disclosures that require investor processing, we do not expect to find evidence of a positive interactive effect between disagreement and changes in attention. In Table 9, column (1) we start with Equation (8) and find that the coefficients on both $PDIS$ ($\beta = 0.017$, t = 8.69) and ΔDIS ($\beta = 0.08$, t = 7.70) are positive and significant. In column (2), we include controls for processing costs around our non-event days and find that $PDIS$ becomes insignificant ($\beta = 0.002$, t = 1.49), while both $PATTN$ ($\beta = 0.022$, t = 15.155) and $\Delta ATTN$ ($\beta = 0.006$, t = 8.96) are positive and significant. We include interactions in column (3). We find that the coefficient on $PDIS \times \Delta ATTN$ is insignificant, and the coefficient on $\Delta DIS \times \Delta ATTN$ is negative and significant at the $p < 0.05$ level ($\beta = -0.002$, t = -2.07). These results suggest that the interactive effects of attention we observe around earnings announcements are not seen on regular trading days, consistent with the disclosure processing costs literature, as a sample of earnings announcements focuses on days that systematically include disclosures that require processing.

VI. CONCLUSION

We examine the association between disagreement and trading volume around earnings announcements using a social media-based measure of disagreement. Our measure has the benefit of incorporating the opinions of individual investors about firm value directly from their posts on the social media network StockTwits. This setting is beneficial as it provides a direct source of information about a broad set of individuals' opinions that are observed at a high frequency, along with other information on the characteristics of the individuals posting. In this social media setting, we find support for the two reasons investor disagreement is expected to increase trading volume around earnings announcements: (1) because individual investors hold different preannouncement beliefs, and (2) due to different interpretations of the earnings news. We next investigate predictions that stem from recent theories that assume individual investors face processing costs when new information arrives. Using measures of investor attention and investor heterogeneity, we find evidence consistent with lower disclosure processing costs amplifying both forms of disagreement.

Our measure is distinct from prior analyst-based measures of disagreement that focus on dispersion in earnings forecasts but does not subsume it in our regression analysis. That both social media measures and analyst forecast measures of disagreement contribute to the explanation of trading volume around earnings announcements suggests that disagreement likely incorporates both disagreement about short-term earnings outcomes as well as disagreements about other aspects of firm value and nonfundamental information such as conflicting signals from technical analysis-based trading strategies.

Our results must be interpreted with the caveat in mind that our measure of investor disagreement is based on the opinions of individuals on the social media website StockTwits and may not be generalizable to disagreement between all stock market participants. Specifically, as StockTwits caters to individual investors with varying levels of financial sophistication, our measure incorporates a diverse set of investor opinions but does not explicitly capture the opinions of institutional traders expected to influence capital market outcomes such as price.

Finally, our results have implications for future research that investigates theories on individual trading behaviors. Processing costs likely play an important role in explaining variation in other theoretical constructs relating to the trading behaviors of individuals and market outcomes arising from these behaviors. We leave it to future research to examine whether these traders simply provide liquidity to the market or can influence other market outcomes, such as price.

REFERENCES

Ahmed, A., S. R. A. Schneible, Jr., 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. https://doi.org/10.1506/N2XD-TF8Y-JT4L-L6V0

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. https://doi.org/10.1111/j.1540-6261.2004.00662.x

Bagnoli, M., M. D. Beneish, and S. G. Watts. 1999. Whisper forecasts of quarterly earnings per share. *Journal of Accounting and Economics* 28 (1): 27–50. https://doi.org/10.1016/S0165-4101(99)00018-X

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. https://doi.org/10.1111/j.1911-3846.2010.01061.x

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.

