

Wisdom of Crowds: Cross-sectional Variation in the Informativeness of Third-Party-Generated Product Information on Twitter

By

Vicki Wei Tang
McDonough School of Business
Georgetown University
Email: wt29@georgetown.edu

July 2017

Accepted by Philip Berger. I would like to thank an anonymous referee for her/his insightful comments and suggestions. This study benefits from the discussions with seminar participants at Georgetown University, the University of Hong Kong, and the U.S. Securities and Exchange Commission. This study also benefits from comments from William Baber, John Core, Rohan Williamson, Prem Jain, Lee Pinkowitz, Allan Eberhart, James Sinclair, Jennie Bai, and Jess Cornaggia. I would also like to thank the Center for Financial Markets and Policy at McDonough School of Business for its financial support. Likefolio.com provides the data on Twitter comments.

Abstract

This paper examines whether third-party-generated product information on Twitter, once aggregated at the firm level, is predictive of firm-level sales, and if so, what factors determine the cross-sectional variation in the predictive power. First, the predictive power of Twitter comments increases with the extent to which they fairly represent the broad customer response to products and brands. The predictive power is greater for firms whose major customers are consumers rather than businesses. Second, the word-of-mouth effect of Twitter comments is greater when advertising is limited. Third, a detailed analysis of the identity of the tweet handles provides the additional insights that the predictive power of the volume of Twitter comments is dominated by “the wisdom of crowds,” whereas the predictive power of the valence of Twitter comments is largely attributable to expert comments. Furthermore, Twitter comments not only reflect upcoming sales, but also capture an unexpected component of sales growth.

Keywords: wisdom of crowds; social media; product information; word of mouth; Twitter; fundamental analysis

JEL codes: D83; G14; M41; O33

I. Introduction

One major shortcoming of the current corporate financial reporting regulatory regime is that it does not require adequate disclosure by listed firms of nonfinancial information, such as product information and customer satisfaction, that would help investors and creditors make informed decisions (Amir and Lev [1996]). This paper examines whether third-party-generated comments about products and brands on Twitter, once aggregated at the firm level, provide information that is useful in forecasting firm-level fundamentals. This study further explores what factors determine the cross-sectional variation in the predictive power of product information on Twitter. Conceptually, the research question extends the scope of the investigation from the *average* predictive power in existing “now-casting” studies to the *cross-sectional variation* in the predictive power of online information. This study differentiates itself from most studies that extract information about sentiment from electronic platforms by focusing on firm fundamentals rather than stock prices. Furthermore, this study distinguishes itself from earlier studies on social media by focusing on third-party-generated information rather than company-initiated information.

This study chooses Twitter as the setting in examining the cross-sectional variation in the information content primarily because of the level of aggregation of product information. Though information about products and brands is available at the *product* level from alternative sources,¹ the assignment of various products and brands to the businesses that own them imposes a significant empirical challenge. The data provider uses its proprietary information to achieve a reliable mapping between products

¹ Examples of alternative online sources include Google, Amazon, and Yelp.

and brands and the entities that own them and, therefore, is able to aggregate Twitter comments about products and brands at the firm level. The aggregation of product information at the firm level provides a significant empirical edge over other settings because the natural unit for fundamental analysis is the firm. Twitter is also one of the two social media platforms the Securities and Exchange Commission allows companies to use to communicate with investors. Accordingly, using Twitter as the setting has the additional benefit of juxtaposing third-party-generated product information with company-initiated disclosure on the same platform.

This study examines the cross-sectional variation in the predictive power of product information on Twitter with respect to firm-level accounting fundamentals. Accordingly, the target of Twitter comments is limited to products and brands, and tweet handles (the holders of tweet comments) are limited to third parties, excluding the company itself. The selected Twitter comments are then aggregated at the firm level using the data provider's proprietary knowledge in mapping from products and brands to the entities that own them. Two statistics are used to summarize the volume and valence of Twitter comments about products and brands. The first statistic (PURCHASE) is defined as the total number of tweets that mention an actual purchase of a product or brand in the past or a forward-looking intent to purchase. PURCHASE maps Twitter comments directly into a recent past sale or a potential sale in the future. The valence of each tweet is classified as positive, negative, or neutral. The second statistic (POSITIVE) is defined as the ratio of the number of tweets that convey a positive assessment of products and brands over the number of tweets that convey a non-neutral (either positive or negative) assessment of products and brands. POSITIVE summarizes the collective

customer satisfaction or dissatisfaction with a company's products and brands.

Product information on Twitter could reflect firm-level sales through a combination of two effects. First, the two statistics summarize Twitter users' responses to products or brands and, therefore, provide easily accessible signals of the broad customer response. This is labeled as the pure *signal* effect of Twitter comments. Second, Twitter comments could spur additional sales through a *word-of-mouth* effect.

From a pure signal perspective, the ability of Twitter comments to reflect firm-level sales depends on whether those tweets are representative of the broad customer response to the company's products and brands. As Twitter is largely a social platform for leisure rather than business activities, individual consumers are more likely to share their product experiences on Twitter than are businesses. Accordingly, Twitter comments are more representative of the broad customer response for companies whose major customers are consumers. Therefore, the predictive power of Twitter comments is expected to be greater among those companies. Empirically, the predictive power of PURCHASE with respect to upcoming sales is more pronounced for companies whose major customers are consumers than otherwise. The second summary statistic, POSITIVE, by construction, factors in only non-neutral (either positive or negative) tweets. To the extent that only extremely satisfied (dissatisfied) customers initiate positive (negative) comments, POSITIVE is susceptible to a higher level of extremity bias. Not surprisingly, the predictive power of POSITIVE with respect to upcoming sales is rather limited.

From the word-of-mouth perspective, the ability of Twitter comments to spur additional sales varies with advertising. Advertising targets a wide audience and seeks to

increase sales by increasing brand awareness. The ability of Twitter comments to spur more sales works through a mechanism similar to that of advertising: A high volume of tweets increases brand awareness through the connected network on Twitter. To the extent that consumer-generated brand awareness on Twitter *substitutes* for producer-generated brand awareness of advertising, the ability of PURCHASE to predict upcoming sales is more pronounced when advertising is limited.

The ability of Twitter comments to spur additional sales also varies with the identity of tweet handles. Tweets generated by various types of tweet handle are perceived with varying degrees of credibility and influence. Empirically, on a stand-alone basis, PURCHASE initiated by product experts, PURCHASE initiated by the media, or PURCHASE initiated by the crowd is predictive of upcoming sales. However, when all three categories of tweet handle are examined jointly, the predictive power of PURCHASE is dominated by those initiated by the crowd, which is consistent with the notion of the wisdom of crowds.

More significantly, this study finds that Twitter comments about products and brands are not only reflective of what we know about upcoming sales through other information sources, but are also predictive of unexpected sales growth. Unexpected sales are defined as realized sales relative to analyst forecasts. On a stand-alone basis, PURCHASE initiated by the crowd or PURCHASE initiated by the media predicts unexpected sales growth. In contrast, the predictive power of POSITIVE with respect to unexpected sales is largely attributable to comments initiated by product experts. The last set of results suggests that product information on Twitter is incrementally informative and that Twitter comments capture an unexpected component of sales growth.

This study contributes to both the academic literature and the needs of practitioners. First, to my knowledge, this is the first study to examine the information content of third-party-generated voluntary disclosure on Twitter with respect to firm fundamentals and the determinants of the cross-sectional variation in its predictive power. Social media enables its users, including customers of products and brands, to disseminate their opinions and recommendations in a manner that is not possible through traditional information outlets (Miller and Skinner [2015]). In contrast to Sunstein [2008], who argues that the blogosphere cannot serve as a marketplace for information, this study finds that product information on Twitter, once summarized at the firm level, is incrementally informative about fundamentals, especially for firms whose major customers are consumers and when advertising is limited.

Second, this is the first study to explore the cross-sectional variation in the predictive power of Twitter comments based on the identity of tweet handles. More interestingly, the predictive power of Twitter comments with respect to upcoming sales is dominated by comments initiated by the crowd rather than those initiated by the media or product experts. The finding provides empirical support for the notion of the wisdom of crowds.

Third, the finding that product information on Twitter captures an unexpected component of sales growth relative to that provided by professionals in the capital markets, such as analysts, is economically important because it works to the advantage of individual investors. An alleged market friction is that some informative signals are inaccessible to individual investors because of their high cost. Anecdotal evidence suggests that hedge funds spend huge amounts of money on satellite photos of parking

lots of retailers to get an early indicator of upcoming sales. If an equally effective leading indicator is obtainable by individual investors from Twitter at virtually no cost, social media levels the playing field between institutional and individual investors. Furthermore, third-party-generated disclosure on social media improves a firm's overall information environment by providing incrementally informative product information over and above that provided by the company itself.

Section II discusses the institutional background and develops hypotheses. Section III presents the data and discusses the research design. Section IV presents the empirical results. Section V concludes the study.

II. Related literature, background, and hypothesis development

2.1 Related literature

This study broadly falls into the emerging now-casting literature. Now-casting has recently become popular in economics because standard measures used to assess the state of an economy, such as gross domestic product, are determined only after a long delay. Earlier now-casting studies examine whether information from web searches can, *on average*, predict contemporaneous economy-level or industry-level economic variables. Goel et al. [2010] provide a detailed review on studies that use web search data to predict contemporaneous information at the economy level. Ettredge et al. [2005] is the first study that suggests the usefulness of web search data in forecasting the U.S. unemployment rate, and Huang and Penna [2009] examine the usefulness of web search data for measuring economy-wide consumer sentiment. McLaren and Shanbhoge [2011] summarize how web search data are used for economic forecasting by central banks. A

few other studies examine how the volume of web searches predicts contemporaneous demand at the *industry* level. For instance, Choi and Varian [2012] find that the volume of Google queries is helpful in forecasting contemporaneous sales in automobiles and tourism, whereas Vosen and Schmidt [2011] and Wu and Brynjolfsson [2013] find its usefulness in forecasting contemporaneous sales in the retail and housing sectors, respectively.

More recent studies, by both academics and practitioners, use information from electronic markets and social media to generate *political* predictions. For instance, Berg et al. [1997] use data from the Iowa Electronic Markets to study factors associated with the ability of the markets to predict the outcome of political elections. Gayo-Avello [2013] provides a comprehensive review of electoral predictions based on Twitter data and concludes that Twitter's presumed predictive power in this regard has been somewhat exaggerated. In practice, using an analysis of Twitter activity, Tweetcast incorrectly predicted the outcome of the 2016 U.S. presidential election.

A number of prior studies have extracted information about sentiment from electronic platforms and examined the predictive power of information on a diverse set of electronic platforms, including Internet message boards, investing websites and social platforms, and Google, with respect to stock prices. For instance, Tumarkin and Whitelaw [2001], Antweiler and Frank [2004], and Das and Chen [2007] examine whether conversations on Internet message boards are associated with stock returns and find either statistically insignificant results or economically marginal effects. Hirschey et al. [2000] find that buy-sell stock recommendations posted on the website Motley Fool generate abnormal stock returns. Da et al. [2011a] find that the volume of Google queries

on ticker symbols measures the attention of retail investors and an increase in search volume predicts higher stock prices in the next two weeks and an eventual reversal within the year. Drake et al. [2012] use the volume of Google queries on ticker symbols as a proxy for investors' demand for financial information and find that, when investors search for more information in the days just prior to an earnings announcement, price and volume changes in the preannouncement period reflect more of the upcoming earnings news; thus, there is less price and volume response at the announcement.

In contrast, this study distinguishes itself from most studies that extract information about sentiment from electronic platforms by focusing on firm-level *accounting fundamentals* rather than stock prices. With respect to predicting firm-level fundamentals, this study is perhaps most closely related to both Da et al. [2011b] and Chen et al. [2014] but provides important insights by extending the scope of the investigation from the average predictive power to the *cross-sectional variation* in the predictive power. In their working paper entitled "In Search of Fundamentals," Da et al. [2011b] examine whether an increase in the search volume of a firm's most popular product predicts positive earnings surprises at the firm level. This study differs from Da et al. [2011b] in two major dimensions. First, on a conceptual level, the volume of Google queries largely captures information demand, whereas product information on Twitter captures information supply, which is defined jointly by volume and valence. Second, the identity of information suppliers is available on Twitter, and identification enables a more in-depth analysis based on the type of tweet handle. This analysis yields additional insights on the cross-sectional variation in the predictive power of Twitter, which is not feasible through Google. Chen et al. [2014] summarize *investors'* opinions

about stocks from the investment-related website, Seeking Alpha, and examine whether the proportion of negative words is predictive of future stock returns and earnings surprises. Unlike investors' opinions about companies in Chen et al. [2014], customers' opinions summarized in this study do not speak to the company's cost structures and other important corporate decisions, and, therefore, this study focuses on sales and sales surprises, rather than earnings surprises, as the predicted accounting variable.

This study also distinguishes itself from earlier social media studies by focusing on third-party-generated information rather than company-initiated information. Prior studies examine the dissemination of financial information through social media. For example, Blankespoor et al. [2014] and Jung et al. [2015] examine how companies use social media to disseminate firm-initiated disclosure and communicate with investors and the economic consequences of that use. Although those studies emphasize the dissemination of financial information through social media, this study is, to the best of my knowledge, the first that examines the information content of third-party-generated information about products and brands of a given firm on Twitter and the cross-sectional variation in its predictive power. The research question is especially important to both fundamental analysis and the disclosure literature because the existing corporate financial reporting regulatory regime does not require listed companies to disclose adequate product information that would help investors and creditors make informed decisions (Amir and Lev [1996]; Miller and Skinner [2015]).

To summarize, this study differs from prior studies on both a conceptual and an empirical level. Conceptually, this study extends the scope of the investigation from the average predictive power of online information to the cross-sectional variation in its

predictive power with respect to firm-level accounting fundamentals. Empirically, this study summarizes third-party-generated product information on Twitter and explores whether the predictive power of Twitter comments with respect to upcoming sales and sales surprises varies with firm characteristics and the identity of tweet handles.

2.2. Background and hypothesis development

The launch of websites such as Facebook (February 2004), YouTube (February 2005), Reddit (June 2005), and Twitter (March 2006), enabled people to share and view user-created content on a level previously unseen. These websites all fall under the blanket term “social media.” Twitter has an influential presence in social media with over 319 million active monthly users as of the fourth quarter of 2016. Twitter is also one of the two social media platforms the Securities and Exchange Commission allows companies to use to communicate with investors. By 2013, about 47% of Standard & Poor’s 1500 companies had used Twitter to communicate with investors (Jung et al. [2015]). The unit of information on Twitter is a tweet, a small blurb of up to 140 characters. Anyone that signs up on Twitter can write a tweet and view tweets written by other users. As the stream on Twitter is unfiltered, what a user signs up for is what he or she sees.

In particular, this study examines whether product information on Twitter is informative about firm-level *sales* and *sales surprises* incremental to existing information, and what factors determine the cross-sectional variation in the predictive power. The predictive power of third-party-generated product information on Twitter originates from two related sources. First, the two statistics summarize Twitter users’ *responses* to

products or brands and therefore provide easily accessible signals of the broad customer response to products and brands. I refer to this as the pure *signal* effect of product information on Twitter. Ceteris paribus, the more self-reported past purchase actions or indications of intent to purchase in the future on Twitter, the higher the contemporaneous sales or sales in the near future. The more satisfied existing customers are, the more likely they are to continue to purchase the particular product or brand (Ittner et al. [2003]).

Second, Twitter comments could spur or discourage more sales through a word-of-mouth effect. Word-of-mouth refers to the dissemination of information, such as opinions and recommendations, through individual-to-individual communications. Twitter provides a friendly platform for users to communicate with their connected audience. The influence of offline word-of-mouth is limited to a local social network; online word-of-mouth, as that on Twitter, reaches beyond the local network and could spread to a much wider audience in cyberspace. The two most important attributes of word-of-mouth communication are volume and valence. Prior studies in marketing find that, when the unit of analysis is at the product level, a high volume of tweets about the particular product is likely to increase the degree of consumer awareness and the number of informed consumers. For instance, Liu [2006] and Asur and Huberman [2010] find that the rates at which chats were initiated on Yahoo websites or tweets were generated were strong indicators of a movie's box office success. In addition, when the unit of analysis is at the product level, positive tweets are likely to sway customers' assessments in favor of the company's products or brands. For instance, Chevalier and Mayzlin [2006] find that the valence of online book reviews influences book sales. The word-of-

mouth effect of Twitter comments at the product level can easily be extrapolated to the firm level because the firm's total sales is the sum of the sales of all of its products and brands.²

From a pure signal perspective, the ability of product information on Twitter to reflect firm-level sales depends on whether the tweets are representative of the broad customer response to the company's products and brands. When a firm's customer base consists predominantly of consumers (a business-to-consumer firm), its representative customer is an individual consumer. When a firm's major customers are business clients (a business-to-business firm), its representative customer is a business. As Twitter is largely a social platform for leisure rather than business activity, individuals are more likely to share their product experiences on Twitter than are businesses. Accordingly, when a company's major customers are consumers, Twitter comments are more representative of the broad consumer response to the company's products and brands. Thus, *ceteris paribus*, the predictive power of product information on Twitter is expected to be greater for firms whose major customers are consumers than for those whose major customers are businesses. This leads to the first hypothesis:

H1: The predictive power of product information on Twitter with respect to firm-level sales is more pronounced for firms whose major customers are consumers.

From the perspective of the word-of-mouth effect, the ability of product information on Twitter to spur more sales varies with the promotional activities initiated

² In Da et al. (2011b), the volume of Google searches is measured only at the product level, in particular for a firm's most popular product. For some firms, the revenue source could come from hundreds of products and brands. For others, the revenue source is limited to only a few products and brands. Therefore, it is no surprise that they document a rather limited predictive power of the product-level volume of Google searches with respect to firm-level fundamentals.

by the company itself. A notable example is advertising, which seeks to increase sales by increasing brand awareness among potential customers. I refer to this mechanism as producer-generated brand awareness. The ability of the volume of Twitter comments to spur more sales works through a mechanism similar to that of advertising: A customer's tweet about a past or potential purchase of a particular product or brand increases the product or brand awareness among her or his connected parties on Twitter. I refer to this mechanism as customer-generated brand awareness. To the extent that customer-generated brand awareness on Twitter *substitutes* for producer-generated brand awareness of advertising, the ability of PURCHASE to spur more sales is more pronounced when advertising is limited. Anecdotaly, in a case study, Reinstein and Snyder [2005] find that third-party reviews have a significant effect on opening weekend box office revenue for narrowly released movies, but not for widely released movies. In summary, the substitution between the producer-generated brand awareness and the customer-generated brand awareness implies that the predictive power of PURCHASE with respect to upcoming sales decreases with the intensity of advertising activities. This leads to the second hypothesis:

H2: Ceteris paribus, the predictive power of PURCHASE with respect to upcoming sales decreases with the level of advertising.

Furthermore, the ability of product information on Twitter to spur more sales could also be related to the concept of social proof. People will conform to the actions of others under the assumption that those actions reflect proper behavior. Social proof can be broadly categorized into expert social proof, media social proof, and wisdom-of-crowds social proof. Expert social proof suggests that a consumer trusts the

professionalism of expert opinion and is likely to purchase a product recommended by experts. Media social proof suggests that a consumer trusts the independence and objectivity of the media and is likely to purchase a product recommended by the media. The wisdom-of-crowds social proof is best put by Kirby [2000] in the *San Francisco Chronicle*: A consumer “may not trust just one non-expert, ... but if 9 out of 10 non-experts agree, it’s probably worth buying.”

Accordingly, this study categorizes product tweets into three groups based on the identity of tweet handles: The first group consists of tweets initiated by the media, the second group consists of tweets initiated by product experts, and the last group consists of tweets initiated by the crowd, that is, all tweets other than those initiated by the media and product experts. The three groups of tweets differ in many dimensions. Product tweets initiated by the crowd are, by definition, based on personal experience and often describe products in terms of their ability to match a consumer’s particular preferences and usage conditions. Accordingly, product tweets by the crowd largely capture the popularity of the product, especially the extent to which it meets the needs of various customers. On the other hand, product tweets initiated by experts are usually product-oriented and based on lab testing, and they often describe product attributes in terms of technical specifications, performance, and reliability (e.g., Chen and Xie [2008]). For instance, *Consumer Reports* and CNET.com, both of which have Twitter accounts, are among the most widely accessed sources for expert product opinion. Expert reviews typically cover every brand within a product category, and, therefore, positive product tweets generated by experts are less subject to extremity bias and are perceived to have a higher level of reliability. Furthermore, given the same content of tweets, product tweets

initiated by the media are expected to have a greater potential of spurring additional sales because of their large number of followers on Twitter. Therefore, the predictive power of Twitter comments with respect to upcoming sales is expected to vary with the identity of tweet handle. This leads to the third hypothesis:

H3: The predictive power of Twitter comments with respect to upcoming sales varies with the type of tweet handle.

III. Data, validity test, and research design

3.1. Summary statistics on product information on Twitter

I use an independent company, Likefolio, to provide the data on Twitter comments because it has proprietary information on the mapping between products and their business owners. Likefolio identifies the holder (handle) of a given tweet (that is, the person or company who initiates it) and the target (hashtag) of a given tweet (that is, the entity the tweet is discussing). Because the study is interested in the incremental information content of third-party-generated comments about a company's products and brands, the keyword or hashtag of the selected tweets is limited to products and brands and the holder or handle of the selected tweet is limited to third parties, not the company itself. Retweets are excluded from the selection. Next, Likefolio uses a combination of knowledge-based techniques and statistical methods to classify the content of each selected tweet.³ Content analysis refers to the use of natural language processing, text analysis, and computational linguistics to identify and extract subjective information in source material. The first task of content analysis is classifying whether a tweet mentions

³ An alternative method would be to use open source software tools that deploy machine learning, statistics, and natural language processing techniques to automate content analysis on those selected tweets myself. Because the volume of selected tweets in the sample periods is in the millions daily, the computing power required is beyond my capacity.

a recent past purchase of a given product or brand or an intention to do so in the future. The content of each tweet is classified as either “with mentioning of purchase” or “without mentioning of purchase” (see exhibit 1 for examples). The second task is classifying the *valence* of a given tweet—whether the opinion expressed is positive, negative, or neutral (see exhibit 2 for examples under each category).

Likefolio then uses its proprietary information to map various products and brands to the businesses that own them and summarizes selected tweets at the *company* level. The first statistic (PURCHASE) is measured as the total number of tweets that explicitly indicate a recent past purchase of a company’s products and brands or an intention to do so in the future. According to the Pew Research Center, 13% of online adults used Twitter at the beginning of 2012, and 23% of online adults used Twitter at the end of 2015. The user statistics imply that a greater proportion of the online population is using Twitter as the platform to voice their opinions over the sample period, and therefore, an increase in PURCHASE over time could capture the increased use of Twitter among customers rather than an increase in sales. Accordingly, I normalize PURCHASE by the total number of tweets circulated on Twitter (in millions) for that particular day to control for the temporal growth of Twitter while maintaining the cross-sectional variation across firms. The normalization procedure results in a variable that is analogous to market share, which captures, for every million tweets, the share that mentions purchase intent or actions for a given company’s products and brands. The second statistic (POSITIVE) is measured as the ratio of the total number of tweets that convey a positive assessment of a company’s products and brands over the total number of tweets that convey a non-neutral (positive or negative) assessment of a company’s

products and brands. The two statistics on Twitter, PURCHASE and POSITIVE, are summarized for each company on a daily basis from January 1, 2012, to December 31, 2015. Furthermore, in order to investigate the differential predictive power of product tweets initiated by different tweet handles, I summarize both PURCHASE and POSITIVE within each category of tweet handle: PURCHASE or POSITIVE by the media, PURCHASE or POSITIVE by product experts, and PURCHASE or POSITIVE by the crowd.

3.2. Sample formation and validity check on the summary statistics

To account for seasonality in quarterly sales, the dependent variable is same-quarter sales growth ($\text{SAMEQUARTER_SALESGROWTH}_{i,q}$), which is calculated as the percentage change in the quarterly sales that are disclosed after the end of the fiscal quarter q ($\text{SALES}_{i,q}$) relative to that for the same quarter in the previous fiscal year ($\text{SALES}_{i,q-4}$). As the dependent variable is measured at the firm-quarter level, explanatory variables, including statistics from Twitter, should be measured at the firm-quarter level. Therefore, I average daily values of PURCHASE and POSITIVE over the corresponding quarter.⁴ In particular, AVG_PURCHASE (AVG_POSITIVE) averages daily PURCHASE (POSITIVE) over quarter q . If PURCHASE (POSITIVE) is missing for any given day during quarter q , AVG_PURCHASE (AVG_POSITIVE) is missing.

The sample with Twitter comments consists of 30,992 firm-quarter observations that cover 1,937 unique companies. Information from either Compustat or AVG_PURCHASE is missing for 4,656 firm-quarter observations. This yields 26,336

⁴ The results are similar if the two summary statistics are averaged over the entire fiscal year and sales growth is measured at annual interval.

firm-quarter observations that cover 1,840 unique firms. For many observations, the number of non-neutral (positive and negative) tweets on a given day is zero, and therefore, a large proportion of `AVG_POSITIVE` is missing. The number of firm-quarter observations where both `AVG_PURCHASE` and `AVG_POSITIVE` are available is 10,668, which covers 1,088 unique firms.

The classification of a firm's major customer base follows a two-step approach. First, Likefolio identifies business-to-consumer firms based on their understanding of the firm's business model. Second, I use various sources to cross-examine the validity of its classification, including the business description section and the detailed disclosure of major customers in annual reports. If the information from the annual report indicates otherwise, I remove the specific firm from the business-to-consumer subsample. The two-step approach yields 166 unique business-to-consumer firms and 2,419 firm-quarter observations where both `AVG_PURCHASEi,q` and `AVG_POSITIVEi,q` are available. Business-to-consumer firms, as a group, have an economically dominant presence on Twitter: The volume of Twitter comments discussing the products and brands from the group of business-to-consumer firms combined accounts for 90% of the total volume of Twitter comments about products and brands from all sample firms. The other subsample has 8,249 firm-quarter observations that cover 922 unique business-to-business firms.