Barber, B. M., and T. Odean. 2008. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies* 21 (2): 785–818. https://doi.org/10.1093/rfs/hhm079

Barron, O. E., and J. M. Karpoff. 2004. Information precision, transaction costs, and trading volume. *Journal of Banking and Finance* 28 (6): 1207–1223. https://doi.org/10.1016/S0378-4266(03)00115-8

Barron, O. E., R. A. Schneible, Jr., and D. E. Stevens. 2018. The changing behavior of trading volume reactions to earnings announcements: Evidence of the increasing use of accounting earnings news by investors. *Contemporary Accounting Research* 35 (4): 1651–1674. https://doi.org/10.1111/1911-3846.12371

Bartov, E., L. Faurel, and P. S. Mohanram. 2018. Can Twitter help predict firm-level earnings and stock returns? *The Accounting Review* 93 (3): 25–57. https://doi.org/10.2308/accr-51865

Blankespoor, E., E. deHaan, and I. Marinovic. 2020. Disclosure processing costs, investors' information choice, and equity market outcomes: A review. *Journal of Accounting and Economics* 70 (2–3): 101344–101346. https://doi.org/10.1016/j.jacceco.2020.101344

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. https://doi.org/10.1007/s11142-017-9422-2

Bollerslev, T., J. Li, and Y. Xue. 2018. Volume, volatility, and public news announcements. *The Review of Economic Studies* 85 (4): 2005–2041. https://doi.org/10.1093/restud/rdy003

Carlin, B. I., F. A. Longstaff, and K. Matoba. 2014. Disagreement and asset prices. *Journal of Financial Economics* 114 (2): 226–238. https://doi.org/10.1016/j.jfineco.2014.06.007

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. https://doi.org/10.1093/rfs/hhu001

Cho, M., and S. Y. Kwon. 2014. Trading volume and investor disagreement around management forecast disclosures. *Journal of Accounting, Auditing and Finance* 29 (1): 3–30. https://doi.org/10.1177/0148558X13516977

Cookson, J. A., and M. Niessner. 2020. Why don't we agree? Evidence from a social network of investors. *The Journal of Finance* 75 (1): 173–228. https://doi.org/10.1111/jofi.12852

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, U.K.: Optirisk Systems Ltd.

Da, Z. H. I., J. Engelberg, and P. Gao. 2011. In search of attention. *The Journal of Finance* 66 (5): 1461–1499. https://doi.org/10.1111/j.1540-6261.2011.01679.x

Giannini, R., P. Irvine, and T. Shu. 2019. The convergence and divergence of investors' opinions around earnings news: Evidence from a social network. *Journal of Financial Markets* 42: 94–120. https://doi.org/10.1016/j.finmar.2018.12.003

Guay, W., D. Samuels, and D. Taylor. 2016. Guiding through the fog: Financial statement complexity and voluntary disclosure. *Journal of Accounting and Economics* 62 (2–3): 234–269. https://doi.org/10.1016/j.jacceco.2016.09.001

Hirshleifer, D. 2015. Behavioral finance. *Annual Review of Financial Economics* 7 (1): 133–159. https://doi.org/10.1146/annurev-financial-092214-043752

Hirshleifer, D., and S. H. Teoh. 2003. Limited attention, information disclosure, and financial reporting. *Journal of Accounting and Economics* 36 (1–3): 337–386. https://doi.org/10.1016/j.jacceco.2003.10.002

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. https://doi.org/10.1086/262005

Kim, O., and R. E. Verrecchia. 1991. Market reaction to anticipated announcements. *Journal of Financial Economics* 30 (2): 273–309. https://doi.org/10.1016/0304-405X(91)90033-G

Kim, O., and R. E. Verrecchia. 1997. Pre-announcement and event-period private information. *Journal of Accounting and Economics* 24 (3): 395–419. http://dx.doi.org/10.1016/S0165-4101(98)00013-5

Le, Q., and T. Mikolov. 2014. *Distributed representations of sentences and documents*. Proceedings of the 31st International Conference on Machine Learning, Beijing, China, June 21–26.

Li, F. 2010. Textual analysis of corporate disclosures: A survey of the literature. *Journal of Accounting Literature* 29: 143–165.

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. https://doi.org/10.1111/j.1540-6261.2010.01625.x

Odean, T. 1998. Volume, volatility, price, and profit when all traders are above average. *The Journal of Finance* 53 (6): 1887–1934. https://doi.org/10.1111/0022-1082.00078

Shannon, C. E. 1948. A mathematical theory of communication. *Bell System Technical Journal* 27 (3): 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

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. https://doi.org/10.1016/0167-2681(89)90076-0

Veldkamp, L. L. 2011. *Information Choice in Macroeconomics and Finance*. Princeton NJ: Princeton University Press.