`MEDIA_AVG_PURCHASE (POSITIVE)` is the average of `PURCHASE (POSITIVE)` initiated by the media over quarter q . If `PURCHASE (POSITIVE)` initiated by the media is missing for any given day during quarter q , `MEDIA_AVG_PURCHASE (POSITIVE)` is missing. `EXPERT_AVG_PURCHASE (POSITIVE)` is the average of

PURCHASE (POSITIVE) initiated by product experts over quarter q . If PURCHASE (POSITIVE) initiated by product experts is missing for any given day during quarter q , the variable is missing. CROWD_AVG_PURCHASE (POSITIVE) is the average of PURCHASE (POSITIVE) initiated by the crowd (those tweet handles other than the media or product experts) over quarter q . If PURCHASE (POSITIVE) generated by the crowd is missing for any given day during quarter q , CROWD_AVG_PURCHASE (POSITIVE) is missing. Seven (138) firm (firm-quarter) observations are either not covered by the Institutional Brokers Estimate System (I/B/E/S) or have no information on AVG_PURCHASE by different types of tweet handle. Accordingly, the sample used to test the predictive power of Twitter comments by the identity of tweet handle includes 2,281 firm-quarter observations that cover 159 unique firms. The sample formation process is presented in table 1.

I perform a series of validity tests on the data. First, I examine whether Likefolio has tracked all products and brands for a given company. Because the mapping of products or brands to their business owners is proprietary, it is impossible to obtain the list of products and brands for each of the sample firms. However, Likefolio shares the list of products and brands for 25 randomly selected firms whose major customers are individual consumers. (See the list of products and brands for the selected companies in exhibit 3.) For example, Lulu owns two brands, whereas Volkswagen owns more than eighty products or brands. I find no omission from the list of products and brands for each of the selected companies by cross-examination with information provided by Nielsen and other online sources.

Second, I explore the cross-sectional determinants of PURCHASE and POSITIVE to examine whether the two summary statistics are correlated with firm characteristics as predicted. The first statistic ($AVG_PURCHASE_{i,q}$) summarizes the incidences of past and potential purchases of a company's products and brands as reported on Twitter. Accordingly, $AVG_PURCHASE$ is expected to be larger when the level of recent past sales is higher. I use the most recent quarter sales as the proxy for recent past sales ($SALES_{i,q-1}$). The second statistic ($AVG_POSITIVE_{i,q}$) summarizes the collective customer satisfaction or dissatisfaction with the quality of a company's products and brands. This statistic measures customer feedback after his or her experience with the product or brand. For customer feedback on products and services, the relevant information is not product information provided by the firm through advertising activities, but rather consumers' satisfaction with the product itself. As a result, $AVG_POSITIVE$ is expected to be unrelated to advertising. As advertising information is available only on annual basis from Compustat, I prorate the annual advertising expense by the proportion of sales volume each quarter in order to maintain a consistent quarterly measurement window for both the dependent variable and the explanatory variables. The prorated quarterly advertising expense is measured as the annual advertising expense multiplied by the proportion of quarterly sales over annual sales during the most recent fiscal year. I use the ratio of the prorated advertising expense over sales during the most recent fiscal quarter to proxy for the intensity of advertising activities ($ADVERTISE_{i,q-1}$).⁵ As reported in table 4, $AVG_PURCHASE$ increases with

⁵ If the intensity of advertising activities is measured as the ratio of the annual advertising expense over annual sales during the most recent fiscal year, the results are virtually identical. This is because the assumption underlying the prorating procedure is that the ratio of quarterly advertising expense to annual advertising expense is the same as the ratio of quarterly sales to

the volume of past sales, but AVG_POSITIVE does not vary with the intensity of advertising activities.

3.3. Research design on the cross-sectional variation in the predictive power of Twitter

First, I use the following specification to examine whether the predictive power of Twitter comments about products and brands with respect to sales growth is more pronounced for firms whose major customers are consumers:

$$\text{SAMEQUARTER_SALESGROWTH}_{i,q} = \alpha + \beta_1 \text{B2B}_i + \beta_2 \text{AVG_PURCHASE(POSITIVE)}_{i,q} + \beta_3 \text{B2B}_i * \text{AVG_PURCHASE(POSITIVE)}_{i,q} + \beta_4 \ln(\text{SALES}_{i,q-4}) + \beta_5 \text{SAMEQUARTER_SALESGROWTH}_{i,q-4} + \beta_6 \text{CHG_BACKLOG}_{i,q-1} + \beta_7 \text{ADVERTISE}_{i,q-1} + \varepsilon_{it} \quad \text{Model (1)}$$

The variable of interest is the interaction term between B2B and AVG_PURCHASE (AVG_POSITIVE). The indicator variable B2B is defined as 1 if a firm's major customers are businesses, and 0 if a firm's major customers are consumers. As the predictive power of Twitter comments is more pronounced for the subsample of firms whose major customers are consumers, I expect the slope coefficient on the interaction term to be negative and statistically significant. Given that the distributions of Twitter comments follow the power law distribution, I use the maximum likelihood estimation to estimate models throughout the study. Furthermore, the industry fixed effect is added in the models throughout the study when estimating the slope coefficients to mitigate the omitted-correlated-variable problem.

I include sales in the same quarter of the previous fiscal year ($\text{SALES}_{i,q-4}$) as the first control variable in order to account for mean reversion in sales growth. The lagged sales variable also captures the size of the customer base because higher past sales result

annual sales. Therefore, the ratio of advertising expense over sales is identical for all four quarters within the same fiscal year. The ratio for the quarterly interval is also identical to the ratio for the annual interval.

in a larger customer base, and potentially, a greater number of follows on Twitter. The second control variable ($\text{SAMEQUARTER_SALESGROWTH}_{i,q-4}$) captures the past trend in the same-quarter sales growth, which is calculated as the same-quarter sales growth for quarter $q-4$. The next two control variables are financial variables: the intensity of advertising activities and the change in deferred revenue (backlog). Lev and Thiagarajan [1993] find that both variables are predictive of upcoming sales. The intensity of advertising activities ($\text{ADVERTISE}_{i,q-1}$) is measured as the ratio of prorated advertising expense over sales in the previous fiscal quarter. The change in deferred revenue ($\text{CHG_BACKLOG}_{i,q-1}$) is measured as the ratio of the change in deferred revenue over sales at the beginning of the fiscal quarter.

Second, I use the following specification to examine whether the predictive power of Twitter comments varies with the intensity of advertising activities:

$$\begin{aligned} \text{SAMEQUARTER_SALESGROWTH}_{i,q} = & \alpha + \beta_1 \text{AVG_PURCHASE}_{i,q} + \beta_2 \text{AVG_PURCHASE}_{i,q} * \text{ADVERTISE}_{i,q-1} \\ & + \beta_3 \text{AVG_POSITIVE}_{i,q} + \beta_4 \text{AVG_POSITIVE}_{i,q} * \text{ADVERTISE}_{i,q-1} + \beta_5 \ln(\text{SALES}_{i,q-4}) \\ & + \beta_6 \text{SAMEQUARTER_SALESGROWTH}_{i,q-4} + \beta_7 \text{CHG_BACKLOG}_{i,q-1} + \beta_8 \text{ADVERTISE}_{i,q-1} + \varepsilon_{it} \end{aligned}$$

Model (2)

The variables of interest from model (2) are the slope coefficients on the interactions between advertising and AVG_PURCHASE . The slope coefficient on the interaction (β_2) is expected to be negative and statistically significant. In exploring whether the predictive power of Twitter comments varies with the identity of tweet handle, PURCHASE is replaced by $\text{MEDIA_AVG_PURCHASE}$, $\text{EXPERT_AVG_PURCHASE}$, and $\text{CROWD_AVG_PURCHASE}$. Similarly, POSITIVE is replaced by $\text{MEDIA_AVG_POSITIVE}$, $\text{EXPERT_AVG_POSITIVE}$, and $\text{CROWD_AVG_POSITIVE}$.

3.4. Twitter comments and unexpected sales

Unexpected sales is measured relative to analyst forecasts, which allows for a direct examination of whether Twitter comments capture sales information that is not identified by analysts. Specifically, unexpected sales is measured as realized sales relative to the consensus rather than *individual* analyst forecasts. Both consensus analyst forecasts and realized sales are collected from the summary tape of the I/B/E/S. The mean analyst forecast from the summary tape is used as a proxy for the consensus. The consensus analyst forecast is revised once a month in I/B/E/S. In calculating unexpected sales, I use the last available consensus analyst forecast for a given fiscal quarter as the market expectation. To avoid a look-ahead bias in analyst forecasts, the window that measures the volume and valence of Twitter comments starts from the first day of quarter q but ends three days prior to the reported date of the last consensus analyst forecast if the report date of the last consensus forecast on I/B/E/S is prior to the fiscal quarter-end. If a firm's last available consensus analyst forecast for quarter q is reported after the fiscal quarter-end on I/B/E/S, the window that measures the volume and valence of Twitter comments starts from the first day of quarter q and ends on the last day of quarter q . In particular, I use the following specification to examine whether Twitter comments about products and brands are predictive of the component of sales growth that is not anticipated by analysts:

$$\begin{aligned} \text{UNEXPECTED_SALESGROWTH}_{i,q} = & \alpha + \beta_1 \text{AVG_PURCHASE (AVG_POSITIVE)}_{i,q} + \beta_2 \text{Ln} \\ & (\text{NUM_FORECAST}_{i,q}) + \beta_3 \text{ACTUAL_FORECAST_DAYS}_{i,q} + \beta_4 \text{Ln} (\text{SALES}_{i,q-1}) \\ & + \beta_5 \text{SAMEQUARTER_SALESGROWTH}_{i,q-4} + \beta_6 \text{CHG_BACKLOG}_{i,q-1} + \beta_7 \text{ADVERTISE}_{i,q-1} + \varepsilon_{it} \end{aligned} \quad \text{Model (3)}$$

In model (3), the dependent variable is unexpected same-quarter sales growth ($\text{UNEXPECTED_SALESGROWTH}_{i,q}$), which is measured as the difference between

realized sales and the consensus analyst forecast divided by sales in the same quarter during the previous year ($SALES_{i,q-4}$). The variables of interest are the slope coefficient on $AVG_PURCHASE$ and that on $AVG_POSITIVE$.

The first set of control variables captures the characteristics of the consensus forecast. First, I include the natural log of the number of individual forecasts ($NUM_FORECAST_{i,q}$). Second, I include the number of calendar days between the report date of realized sales and the report date of the consensus forecast ($ACTUAL_FORECAST_DAYS_{i,q}$) because prior studies suggest that the time interval between the forecast date and the announcement date influences unexpected sales (Barron et al. [1998]). To ensure that the relationship between unexpected sales growth and Twitter comments is not subsumed by other available information, the second set of control variables includes four financial variables: (1) the natural log of sales during the most recent quarter ($SALES_{i,q-1}$); (2) the past trend in the same-quarter sales growth ($SAMEQUARTER_SALESGROWTH_{i,q-4}$); (3) advertising expense as a percentage of sales for the previous quarter ($ADVERTISE_{i,q-1}$); and (4) the percentage change in backlog at the beginning of the fiscal quarter ($CHG_BACKLOG_{i,q-1}$). Figure 1 graphs the timeline for measuring information on Twitter and analyst forecasts of upcoming sales.

IV. Empirical Results

4.1. Descriptive statistics and correlations

Table 2 provides the descriptive statistics. As shown in panel A of table 2, out of every million of tweets circulated on Twitter, the average number of tweets mentioning purchase actions or intent on a daily basis ($AVG_PURCHASE$) is 225.78 and the median is 2.92. The average ratio of the number of positive tweets over the number of non-

neutral (positive and negative) tweets on a daily basis is 88% and the median ratio is 92%. The standard deviation of AVG_PURCHASE is 2085.75. The maximum AVG_PURCHASE is 57,442 and the minimum is 0. The standard deviation of AVG_POSITIVE is 13%. The maximum AVG_POSITIVE is 100% and the minimum is 0%. The comparison of the two subsamples suggests that the volume of tweets mentioning purchase is higher, but customer assessment is less positive for business-to-consumer firms.

Panel B of table 2 provides the descriptive statistics on the two summary statistics by tweet handle for the subsample of business-to-consumer firms. First, the percentage of tweets initiated by the media is only 0.04% on average and the median is only 0.01% of all tweets. There is a significant cross-sectional variation in the percentage of tweets initiated by the media with a maximum of 3.81% and a minimum of 0%. The pattern is similar for the percentage of tweets initiated by product experts: the mean is 0.01% and the median is 0%. The two statistics combined suggest that, on average, the crowd initiated 99.95% of tweets about product and brands for the subsample. Furthermore, the media or product experts rarely mention a past purchase or an intention to purchase a particular product or brand in the future. The median AVG_PURCHASE and the median CROWD_AVG_PURCHASE are virtually identical for this particular subsample. As the majority of observations have missing value on positive or negative tweets initiated by the media or product experts, there are only 244 firm-quarter observations for which AVG_POSITIVE by all three types of tweet handle are available. Interestingly, both MEDIA_AVG_POSITIVE and EXPERT_AVG_POSITIVE have means (medians) above 95%, whereas the mean (median) CROWD_AVG_POSITIVE is 80.9% (82.1%).

Assessments from the media and product experts are more positive than those of the crowd on Twitter.

Panel A of table 3 provides the correlations between the two summary statistics on Twitter and upcoming sales growth for the entire sample. The Pearson (Spearman) correlation between AVG_PURCHASE and SAMEQUARTER_SALESGROWTH is -0.040 (-0.055) and statistically significant (p -value = 0.001). Similarly, the Spearman correlation between AVG_POSITIVE and SAMEQUARTER_SALESGROWTH is 0.026 and statistically significant (p -value = 0.006). Panel B of table 3 provides the correlations between Twitter comments and unexpected sales growth. Neither AVG_PURCHASE nor AVG_POSITIVE is statistically correlated with unexpected sales growth.

Table 4 presents the empirical results on the cross-sectional variation on the volume and valence of Twitter comments on products and brands. Consistent with the underlying constructs, AVG_PURCHASE increases with sales in the prior quarter, but AVG_POSITIVE does not vary with the intensity of advertising activities. The correlation structure validates the empirical measures for the volume and valence of Twitter comments.

4.2. Cross-sectional variation in the predictive power of Twitter comments

Table 5 presents the empirical results on whether Twitter comments are informative about upcoming sales and the cross-sectional variation in their predictive power. The benchmark specification, in which only available financial information is included to predict upcoming sales, is presented in column 1. As shown in the second

column of table 5, when AVG_PURCHASE is used as the explanatory variable, the slope coefficient is 0.066 and statistically significant (p -value = 0.001). As shown in the fourth column of table 5, the slope coefficient on the interaction between AVG_PURCHASE and B2B is -0.079 and statistically significant (p -value = 0.06). To summarize, the number of tweets with purchase action or intent is informative about upcoming sales incremental to financial information, and the predictive power is more pronounced for firms whose major customers are individual consumers. However, as shown in columns 3 and 5 of table 5, neither the main effect of AVG_POSITIVE nor the interaction effect between AVG_POSITIVE and B2B is statistically significant. To the extent that only extremely satisfied customers are likely to initiate positive comments and only extremely dissatisfied customers are likely to initiate negative comments, the summary statistic, POSITIVE, is more susceptible to extremity bias and thus may not fairly represent the broad customer response to the company's products and brands. Lack of the predictive power is consistent with the interpretation that AVG_POSITIVE is more susceptible to extremity bias. The last column of table 5 suggests that, when both AVG_PURCHASE and AVG_POSITIVE are included jointly as explanatory variables, AVG_PURCHASE and its interaction with B2B are statistically significant, whereas AVG_POSITIVE and its interaction with B2B continue to be statistically insignificant.

Given that the predictive power of Twitter comments is more pronounced for the subsample of firms whose major customers are consumers, table 6 presents the results on how the predictive power of Twitter comments about products and brands varies with the intensity of advertising activities for this particular subsample. As shown in the first column of table 6, given that the standard deviation in the natural log of

AVG_PURCHASE is 1.779, the slope coefficient on $\ln(\text{AVG_PURCHASE})$ implies that an increase of one standard deviation in the number of tweets that mention purchase intent or actions for a company's products or brands is associated with an increase of 6.4% in sales growth. As shown in the second column of table 6, the slope coefficient on the interaction between AVG_PURCHASE and ADVERTISE is -0.411 and statistically significant ($p\text{-value} = 0.001$), suggesting that consumer-generated brand awareness substitutes for producer-generated brand awareness. As shown in the last column of table 6, when both summary statistics are included jointly in the model, the slope coefficient on the interaction between AVG_PURCHASE and ADVERTISE is -0.369 and statistically significant, and the slope coefficient on the interaction between AVG_POSITIVE and ADVERTISE is statistically insignificant.

Table 7 presents the predictive power of Twitter comments by the type of tweet handle with respect to sales growth. As shown in the first column of panel A of table 7, the slope coefficient on MEDIA_AVG_PURCHASE is 0.543 and statistically significant ($p\text{-value} = 0.001$). As shown in the second column, the slope coefficient on EXPERT_AVG_PURCHASE is 2.255 and statistically significant ($p\text{-value} = 0.04$). As shown in the third column, the slope coefficient on CROWD_AVG_PURCHASE is 0.034 and statistically significant ($p\text{-value} = 0.001$). The results collectively suggest that PURCHASE initiated by the three types of tweet handle is predictive of upcoming sales growth on a stand-alone basis. The magnitude of the slope coefficients suggests that the marginal effect of one additional mention of purchase intent or actions is higher for those initiated by the media and product experts than those initiated by the crowd. Finally, as shown in the last column of panel A of table 7, when PURCHASE by three types of tweet

handle are included jointly as explanatory variables, the slope coefficient on CROWD_AVG_PURCHASE remains to be positive and statistically significant (p -value = 0.02), whereas the slope coefficients on MEDIA_AVG_PURCHASE and EXPERT_AVG_PURCHASE become statistically insignificant. This indicates that the predictive power of PURCHASE with respect to upcoming sales growth is dominated by comments initiated by the crowd, consistent with the interpretation of the wisdom of crowds. Given that the standard deviation of the natural log of the number of purchase intent or actions initiated by the crowd is 1.812, the slope coefficient on $\ln$ (CROWD_AVG_PURCHASE) implies that an increase of one standard deviation in PURCHASE initiated by the crowd is associated with an increase of 6.2% in sales growth.

As shown in panel B of table 7, neither POSITIVE by the media nor POSITIVE by product experts is predictive of upcoming sales growth. POSITIVE initiated by the crowd is also not predictive of upcoming sales growth.

4.3. Twitter information and unexpected sales growth

Table 8 presents the results on whether the volume and valence of Twitter comments capture the unexpected component of sales growth. As shown in the first column of table 8, the slope coefficient on AVG_PURCHASE is 0.004 and statistically significant (p -value = 0.077). As shown in the second column, the slope coefficient on AVG_POSITIVE is negative but statistically insignificant. The third and fourth columns of table 8 convey the same message: The slope coefficient on AVG_PURCHASE is positive and statistically significant, whereas the slope coefficient on AVG_POSITIVE is

statistically insignificant. The results from table 8 suggest that the volume of Twitter comments captures “new” information about sales that is not identified by analysts. The slope coefficient on $\text{Ln}(\text{AVG_PURCHASE})$ implies that an increase of one standard deviation in the number of tweets that mention purchase intent or action for a company’s products or brands is associated with an increase of 0.71% in sales growth that is not anticipated by analysts. The predictive power of PURCHASE with respect to the component of sales growth that is not anticipated by analysts is about one-ninth of that for sales growth itself in economic significance. The lack of explanatory power for AVG_POSITIVE when the predicted variable is unexpected sales growth is not surprising given that AVG_POSITIVE is not informative about upcoming sales growth, as shown in table 6.

Panel A and panel B of table 9 explore the source of the predictive power of Twitter comments with respect to unexpected sales growth. On a stand-alone basis, as shown in the first three columns of panel A of table 9, the slope coefficient on $\text{MEDIA_AVG_PURCHASE}$ is 0.048 and statistically significant. Given that the standard deviation in the natural log of $\text{MEDIA_AVG_PURCHASE}$ is 0.07, the slope coefficient on $\text{MEDIA_AVG_PURCHASE}$ implies that an increase of one standard deviation in PURCHASE initiated by the media is associated with an increase of 0.33% in sales growth that is not anticipated by analysts. The explanatory power of $\text{MEDIA_AVG_PURCHASE}$ is consistent with the interpretation that the media reach a much wider audience because of their greater number of follows on Twitter and, therefore, potentially trigger a greater than expected word-of-mouth effect. On a stand-alone basis, the slope coefficient on $\text{CROWD_AVG_PURCHASE}$ is 0.004 and

statistically significant. The virtually identical slope coefficients on AVG_PURCHASE from table 8 and on CROWD_AVG_PURCHASE from table 9 indicate that the predictive power of AVG_PURCHASE with respect to unexpected sales growth is dominated by tweets initiated by the crowd, again consistent with the “wisdom of crowds”. As shown in the second column, the slope coefficient on EXPERT_AVG_PURCHASE is not statistically significant. When the summary statistics by the three types of tweet handle are included jointly as explanatory variables, as shown in the last column, none of the three slope coefficients on AVG_PURCHASE by the three types of tweet handles is statistically significant.

As shown in panel B of table 9, on a stand-alone basis, the slope coefficient on EXPERT_AVG_POSITIVE is 0.005 and statistically significant (p -value = 0.001). In contrast, the slope coefficients on both MEDIA_AVG_POSITIVE and CROWD_AVG_POSITIVE are statistically insignificant. As shown in the last column, when POSITIVE by the three types of tweet handle are included jointly as explanatory variables, the slope coefficient on EXPERT_AVG_POSITIVE is 0.007 and continues to be statistically significant. The results suggest that the predictive power of the valence of Twitter comments with respect to unexpected sales growth is largely attributable to comments initiated by experts. Given that the standard deviation of EXPERT_AVG_POSITIVE is 8.15%, the slope coefficient implies that an increase of one standard deviation in the ratio of the number of positive comments initiated by product experts relative to the number of non-neutral comments is associated with an increase of 0.06% in sales growth that is not anticipated by analysts. The differential predictive power of POSITIVE by the three types of tweet handle is expected given that

expert reviews typically cover every brand within a product category and, therefore, positive product tweets generated by experts are less subject to extremity bias and are perceived to have a higher level of authority.

4.4 Robustness checks

First, in the main empirical specifications, lagged revenue is the denominator of the dependent variable, a stand-alone control variable, and the numerator of another control variable. Measurement errors in the repeated variable can generate an extreme level of bias in each coefficient, including those on the twitter variables. I do not believe this is an issue because revenue is not measured with error. Furthermore, I use firm fixed effects to handle the potential “ratio variable problem.” Specifically, firm fixed effects are included in lieu of lagged sales and sales growth to account for mean reversion in sales growth. The results with firm fixed effects are qualitatively similar to, but quantitatively smaller than, those reported in table 6 and table 8. For instance, in untabulated results, when the predicted variable is the same-quarter sales growth, with firm fixed effects, the slope coefficients on both AVG_PURCHASE and AVG_POSITIVE are positive and statistically significant on a stand-alone basis. When both variables are included jointly as explanatory variables, the slope coefficient on AVG_PURCHASE is 0.030 and statistically significant (p -value = 0.02) and that on the interaction between AVG_PURCHASE and advertising expense is -0.519 and statistically significant (p -value = 0.04). When the predicted variable is unexpected sales growth, on a stand-alone basis, the slope coefficient on AVG_PURCHASE is 0.002 and statistically significant (p -value = 0.06).

Second, I address whether the results are robust to measures of standardized unexpected sales. Following Da et al. (2011b), standardized unexpected sales relative to the previous quarter (the same quarter in the previous fiscal year) is defined as the difference in sales in the current quarter and that in the previous quarter (in the same quarter in the previous fiscal year) divided by the standard deviation of sales during the past eight quarters. The results are robust to both measures of standardized unexpected sales. In untabulated results, when the dependent variable in model (3) is standardized unexpected sales relative to prior quarter (the same quarter in the previous fiscal year), the slope coefficient on AVG_PURCHASE is 0.062 (0.136) and statistically significant with a p -value of 0.001 (0.001), whereas the slope coefficient on AVG_POSITIVE is 0.487 (0.679) but statistically insignificant. When the volume and valence of Twitter comments are included as explanatory variables jointly in model (3), the slope coefficient on AVG_PURCHASE is 0.075 (0.138) and continues to be statistically significant.

Finally, in the main specification, PURCHASE is normalized by the total number of tweets circulated on Twitter for that particular day to control for the temporal growth of Twitter while maintaining the cross-sectional variation across firms. An alternative approach is to normalize PURCHASE by the number of tweets that discuss the same company's products and brands. This alternative approach results in a ratio-like variable (PURCHASERATIO), which also controls for the temporal growth in Twitter. However, PURCHASERATIO ignores the cross-sectional variation in the level of interest in a given company's products and brands. I expect that the correlation between PURCHASERATIO and sales is ambiguous in sign. To illustrate this point numerically, firm A 's products generated very limited interest among customers and received only one

customer comment, which mentioned his recent purchase of firm *A*'s products. Firm *B*'s and firm *C*'s products generated significant interest among customers and each received 100 comments, of which 50 tweets mentioned purchase of firm *B*'s products and 100 tweets mentioned purchase of firm *C*'s products. Accordingly, PURCHASERATIO is 50% for firm *B* and 100% for both firm *A* and firm *C*. Under the reasonable assumption that the number of tweets mentioning purchase intent or actions is proportional to the total volume of sales and the proportion is similar across firms at a given point in time, the total volume of sales for firm *B* is 50 times of that for firm *A*, but is only half of firm *C*'s. Consistent with the conjecture, in untabulated results, if AVG_PURCHASERATIO replaces AVG_PURCHASE in model (2), the slope coefficients on AVG_PURCHASERATIO are all insignificant for the subsample of firms whose major customer base is consumers.

V. Conclusion

This study suggests that there is “new” information on social media that can be used to predict firm fundamentals. In particular, third-party-generated product information on Twitter, once aggregated at the firm level, is predictive of both upcoming sales and the unexpected component of sales growth at the firm level. A direct follow-up research question would be an examination of whether this predictive power results in abnormal returns around the announcement date of upcoming quarterly sales conditional on the volume and valence of tweets discussing products and brands. Another direction of future research would be an investigation of the interplay between third-party generated comments and firm-initiated disclosure on social media.