Yang, Z., Z. Dai, Y. Yang, J. Carbonell, R. Salakhutdinov, and Q. V. Le. 2019. *Xlnet: Generalized autoregressive pretraining for language understanding.* Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

APPENDIX A

Variable Definitions

Variable	Definition	Database
ADJTURN	The median-adjusted turnover calculated as the three-day earnings turnover less the median; three-day turnover in the previous 249 trading days.	CRSP
PDIS	The standard deviation of user sentiment in the 24 hours before earnings are announced.	StockTwits
ΔDIS	The standard deviation of sentiment in the 24 hours after earnings less the standard deviation of sentiment in the 24 hours before earnings.	StockTwits
PATTN (log)	The natural log of the number of active users in the 24 hours before earnings are announced less the average daily attention in days −10 to −30 relative to the announcement.	StockTwits
ΔATTN (log)	The natural log of the number of active users in the 24 hours after earnings less the number of active users in the 24 hours after earnings are announced.	StockTwits
PATTN (rank)	The within-sample percentage rank of *PATTN (log)*.	StockTwits
ΔATTN (rank)	The within-sample percentage rank of *ΔATTN (log)*.	StockTwits
PRC	The natural logarithm of the market price of a share of common stock two days before the earnings announcement.	CRSP
ARET	The absolute value of returns in the three trading days centered on the earnings announcement.	CRSP
MKTVOL	The turnover for all firms in the CRSP dataset in the three days centered on earnings.	CRSP
Heterogeneity	The sum of Shannon's Entropies of each user characteristic in the given timespan.	StockTwits
Engagement	Attention from one of the top 1,000 most central users (constructed by year).	StockTwits
Followers	Attention from one of the top 1,000 most followed users (constructed by year).	StockTwits
High News Pre	Dummy variable equal to 1 for observations for which the number of news articles in the 24 hours before earnings was in the top quintile.	RavenPack
High News Post	Dummy variable equal to 1 for observations for which the number of news articles in the 24 hours after earnings was in the top quintile.	RavenPack
High NI Dispersion	Dummy variable equal to 1 for observations in which the standard deviation of net income divided by book value was in the top quintile.	Compustat
High Attention Pre	Dummy variable equal to 1 for observations in which the level of attention in the 24 hours before the announcement was in the top quintile.	StockTwits
High Attention Post	Dummy variable equal to 1 for observations in which the level of attention in the 24 hours after the announcement was in the top quintile.	StockTwits
Absolute Earnings Surprise	The absolute value of the actual earnings per share less the median analyst estimate scaled by price two days before earnings are announced.	I/B/E/S
Signed Earnings Surprise	The actual earnings per share less the median analyst estimate scaled by price two days before earnings are announced.	I/B/E/S
Market Value of Equity	The price per share of common stock multiplied by the number of shares outstanding.	CRSP

(continued on next page)

APPENDIX A (continued)

Variable	Definition	Database
Analyst Following	The number of analysts contributing to the consensus estimate.	I/B/E/S
Institutional Ownership	The percentage of shares owned by institutional investors that report with a Form 13 F.	Thomson Reuters 13 F, CRSP, Compustat
Prior Earnings Volatility	The variance of net income divided by book value for the previous four years.	Compustat
RavenPack Event Volume (−10,−2)	The count of news sources mentioning a firm in the −10 to −2 days relative to earnings.	RavenPack
Number of Earnings Announcements on the Same Day	The number of earnings announcements in the full sample of announcements covered by the IBES database.	I/B/E/S
Number of Active Users in 24 Hours Prior to Earnings	The number of users posting in the 24 hours before earnings are announced.	StockTwits
Change in Number of Users in the 48 Hours around Earnings	The number of users posting in the 24 hours after earnings are announced less the number of users posting in the 24 hours before earnings are announced.	StockTwits