This paper also finds that the predictive power of the volume of Twitter comments is dominated by the wisdom of crowds, but the predictive power of the valence of Twitter comments is largely attributable to expert comments. Future work can explore the causes and consequences of the differential predictive power of the wisdom of crowds and expert comments. My conjecture is that the wisdom of crowds would be most useful when the underlying construct is quantitative and expert opinion would be most useful when the underlying construct is qualitative. Given the Securities and Exchange Commission's recent crackdown on undisclosed "paid stock promotions" on financial websites, the finding from this study implies that targeting expert writers with a high number of followers on social media would be most effective in achieving the intended objective.

References

- ANTWEILER, W., and M. FRANK. 'Is All that Talk Just Noise? The Information Content of Internet Stock Message Boards.' *Journal of Finance* 59 (2004): 1259–94.
- AMIR, E., and B. LEV. 1996. 'Value-Relevance of Nonfinancial Information: The Wireless Communications Industry.' *Journal of Accounting and Economics* 22 (1996): 3–30.
- ASUR, S., and B. HUBERMAN. 'Predicting the Future with Social Media.' *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology* 1 (2010): 492–99.
- BARRON, O.E., O. KIM, S.C. LIM, and D.E. STEVENS. 'Using Analysts' Forecasts to Measure Properties of Analysts' Information Environment.' *Accounting Review* 73 (2003): 421–33.
- BERG, J., R. FORSYTHE, and T. RIETZ. 'What makes markets predict well? Evidence from the Iowa Electronic Markets.' In *Understanding Strategic Interaction*, 444–63. Springer Berlin Heidelberg, 1997.
- BLANKESPOOR, E., G. MILLER, and H. WHITE. 'The Role of Dissemination in Market Liquidity: Evidence from Firms' Use of Twitter.' *Accounting Review* 89 (2014): 79–112.
- CHEN, H., P. DE, Y. HU, and B. HWANG. 'Wisdom of Crowds: The Value of Stock Opinions Transmitted Through Social Media.' *Review of Financial Studies* 27 (2014): 1367–1403.
- CHEN, Y., and J. XIE. 'Online Consumer Review: Word-of-Mouth as a New Element of Marketing Communication Mix.' *Management Science* 54 (2008): 477–91.
- CHEVALIER, J.A., and D. MAYZLIN. 'The Effect of Word of Mouth on Sales: Online Book Reviews.' *Journal of Marketing Research* 43 (2006): 345–54.
- CHOI, H., and H. VARIAN. 'Predicting the Present with Google Trends.' *Economic Record* 88 (2012): 2–9.
- DA, Z., J. ENGELBERG, and P. GAO. 'In Search of Attention.' *Journal of Finance* 5 (2011a): 1461–99.
- DA, Z., J. ENGELBERG, and P. GAO. 'In Search of Fundamentals.' AFA 2012 Chicago Meetings Paper, 2011b. Available at SSRN: <https://ssrn.com/abstract=1589805>.

DAS, S., and M. CHEN. 'Yahoo! for Amazon: Sentiment Extraction from Small Talk on the Web.' *Management Science* 53 (2007): 1375–88.

DRAKE, M.S., D. ROULSTONE, and J.R. THORNOCK. 'Investor Information Demand: Evidence from Google Searches Around Earnings Announcements.' *Journal of Accounting Research* 50 (2012): 1001–40.

ETTREDGE, M., J. GERDES, and G. KARUGA. 'Using Web-based Search Data to Predict Macroeconomic Statistics.' *Communications of the ACM* 48 (2005): 87–92.

GAYO-AVELLO, D. 'A Meta-Analysis of State-of-the-Art Electoral Prediction from Twitter Data.' *Social Science Computer Review* 31 (2013): 649–79.

GOEL, S., J.M. HOFMAN, S. LAHAIE, D.M. PENNOCK, and D.J. WATTS. 'Predicting Consumer Behavior with Web Search.' *Proceedings of the National Academy of Sciences* 107 (2010): 17486–90.

HIRSCHEY, M., V. RICHARDSON, and S. SCHOLZ. 'Stock-Price Effects of Internet Buy-Sell Recommendations: The Motley Fool Case.' *Financial Review* 35 (2000): 147–74.

HUANG, H., and N.D. PENNA. 'Constructing Consumer Sentiment Index for U.S. Using Google Searches.' Working paper, University of Alberta, 2009. Available at http://ideas.repec.org/p/ris/albaec/2009_026.html.

ITTNER, C.D., D.F. LARCKER, and M.W. MEYER. 'Subjectivity and the Weighting of Performance Measures: Evidence from a Balanced Scorecard.' *Accounting Review* 78 (2003): 725–58.

JUNG, M., J. NAUGHTON, A. TAHOUN, and C. WANG. 'Corporate Use of Social Media.' Unpublished paper, New York University, 2015. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2588081.

KIRBY, C. 'Everyone's a Critic: Web Sites Hope Online Reviews of Products Lead to Online Buying.' *San Francisco Chronicle* (January 22, 2000): E1.

LEV, B., and S.R. THIAGARAJAN. 'Fundamental Information Analysis.' *Journal of Accounting Research* 31 (1993): 190–215.

LIU, Y. 'Word of Mouth for Movies: Its Dynamics and Impact on Box Office Revenue.' *Journal of Marketing* 70 (2006): 74–89.

MCLAREN, N., and R. SHANBHOGHE. 'Using Internet Data as Economic Indicators.' *Bank of England Quarterly Bulletin* Q2 (2011): 134–40.

MILLER, G., and D. SKINNER. 'The Evolving Disclosure Landscape: How Changes in Technology, the Media, and Capital Markets Are Affecting Disclosure.' *Journal of Accounting Research* 53 (2015): 221–39.

REINSTEIN, A., and C. SNYDER. 'The Influence of Expert Reviews on Consumer Demand for Experience Goods: A Case Study of Movie Critics.' *Journal of Industrial Economics* 53 (2005): 27–50.

VOSEN, S., and T. SCHMIDT. 'Forecasting Private Consumption: Survey-based Indicators vs. Google Trends.' *Journal of Forecasting* 30 (2011): 565–78.

WU, L., and E. BRYNJOLFSSON. 'The Future of Prediction: How Google Searches Foreshadow Housing Prices and Sales.' Working paper, University of Pennsylvania, 2013. Available at <http://ssrn.com/abstract=202293>.

SUNSTEIN, C. *Infotopia: How Many Minds Produce*. New York: Oxford University Press, 2008.

TUMARKIN, R., and R.F. WHITELOW. 'News or Noise? Internet Postings and Stock Prices.' *Financial Analysts Journal* (2001): 41–51.

Exhibit 1
Examples of Tweets Containing Purchase Intent or Action

1. I then got this iPhone 6 Plus trying to be big baller and absolutely hate it?
2. I think I'm just going to get an iPhone 6 instead of the S7 Edege?
3. Got the iPhone 6 Plus and I feel like I'm holding a tv
4. My dad ordered me a new phone, he ordered the iPhone 6s instead of the IPhone 6. Ain't even mad?
5. My uncle just got an iPhone and i taught him how to FaceTime now he FaceTimes me all the time. At work at shoprite... Love that guy

Exhibit 2

Examples of the Valence of Tweets

A. Positive tweets

1. This technology is amazing. Talked to my mother in-law tonight on my, new to me, iPhone using Facetime. "Wind in off the water 2 weeks now."
2. Always nice to get a convert to the iPhone team. #bluebubblegoodness #welcome
3. I remember back when I was younger. All I had was an iPhone and iMovie. I remember being so excited about the footage. The thrill was great.
4. @KittyKaty_14? I'm glad you're jealous lol. Don't worry about it, you'll be more than fucking happy when you get new iPhone 6s Plus once.
5. @AccuWxBeck love, love, love my iPhone. I will be sad when I leave At&mt and don't have it anymore

B. Neutral tweets

1. Whenever I'm around my mom act like she don't know how to work any electronic known to man, but she got that iPhone 6 tho?
2. Anyone have an extra AT&T iPhone?
3. iPhone owners holding onto their phones longer - CNET <https://t.co/kCxMZNO0KG>
4. Another way to reach us... Have an iPhone and use iMessage now send us messages / photos / videos at columbuzz@icloud.com
5. Jaw-Dropping Scene Captured in Germany From The Weather Channel iPhone App. <https://t.co/k0WFGnMivx> <https://t.co/tfO78W7nC4>

C. Negative tweets

1. I then got this iPhone 6 Plus trying to be big baller and absolutely hate it?
2. Ok, iPhone! Does the reminders work properly? I miss all of them! That's embarrassing and frustrating. Especially in the morning.
3. @LetsRabbit sorry to complain once again? but whenever I use your app the wifi goes out of my iPhone or iPad (whichever I watch it on) why?
4. Why did my iPhone 7 get so slow and laggy in like the last 2 days??
5. My Iphone 6 camera sucks

Exhibit 3
List of products and brands for 25 randomly selected business-to-consumer firms

Company	Products and Brands			
Chipotle Mexican Grill	Chipotle Mexican Grill	Pizzeria Locale	ShopHouse Southeast Asian Kitchen	—
Delta Air Lines Inc.	Delta Air Lines	Delta Sky Magazine	Delta Sky Club	SkyMiles
EBay Inc.	Close5	Decide.com	eBay	Ebay Enterprise
	GSI Commerce	Hunch	Magento	RedLaser
	Sell for Me	Start Tank	Svpply	Twice
	Half.com	Shopping.com	StubHub	—
Fitbit Inc	Alta	Aria	Blaze	Charge
	Fitbit	Flex	Force	One
	Surge	Ultra	Zip	—
General Motors Company	ACDelco	Buick	Cascada	Enclave
	Encore	Envision	Lacrosse	Regal
	Verano	ATS	ATS-V	Cadillac
	CT6	ELR	Escalade	ESV
	SRX	XT5	XTS	Camaro
	Chevrolet	Colorado	Corvette	Cruze
	Equinox	Impala	Malibu	Silverado
	Sonic	Spark	SS Sedan	Suburban
	Tahoe	Traverse	Trax	Volt
	General Motors	Acadia	Canyon	Denali
	GMC	Maven	Savanna	Sierra 1500
	Sierra 2500	Sierra 3500	Terrain	Yukon
	Holden	H2	H3	Hummer
	OnStar	Opel	G3	G5
	G6	G8	Pontiac	Solstice
	Torrent	Vibe	Sidecar	Vauxhall
GoPro, Inc.	Quik	Splice	Vemory	GoPro
	Hero	Hero Session	HeroCast	Karma
	Omni	Kolor	Stupeflix	—
Groupon, Inc.	Groupon	Groupon Getaways	Groupon Pages	OrderUp
	Savored	Snap	Pretty Quick	—
Hilton Worldwide	Canopy Hotel	Conrad Hotels & Resorts	Curio Hotel	DoubleTree
	Embassy Suites Hotels	Hampton	Hilton Garden Inn	Hilton Grand Vacations
	Hilton Hotels & Resorts	Hilton Hhonors	Home 2 Suites	Homewood Suites
	Parc 55	Tru Hotel	Waldorf Astoria	—
L Brands, Inc.	Bath & Body Works	Henri Bendel	La Senza	Victoria's Secret
	VS Pink	—	—	—
LinkedIn Corporation	Cardmunch	Compilr	Elevate	LinkedIn
	LinkedIn Job Search	LinkedIn Pulse	Rapportive	Refresh.io
	Lynda.com	Video2Brain	Newsle	LinkedIn Slideshare
Lululemon Athletica	Ivivva	Lululemon Athletica	—	—

Exhibit 3 (continued)

Company	Products and Brands			
Hershey's	Allan Bear Bites	Allan Big Foot	Allan Sour BlueRaspberry	Allan Hot Lips
	Allan Sour Watermelon Slices	Allan Sour Wormies	Allan Wormies	BreathSavers
	Bubble Yum	Good & Plenty	Ice Breakers	Ice Breakers Duo
	Ice Breakers Ice Cubes	Jolly Rancher	Lancaster	PayDay
	Pelon Pelo Rico	Take 5	Twizzlers	ZAGNUT
	Zero	5th Avenue	Almond Joy	barkThins
	Brookside	Cadbury	Dagoba	Heath
	Hershey's	Hershey's Bliss	Hershey's Kisses	Hershey's Syrup
	Kit Kat	Krackel	Milk Duds	Mounds Bars
	Mr. Goodbar	Reese's	ROLO	Scharffen Berger
	SKOR	Symphony	Whatchamacallit	Sofit
	YORK	Krave Jerky	—	—
McDonald's	McDonald's	—	—	—
Match Group	Match.com	OkCupid	Chemistry.com	DateHookup
	Humin	IndiaMatch	LDSPlanet	LoveAndSeek
	OurTime.com	Plenty of Fish	Singlesnet.com	Stepout
	AsianPeopleMeet	BabyBoomerPeople Meet	BBPeopleMeet.com	BlackBabyBoomerPe opleMeet
	BlackChristianPeopleMeet	BlackPeopleMeet	CatholicPeopleMeet	ChinesePeopleMeet
	DemocraticPeopleMeet	DivorcedPeopleMee t	InterracialPeopleMeet	JPeopleMeet
	LatinoPeopleMeet	LittlePeopleMeet	MarriagemindedPeopl eMeet	PetPeopleMeet
	RepublicanPeopleMeet	SeniorBlackPeople Meet	SeniorPeopleMeet	SingleParentMeet
	The Princeton Review	Tinder	Twoo	
New York Times Company	Idea Lab	International New York Times	New York Times Conferences	New York Times Cooking
	New York Times Magazine	NYTimes.com	T Brand Studio	The New York Times
	Times Digest	Times Journeys	Times Talk	—
PepsiCo Inc.	AMP Energy	Cheetos	Cracker Jack	Doritos
	El Isleno	Frito-Lay	Frito's	Funyuns
	Grandma's	Lay's	Maui Style	Miss Vickie's
	Munchies	Munchos	Nut Harvest	Rold Gold
	Ruffles	Sabra	Sabritones	Santitas
	Smartfood Popcorn	Spitz	Stacy's	Sun Chips
	Tostitos	Gatorade	Mountain Dew	7UP
	Brisk	Citrus Blast	IZZE	Mug Root Beer
	Sierra Mist	1893	Pepsi	Aunt Jemima
	Cap'n Crunch	King Vitamin	Life Cereal	Quaker
	Quisp Cereal	Rice-A-Roni	Matador Beef Jerky	Aquafina
	Naked Juice	Ocean Spray	Propel Zero	Pure Leaf
	Sobe	Tropicana	—	—

Exhibit 3 (continued)

Company	Products and Brands			
Starbucks Corp.	Ethos Water	Evolution Fresh	Frappuccino	Hear Music
	La Boulange	Roy Street Coffee & Tea	Seattle's Best	Starbucks
	Tazo	Teavana	Verismo	
SeaWorld Parks & Entertainment	Adventure Island	Aquatica	Busch Gardens	Discovery Cove
	SeaWorld	Sesame Place	Water Country USA	
AT&T Inc	Aio Wireless	AT&T	AT&T Mobility	BellSouth
	Cricket Wireless	Audience Network	DirecTV	Southwestern Bell
	Uverse	—	—	—
Target Corp.	A Bullseye View	ClearRX	Dermstore	Target
Twitter	Cover	Digits	MoPub	Periscope
	SnappyTV	Trendrr	Curator	Madbits
	Marakana	Mitro	Twitter	Vine & Zipdial
Volkswagen Group	A1, A3, A4, A5, A6, A7,A8	Allroad	Audi	Q3, Q5, Q7
	R8, RS7, XL1	S3, S4, S5, S6, S7, S8, SQ5	TT Coupe	TT Roadster
	TTS	Arnage Saloon	Azure Convertible	Azure T
	Bentayga	Brooklands	Continental	Flying Spur
	Mulsanne	State Limousine	Zagato	Bently
	Bugatti	Chiron	Veyron	748;749; 848;718
	911, 916, 996, 998, 999, 1098, 1198	Diavel	Ducati	Hypermotard
	Monster	Multistrada	PaulSmart 1000	Scrambler
	ST	Superbike	SuperSport	Aventador
	Centenario	Diablo	Egoista	Gallardo
	Huracan	Lamborghini	Murcielago	Reventon
	Sesto Elemento	Veneno Roadster	Cayenne	Cayman
	Macan	Panamera	Porsche	Beetle
	CC	EOS	Golf	Jetta
	Passat	Tiguan	Touareg	Volkswagen
Wal-Mart Stores	Sam's Club	Vudu	Wal-Mart	Luvocracy
	Spark Studio	Walmart Labs	Yumprint	—
Yahoo Inc.	Aviate	Beam It	Blink	Cooliris
	Docspad	Luminate	PlayerScale	Polyvore
	Qwiki	RayV	Summly	Vizify
	Yahoo Radar	Yahoo Screen	Zofari	MessageMe
	Yahoo Livetext	Yahoo! Mail	Yahoo! Messenger	Tumblr
	BrightRoll	Yahoo! Advertising	Yahoo! Search Marketing	Yahoo! Web Analytics
	Xobni	Yahoo	Yahoo! Buzz	Yahoo! Answers
	Yahoo! Axis	Yahoo! Developer Network	Yahoo! Directory	Yahoo! Esports
	Yahoo! Finance	Yahoo! Games	Yahoo! Green	Yahoo! Groups
	Yahoo! Kids	Yahoo! Local	Yahoo! Maps	Yahoo! Meme
	Yahoo! Mobile	Yahoo! Movies	Yahoo! Music	Yahoo! News
	Yahoo! Personals	Yahoo! Pipes	Yahoo! Publisher Network	Yahoo! Real Estate
	Yahoo! Screen	Yahoo! Search	Yahoo! Shopping	Yahoo! Sports
	Yahoo! Travel	Yahoo! TV	Yahoo! Voice	—
Yum! Brands Inc.	KFC	Pizza Hut	Taco Bell & Winestreet	U.S. Taco Co. and Urban Taproom

Figure 1
Timeline for Measuring Unexpected Sales and Corresponding Twitter Comments




Table 1
Sample Formation

	Number of firms	Number of firm- quarter observations
Product and brands comments on Twitter	1,937	30,992
Missing information from Compustat and AVG_PURCHASE	(97)	(4,656)
Firm-quarter observations that have information on AVG_PURCHASE and financial information from Compustat	1,840	26,336
Missing information on AVG_POSITIVE	(752)	(15,668)
Final sample to test the predictive power of Twitter comments with respect to upcoming sales	1,088	10,668
Business-to-business subsample of firms whose major customer base is businesses	922	8,249
Business-to-consumer subsample of firms whose major customer base is consumers	166	2,419
Missing information on PURCHASE by the media, PURCHASE by experts, and PURCHASE by the crowd	(7)	(122)
Missing information from I/B/E/S	(0)	(16)
Business-to-consumer subsample to test the predictive power of PURCHASE by the type of tweet handle with respect to unexpected sales	159	2,281
Missing information on POSITIVE by the media, POSITIVE by experts, and POSITIVE by the crowd	(129)	(2,037)
Business-to-consumer subsample to test the predictive power of POSITIVE by the type of tweet handle with respect to unexpected sales	30	244

Table 2
Descriptive Statistics

Panel A: Descriptive statistics of the sample used to test the cross-sectional variation in the predictive power of Twitter

	Overall Sample						Subsample where major customer base is consumers (N = 2,419)		Subsample where major customer base is businesses (N = 8,249)	
Variable	N	Mean	Median	Std. Dev.	Min	Max	Mean	Median	Mean	Median
AVG_PURCHASE _{i,q}	10,668	225.78	2.92	2085.75	0	57442.48	973.47	69.37	6.52***	1.92
AVG_POSITIVE _{i,q}	10,668	0.88	0.92	0.13	0	1.00	0.81	0.82	0.91***	0.96
SAMEQUARTER_SALESGROWTH _{i,q}	10,668	0.20	0.04	2.41	-1.00	119.07	0.09	0.04	0.24	0.05
SALES _{i,q}	10,668	3,373	573	9,368	0	129,760	8,579	2,874	1,847***	302
ADVERTISE _{i,y-1}	10,668	0.03	0.00	0.06	0.00	1.21	0.03	0.02	0.03	0.00
CHG_BACKLOG _{i,q}	10,668	0.04	0.00	2.01	-4.38	196.63	0.01	0.00	0.05	0.00

*** Differences are significant at the 0.01 level between the two subsamples.

Table 2 (continued)

Panel B: Descriptive statistics of the sample used to test whether Twitter comments are predictive of unexpected sales for the subsample of firms whose major customer base is consumers

Variable	N	Mean	Median	Std. Dev.	Minimum	Maximum
MEDIA_PCT _{i,q}	2,281	0.04%	0.01%	0.11%	0.00%	3.81%
EXPERT_PCT _{i,q}	2,281	0.01%	0.00%	0.04%	0.00%	0.79%
MEDIA_PURCHASE _{i,q}	2,281	0.02	0.00	0.10	0.00	1.87
EXPERT_PURCHASE _{i,q}	2,281	0.00	0.00	0.01	0.00	0.23
CROWD_PURCHASE _{i,q}	2,281	922.99	69.38	4146.11	1.00	57442.09
MEDIA_POSITIVE _{i,q}	224	95.39%	100.00%	12.34%	25.00%	100.00%
EXPERT_POSITIVE _{i,q}	244	97.61%	100.00%	8.15%	25.00%	100.00%
CROWD_POSITIVE _{i,q}	244	80.91%	82.14%	9.60%	0.00%	100.00%
UNEXPECTED_SALESGROWTH _{i,q}	2,281	0.01	0.00	0.20	-1.05	6.63
NUM_FORECAST _{i,q}	2,281	17.58	17.00	8.52	1.00	45.00
ACTUAL_FORECAST_DAYS _{i,q}	2,281	16.01	14.00	9.36	-8.00	49.00
SALES _{i,q-1}	2,281	8,658	2,958	15,569	21	131,565
SAMEQUARTER_SALESGROWTH _{i,q-4}	2,281	0.01	0.00	0.18	-0.71	7.57
CHG_BACKLOG _{i,q-1}	2,281	0.03	0.02	0.04	0.00	0.27
ADVERTISE _{i,y-1}	2,281	0.04%	0.01%	0.11%	0.00%	3.81%

Table 2
(continued)

PURCHASE is the number of tweets on Twitter that express purchase intent for a company i 's products and brands on a daily basis. POSITIVE is the ratio of the number of positive tweets over the number of non-neutral (positive and negative) tweets about a company i 's products and brands on a daily basis. $AVG_PURCHASE_{i,q}$ averages daily PURCHASE over quarter q . $AVG_POSITIVE_{i,q}$ averages daily POSITIVE over quarter q . $SAMEQUARTER_SALESGROWTH_{i,q}$ is measured as the percentage change in sales in quarter q relative to sales in quarter $q-4$. B2B is defined as 1 if a firm's major customers are businesses, and 0 if a firm's major customers are consumers. $SALES_{i,q-4}$ is measured as sales in the same quarter of the previous fiscal year. $SALES_{i,q-1}$ is measured as sales in the prior quarter. $SAMEQUARTER_SALESGROWTH_{i,q-4}$ is measured as the same-quarter sales growth in quarter $q-4$. $ADVERTISE_{i,q-1}$ is measured as the ratio of pro-rated advertising expense over sales during the most recent fiscal quarter. $CHG_BACKLOG_{i,q-1}$ is measured as the ratio of the change in deferred revenue over sales during the previous quarter. $MEDIA_PCT_{i,q}$ is the percentage of tweets about products and brands that are initiated by the media. $EXPERT_PCT_{i,q}$ is the percentage of tweets about products and brands that are initiated by product experts. $NUM_FORECAST_{i,q}$ is the number of forecasts included in the consensus forecast for sales in quarter q . $ACTUAL_FORECAST_DAYS_{i,q}$ is measured as the number of calendar days between the report date of realized sales in quarter q and the report date of the consensus forecast. $MEDIA_AVG_PURCHASE_{i,q}$ is $AVG_PURCHASE$ initiated by the media over quarter q . $MEDIA_AVG_POSITIVE_{i,q}$ is $AVG_POSITIVE$ initiated by the media over quarter q . $EXPERT_AVG_PURCHASE_{i,q}$ is $AVG_PURCHASE$ initiated by product experts over quarter q . $EXPERT_AVG_POSITIVE_{i,q}$ is $AVG_POSITIVE$ initiated by product experts over quarter q . $CROWD_AVG_PURCHASE_{i,q}$ is $AVG_PURCHASE$ initiated by the crowd over quarter q . $CROWD_AVG_POSITIVE_{i,q}$ is $AVG_POSITIVE$ initiated by the crowd over quarter q . $UNEXPECTED_SALES_{i,q}$ is measured as actual sales ($SALES_{i,q}$) minus the consensus analyst forecast ($FORECAST_{i,q}$). $UNEXPECTED_SALESGROWTH_{i,q}$ is measured as $UNEXPECTED_SALES_{i,q}$ divided by sales in the same quarter of the previous fiscal year ($SALES_{i,q-4}$).

Table 3
Correlation Table (Pearson Correlations above Diagonal and Spearman Correlations below Diagonal)

Panel A: Correlation of the sample used to examine the cross-sectional variation in the predictive power of Twitter comments

	SAMEQUARTER_SALES GROWTH _{i,q}	Ln (AVG_ PURCHASE _{i,q})	AVG POSITIVE _{i,q}	Ln (SALES _{q-4})	SAMEQUARTER_ SALES GROWTH _{i,q-4}	ADVERTIS E _{i,y-1}	CHG_ BACKLOG _{i,q-1}
SAMEQUARTER SALES_GROWTH _{i,q}	1.000	-.040**	0.019	-.186**	0.016	.040**	0.007
		<i>0.000</i>	<i>0.052</i>	<i>0.000</i>	<i>0.090</i>	<i>0.000</i>	<i>0.461</i>
Ln (AVG_PURCHASE _{i,q})	-.055**	1.000	-.282**	.451**	-.019*	.148**	-0.014
	<i>0.000</i>		<i>0.000</i>	<i>0.000</i>	<i>0.045</i>	<i>0.000</i>	<i>0.159</i>
AVG_POSITIVE _{i,q}	.026**	-.408**	1.000	-.188**	-0.019	0.002	-0.010
	<i>0.006</i>	<i>0.000</i>		<i>0.000</i>	<i>0.051</i>	<i>0.858</i>	<i>0.280</i>
Ln (SALES _{i,q-4})	-.232**	.457**	-.306**	1.000	-.044**	-.083**	-.042**
	<i>0.000</i>	<i>0.000</i>	<i>0.000</i>		<i>0.000</i>	<i>0.000</i>	<i>0.000</i>
SAMEQUARTER_ SALES_GROWTH _{i,q-4}	.292**	-.027**	-0.009	-.108**	1.000	.040**	-0.001
	<i>0.000</i>	<i>0.005</i>	<i>0.337</i>	<i>0.000</i>		<i>0.000</i>	<i>0.950</i>
ADVERTISE _{i,y-1}	0.013	.392**	-.110**	.047**	.050**	1.000	-0.003
	<i>0.177</i>	<i>0.000</i>	<i>0.000</i>	<i>0.000</i>	<i>0.000</i>		<i>0.752</i>
CHG_BACKLOG _{i,q-1}	.092**	.038**	0.004	.022*	.077**	.030**	1.000
	<i>0.000</i>	<i>0.000</i>	<i>0.697</i>	<i>0.024</i>	<i>0.000</i>	<i>0.002</i>	

The correlation coefficient is in bold. *P*-value for correlation coefficients is in italic. **Correlations are significant at 0.01 level. *Correlations are significant at 0.05 level.

Table 3
(continued)

Panel B: Correlation of the sample used to test whether Twitter comments predict unexpected sales growth

	UNEXPECTED_SALES GROWTH _{i,q}	Ln (AVG_PURCHASE _{i,q})	AVG_POSITIVE _{i,q}	ACTUAL_FORECAST_DAYS _{i,q}	Ln (NUM_FORECAST _{i,q})	Ln (SALES _{i,q-1})	SAMEQUARTER_SALES GROWTH _{i,q-4}	CHG_BACKLOG _{i,q-1}	ADVERTISE _{i,q-1}
UNEXPECTED_SALES GROWTH _{i,q}	1.000	0.013	-0.013	0.025	0.009	-.080**	0.028	.083**	0.006
		<i>0.540</i>	<i>0.546</i>	<i>0.235</i>	<i>0.681</i>	<i>0.000</i>	<i>0.187</i>	<i>0.000</i>	<i>0.788</i>
Ln (AVG_PURCHASE _{i,q})	0.007	1.000	-.073**	-0.009	.347**	.187**	.067**	0.016	.163**
	<i>0.737</i>		<i>0.000</i>	<i>0.684</i>	<i>0.000</i>	<i>0.000</i>	<i>0.001</i>	<i>0.448</i>	<i>0.000</i>
AVG_POSITIVE _{i,q}	-0.002	-.089**	1.000	-0.036	-.046*	-.132**	0.010	-.066**	0.026
	<i>0.921</i>	<i>0.000</i>		<i>0.086</i>	<i>0.027</i>	<i>0.000</i>	<i>0.633</i>	<i>0.002</i>	<i>0.217</i>
ACTUAL_FORECAST_DAYS _{i,q}	-0.016	0.021	-0.018	1.000	.068**	.048*	0.015	0.009	0.004
	<i>0.455</i>	<i>0.325</i>	<i>0.379</i>		<i>0.001</i>	<i>0.023</i>	<i>0.479</i>	<i>0.670</i>	<i>0.851</i>
Ln (NUM_FORECAST _{i,q})	.058**	.357**	-0.018	.051*	1.000	.064**	.042*	0.006	-.051*
	<i>0.006</i>	<i>0.000</i>	<i>0.393</i>	<i>0.014</i>		<i>0.002</i>	<i>0.044</i>	<i>0.791</i>	<i>0.015</i>
Ln (SALES _{i,q-1})	-.168**	.148**	-.142**	0.031	.081**	1.000	-.147**	-.064**	-.221**
	<i>0.000</i>	<i>0.000</i>	<i>0.000</i>	<i>0.140</i>	<i>0.000</i>		<i>0.000</i>	<i>0.002</i>	<i>0.000</i>
SAMEQUARTER_SALES GROWTH _{i,q-4}	.086**	.087**	.139**	-0.036	.186**	-.237**	1.000	0.039	-0.039
	<i>0.000</i>	<i>0.000</i>	<i>0.000</i>	<i>0.089</i>	<i>0.000</i>	<i>0.000</i>		<i>0.066</i>	<i>0.063</i>
CHG_BACKLOG _{i,q-1}	.058**	.050*	0.008	-0.003	-0.006	-0.006	.067**	1.000	0.041
	<i>0.005</i>	<i>0.017</i>	<i>0.699</i>	<i>0.899</i>	<i>0.773</i>	<i>0.781</i>	<i>0.001</i>		<i>0.052</i>
ADVERTISE _i	-0.038	.330**	0.009	0.006	-.042*	-.213**	-.077**	-.042*	1.000
	<i>0.072</i>	<i>0.000</i>	<i>0.685</i>	<i>0.775</i>	<i>0.043</i>	<i>0.000</i>	<i>0.000</i>	<i>0.044</i>	

The correlation coefficient is in bold. *P*-value for correlation coefficients is in italic. **Correlations are significant at 0.01 level. *Correlations are significant at 0.05 level.

Table 3
(continued)

PURCHASE is the number of tweets on Twitter that express purchase intent for a company i 's products and brands on a daily basis. POSITIVE is the ratio of the number of positive tweets over the number of non-neutral (positive and negative) tweets about a company i 's products and brands on a daily basis. $AVG_PURCHASE_{i,q}$ averages daily PURCHASE over quarter q . $AVG_POSITIVE_{i,q}$ averages daily POSITIVE over quarter q . $SAMEQUARTER_SALESGROWTH_{i,q}$ is measured as the percentage change in sales in quarter q relative to sales in quarter $q-4$. B2B is defined as 1 if a firm's major customers are businesses, and 0 if a firm's major customers are consumers. $SALES_{i,q-4}$ is measured as sales in the same quarter of the previous fiscal year. $SALES_{i,q-1}$ is measured as sales in the prior quarter. $SAMEQUARTER_SALESGROWTH_{i,q-4}$ is measured as the same-quarter sales growth in quarter $q-4$. $ADVERTISE_{i,q-1}$ is measured as the ratio of pro-rated advertising expense over sales during the most recent fiscal quarter. $CHG_BACKLOG_{i,q-1}$ is measured as the ratio of the change in deferred revenue over sales during the previous quarter. $MEDIA_PCT_{i,q}$ is the percentage of tweets about products and brands that are initiated by the media. $EXPERT_PCT_{i,q}$ is the percentage of tweets about products and brands that are initiated by product experts. $NUM_FORECAST_{i,q}$ is the number of forecasts included in the consensus forecast for sales in quarter q . $ACTUAL_FORECAST_DAYS_{i,q}$ is measured as number of calendar days between the report date of realized sales in quarter q and the report date of the consensus forecast. $MEDIA_AVG_PURCHASE_{i,q}$ is $AVG_PURCHASE$ initiated by the media over quarter q . $MEDIA_AVG_POSITIVE_{i,q}$ is $AVG_POSITIVE$ initiated by the media over quarter q . $EXPERT_AVG_PURCHASE_{i,q}$ is $AVG_PURCHASE$ initiated by product experts over quarter q . $EXPERT_AVG_POSITIVE_{i,q}$ is $AVG_POSITIVE$ initiated by product experts over quarter q . $CROWD_AVG_PURCHASE_{i,q}$ is $AVG_PURCHASE$ initiated by the crowd over quarter q . $CROWD_AVG_POSITIVE_{i,q}$ is $AVG_POSITIVE$ initiated by the crowd over quarter q . $UNEXPECTED_SALES_{i,q}$ is measured as actual sales ($SALES_{i,q}$) minus the consensus analyst forecast ($FORECAST_{i,q}$). $UNEXPECTED_SALESGROWTH_{i,q}$ is measured as $UNEXPECTED_SALES_{i,q}$ divided by sales in the same quarter of the previous fiscal year ($SALES_{i,q-4}$).

Table 4
Validation Check: Cross-sectional Determinants of the Volume and Valence of Twitter Comments

	Dependent Variable = Ln (AVG_PURCHASE _{i,q})			Dependent Variable = AVG_POSITIVE _{i,q}		
	Expected sign	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Expected sign	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)
Intercept		Included	Included		Included	Included
B2B _i	(-)		-3.039*** (234.255)	(?)		0.089*** (105.349)
SIZE _{i,q-1}	(?)	0.044 (0.925)	-0.001 (0.120)	(?)	-0.001** (4.797)	-0.008** (5.083)
Ln (SALES _{i,q-1})	(+)	0.345*** (39.435)	0.145*** (23.077)	(insignificant)	-0.003 (0.265)	0.003 (0.304)
ADVERTISE _{i,q-1}	(?)	6.362*** (17.502)	3.914*** (21.915)	(insignificant)	-0.070 (1.669)	0.002 (0.001)
Industry fixed effects		Yes	Yes		Yes	Yes
Number of observations		26,336	26,336		10,668	10,668
Likelihood ratio		28092.45	14275.20		181.76	171.84

PURCHASE is the number of tweets on Twitter that express purchase intent for a company i 's products and brands on a daily basis. POSITIVE is the ratio of the number of positive tweets over the number of non-neutral (positive and negative) tweets about a company i 's products and brands on a daily basis. AVG_PURCHASE _{i,q} averages daily PURCHASE over quarter q . AVG_POSITIVE _{i,q} averages daily POSITIVE over quarter q . SALES _{$i,q-1$} is measured as sales in the prior quarter. B2B _{i} is defined as 1 if a firm's major customers are businesses, and 0 if a firm's major customers are consumers. SIZE _{$i,q-1$} is measured as the natural log of total assets at the beginning of quarter q . ADVERTISE _{$i,q-1$} is measured as the ratio of pro-rated advertising expense over sales during the most recent fiscal quarter. Standard errors are clustered at the 6-digit NAICS level.

Table 5
Nonfinancial Information on Twitter and Upcoming Sales for the Entire Sample

	Dependent variable = SAMEQUARTER_SALESGROWTH _{i,q}						
	Expected sign	Coefficient (<i>chi</i> -value)	Coefficient (<i>chi</i> -value)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -value)	Coefficient (<i>chi</i> -value)
Intercept		Included	Included	Included	Included	Included	Included
B2B _i	(?)				-0.105 (1.231)	-0.697 (1.595)	-0.362 (0.558)
Ln (AVG_PURCHASE _{i,q})	(+)		0.066*** (9.493)		0.049*** (9.133)		0.048*** (8.967)
AVG_POSITIVE _{i,q}	(+)			-0.300 (1.854)		-0.352 (0.999)	-0.278 (0.640)
Ln (AVG_PURCHASE _{i,q})*B2B _i	(-)				-0.079** (3.932)		-0.078** (4.112)
AVG_POSITIVE _{i,q} *B2B _i	(-)					0.395 (0.557)	0.307 (0.361)
Ln (SALES _{i,q-4})	(-)	-0.191*** (7.740)	-0.216*** (8.483)	-0.194*** (8.364)	-0.219*** (8.763)	-0.221*** (8.523)	-0.219*** (8.821)
SAMEQUARTER_SALESGROWTH _{i,q-4}	(?)	0.002 (0.579)	0.002 (0.660)	0.002 (0.499)	0.002 (0.543)	0.002 (0.591)	0.002 (0.540)
ADVERTISE _{i,q-1}	(+)	1.076 (1.044)	0.660 (0.392)	1.067 (1.033)	0.834 (0.635)	0.796 (0.594)	0.834 (0.621)
CHG_BACKLOG _{i,q-1}	(+)	-0.001 (0.006)	-0.001 (0.016)	-0.001 (0.015)	-0.002 (0.041)	-0.002 (0.036)	-0.002 (0.040)
Industry fixed effects		Yes	Yes	Yes	Yes	Yes	Yes
Number of observations		10,668	10,668	10,668	10,668	10,668	10,668
Likelihood ratio		59,948.9	59,827.2	59,935.1	59,719.2	59,741.6	59,721.4

***Coefficients are significant at 0.01 level. **Coefficients are significant at 0.05 level. *Coefficients are significant at 0.10 level.

Table 5
(continued)

Regression results from model (1):

$$\text{SAMEQUARTER_SALESGROWTH}_{i,q} = \alpha + \beta_1 \text{B2B}_i + \beta_2 \text{AVG_PURCHASE(POSITIVE)}_{i,q} + \beta_3 \text{B2B}_i * \text{AVG_PURCHASE(POSITIVE)}_{i,q} + \beta_4 \ln(\text{SALES}_{i,q-4}) + \beta_5 \text{SAMEQUARTER_SALESGROWTH}_{i,q-4} + \beta_6 \text{CHG_BACKLOG}_{i,q-1} + \beta_7 \text{ADVERTISE}_{i,q-1} + \varepsilon_{it}$$

PURCHASE is the number of tweets on Twitter that express purchase intent for a company i 's products and brands on a daily basis. POSITIVE is the ratio of the number of positive tweets over the number of non-neutral (positive and negative) tweets about a company i 's products and brands on a daily basis. $\text{AVG_PURCHASE}_{i,q}$ averages daily PURCHASE over quarter q . $\text{AVG_POSITIVE}_{i,q}$ averages daily POSITIVE over quarter q . $\text{SAMEQUARTER_SALESGROWTH}_{i,q}$ is measured as the percentage change in sales in quarter q relative to sales in quarter $q-4$. B2B is defined as 1 if a firm's major customers are businesses, and 0 if a firm's major customers are consumers. $\text{SALES}_{i,q-4}$ is measured as sales in the same quarter of the previous fiscal year. $\text{SAMEQUARTER_SALESGROWTH}_{i,q-4}$ is measured as the same-quarter sales growth in quarter $q-4$. $\text{ADVERTISE}_{i,q-1}$ is measured as the ratio of pro-rated advertising expense over sales during the most recent fiscal quarter. $\text{CHG_BACKLOG}_{i,q-1}$ is measured as the ratio of the change in deferred revenue over sales during the previous quarter. Standard errors are clustered at the 6-digit NAICS level.

Table 6
Interplay Between the Word-of-Mouth Effect of Twitter Comments and Advertising
for the Business-to-Consumer Subsample

	Dependent Variable = SAMEQUARTER_SALESGROWTH_{i,q}					
	Expected sign	Coefficient (<i>chi</i> -value)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)
Intercept		Included	Included	Included	Included	Included
Ln (AVG_PURCHASE _{i,q})		0.036*** (10.632)	0.049*** (14.871)			0.048*** (15.550)
Ln (AVG_PURCHASE_{i,q})*ADVERTISE_{i,q-1}	(-)		-0.411** (6.073)			-0.369** (5.869)
AVG_POSITIVE _{i,q}				0.023 (0.038)	-0.129 (0.787)	-0.082 (0.365)
AVG_POSITIVE _{i,q} *ADVERTISE _{i,q-1}	(+)				5.143 (1.892)	5.117 (1.695)
Ln (SALES _{i,q-4})	(-)	-0.096** (5.637)	-0.095** (5.544)	-0.086** (5.014)	-0.087** (5.076)	-0.097** (5.550)
SAMEQUARTER_SALESGROWTH _{i,q-4}	(?)	0.021 (0.244)	0.019 (0.207)	0.033 (0.493)	0.031 (0.454)	0.018 (0.180)
ADVERTISE _{i,q-1}	(+)	-1.122 (2.175)	0.712 (0.421)	-0.771 (1.423)	-5.024 (2.089)	-3.717 (0.919)
CHG_BACKLOG _{i,q-1}	(+)	0.355*** (73.632)	0.369*** (81.268)	0.354*** (57.880)	0.369*** (60.214)	0.376*** (74.701)
Industry fixed effects		Yes	Yes	Yes	Yes	Yes
Number of observations		2,419	2,419	2,419	2,419	2,419
Likelihood ratio		690.2	690.9	699.0	700.3	694.1

***Coefficients are significant at 0.01 level. **Coefficients are significant at 0.05 level. *Coefficients are significant at 0.10 level.

Table 6
(continued)

Regression results from model (2):

$$\text{SAMEQUARTER_SALESGROWTH}_{i,q} = \alpha + \beta_1 \text{AVG_PURCHASE}_{i,q} + \beta_2 \text{ADVERTISE}_{i,q-1} * \text{AVG_PURCHASE}_{i,q} + \beta_3 \text{AVG_POSITIVE}_{i,q} + \beta_4 \text{ADVERTISE}_{i,q-1} * \text{AVG_POSITIVE}_{i,q} + \beta_5 \ln(\text{SALES}_{i,q-4}) + \beta_6 \text{SAMEQUARTER_SALESGROWTH}_{i,q-4} + \beta_7 \text{CHG_BACKLOG}_{i,q-1} + \beta_8 \text{ADVERTISE}_{i,y-1} + \epsilon_{it}$$

PURCHASE is the number of tweets on Twitter that express purchase intent for a company i 's products and brands on a daily basis. POSITIVE is the ratio of the number of positive tweets over the number of non-neutral (positive and negative) tweets about a company i 's products and brands on a daily basis. AVG_PURCHASE _{i,q} averages daily PURCHASE over quarter q . AVG_POSITIVE _{i,q} averages daily POSITIVE over quarter q . SAMEQUARTER_SALESGROWTH _{i,q} is measured as the percentage change in sales in quarter q relative to sales in quarter $q-4$. SALES _{$i,q-4$} is measured as sales in the same quarter of the previous fiscal year. SAMEQUARTER_SALESGROWTH _{$i,q-4$} is measured as the same-quarter sales growth in quarter $q-4$. ADVERTISE _{$i,q-1$} is measured as the ratio of prorated advertising expense over sales during the most recent fiscal quarter. CHG_BACKLOG _{$i,q-1$} is measured as the ratio of the change in deferred revenue over sales during the previous quarter. Standard errors are clustered at the 6-digit NAICS level.

Table 7
Differential Predictive Power of Twitter Comments with respect to Upcoming Sales by Types of Tweet Handle

Panel A: The predictive power of PURCHASE by tweet handle

		DEPENDENT VARIABLE = SAMEQUARTER_SALESGROWTH _{i,q}				
	Expected sign	Coefficient (<i>chi</i> -value)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)
Intercept		Included	Included	Included	Included	Included
Ln (MEDIA_AVG_PURCHASE _{i,q})	(?)	0.543*** (55.567)			0.163 (0.755)	0.159 (0.320)
Ln (EXPERT_AVG_PURCHASE _i)	(?)		2.255** (5.371)		0.190 (0.086)	0.482 (0.196)
Ln (CROWD_AVG_PURCHASE _{i,q})	(?)			0.034*** (10.485)	0.030** (4.288)	0.041** (5.521)
Ln (SALES _{i,q-4})	(-)	-0.090** (4.969)	-0.089** (4.855)	-0.098** (5.246)	-0.097** (5.081)	-0.096** (4.974)
SAMEQUARTER_SALESGROWTH _{i,q-4}	(+)	0.022 (0.307)	0.025 (0.340)	0.017 (0.174)	0.017 (0.173)	0.015 (0.144)
CHG_BACKLOG _{i,q-1}	(+)	0.378*** (134.895)	0.376*** (130.725)	0.371*** (114.528)	0.372*** (120.871)	0.376*** (120.551)
ADVERTISE _{i,q-1}	(+)	-0.913 (1.783)	-0.921 (1.750)	-1.211 (2.295)	-1.186 (2.102)	0.397 (2.689)
Ln (MEDIA_AVG_PURCHASE _{i,q})*ADVERTISE _{i,q-1}	(?)					-0.053 (0.001)
Ln (EXPERT_AVG_PURCHASE _{i,q})*ADVERTISE _{i,q-1}	(?)					-3.521 (0.061)
Ln (CROWD_AVG_PURCHASE _{i,q})*ADVERTISE _{i,q-1}	(?)					-0.355 (2.559)
Industry fixed effects		Yes	Yes	Yes	Yes	Yes
Number of observations		2,281	2,281	2,281	2,281	2,281
Likelihood ratio		685.9	688.1	681.9	686.6	690.6

Table 7
(continued)

Panel B: The predictive power of POSITIVE by tweet handle

		DEPENDENT VARIABLE = SAMEQUARTER_SALESGROWTH _{i,q}				
	Expected sign	Coefficient (<i>chi</i> -value)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)
Intercept		Included	Included	Included	Included	Included
MEDIA_AVG_POSITIVE_{i,q}	(?)	0.001 (1.172)			0.001 (1.117)	-0.001 (0.636)
EXPERT_AVG_POSITIVE_{i,q}	(?)		0.002 (1.243)		0.002 (0.980)	0.004 (1.901)
CROWD_AVG_POSITIVE_{i,q}	(?)			-0.098 (0.352)	-0.097 (0.394)	-0.414 (2.863)
Ln (SALES _{i,q-4})	(-)	-0.045*** (52.366)	-0.044*** (47.834)	-0.043*** (44.177)	-0.044*** (48.532)	-0.042*** (38.216)
SAMEQUARTER_SALESGROWTH _{i,q-4}	(+)	0.388*** (14.375)	0.392*** (16.300)	0.384*** (15.125)	0.389*** (17.114)	0.374*** (15.936)
CHG_BACKLOG _{i,q-1}	(+)	-0.169*** (11.156)	-0.156*** (13.430)	-0.181*** (13.706)	-0.171*** (20.954)	-0.154*** (32.836)
ADVERTISE _{i,q-1}	(+)	-1.340*** (17.032)	-1.384*** (17.214)	-1.330*** (18.789)	-1.346*** (16.892)	-1.509 (0.084)
MEDIA_AVG_POSITIVE _{i,q} *ADVERTISE _{i,q-1}						0.032 (2.444)
EXPERT_AVG_POSITIVE _{i,q} *ADVERTISE _{i,q-1}						-0.078 (2.267)
CROWD_AVG_POSITIVE _{i,q} *ADVERTISE _{i,q-1}						5.814 (2.834)
Industry fixed effects		Yes	Yes	Yes	Yes	Yes
Number of observations		244	244	244	244	244
Likelihood ratio		18.8	18.8	18.8	22.8	28.7

Table 7
(continued)

Regression results from model (2):

$$\text{SAMEQUARTER_SALESGROWTH}_{i,q} = \alpha + \beta_1 \text{AVG_PURCHASE}_{i,q} + \beta_2 \text{ADVERTISE}_{i,q-1} * \text{AVG_PURCHASE}_{i,q} + \beta_3 \text{AVG_POSITIVE}_{i,q} + \beta_4 \text{ADVERTISE}_{i,q-1} * \text{AVG_POSITIVE}_{i,q} + \beta_5 \ln(\text{SALES}_{i,q-4}) + \beta_6 \text{SAMEQUARTER_SALESGROWTH}_{i,q-4} + \beta_7 \text{CHG_BACKLOG}_{i,q-1} + \beta_8 \text{ADVERTISE}_{i,q-1} + \varepsilon_{it}$$

where MEDIA_AVG_PURCHASE(POSITIVE), EXPERT_AVG_PURCHASE(POSITIVE), and CROWD_AVG_PURCHASE(POSITIVE) replace AVG_PURCHASE(POSITIVE).

PURCHASE is the number of tweets on Twitter that express purchase intent for a company i 's products and brands on a daily basis. POSITIVE is the ratio of the number of positive tweets over the number of non-neutral (positive and negative) tweets about a company i 's products and brands on a daily basis. AVG_PURCHASE $_{i,q}$ averages daily PURCHASE over quarter q . AVG_POSITIVE $_{i,q}$ averages daily POSITIVE over quarter q . SAMEQUARTER_SALESGROWTH $_{i,q}$ is measured as the percentage change in sales in quarter q relative to sales in quarter $q-4$. SALES $_{i,q-4}$ is measured as sales in the same quarter of the previous fiscal year. SAMEQUARTER_SALESGROWTH $_{i,q-4}$ is measured as the same-quarter sales growth in quarter $q-4$. ADVERTISE $_{i,q-1}$ is measured as the ratio of pro-rated advertising expense over sales during the most recent fiscal quarter. CHG_BACKLOG $_{i,q-1}$ is measured as the ratio of the change in deferred revenue over sales during the previous quarter. MEDIA_AVG_PURCHASE $_{i,q}$ is AVG_PURCHASE initiated by the media over quarter q . MEDIA_AVG_POSITIVE $_{i,q}$ is AVG_POSITIVE initiated by the media over quarter q . EXPERT_AVG_PURCHASE $_{i,q}$ is AVG_PURCHASE initiated by product experts over quarter q . EXPERT_AVG_POSITIVE $_{i,q}$ is AVG_POSITIVE initiated by product experts over quarter q . CROWD_AVG_PURCHASE $_{i,q}$ is AVG_PURCHASE initiated by the crowd over quarter q . CROWD_AVG_POSITIVE $_{i,q}$ is AVG_POSITIVE initiated by the crowd over quarter q . Standard errors are clustered at the 6-digit NAICS level.

***Coefficients are significant at 0.01 level. **Coefficients are significant at 0.05 level. *Coefficients are significant at 0.10 level.

Table 8
Unexpected Sales and Product Information on Twitter for the Business-to-Consumer Subsample

		Dependent Variable = UNEXPECTED_SALESGROWTH_{i,q}			
	Expected Sign	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)
Intercept		Included	Included	Included	Included
Ln (AVG_PURCHASE_{i,q})	(+)	0.004* (3.132)		0.003* (2.906)	0.004* (2.993)
AVG_POSITIVE_{i,q}	(?)		-0.036 (2.296)	-0.034 (2.039)	0.007 (0.028)
Ln (NUM_FORECAST _{i,q})	(?)	0.001 (0.001)	0.003 (0.284)	0.001 (0.001)	-0.001 (0.001)
ACTUAL_FORECAST_DAYS _{i,q}	(?)	0.001* (2.807)	0.001 (2.596)	0.001* (2713)	0.001* (3.000)
Ln (SALES _{i,q-1})	(?)	-0.011** (4.612)	-0.01** (5.198)	-0.012** (4.884)	-0.011** (4.240)
SAMEQUARTER_SALESGROWTH _{i,q-4}	(?)	0.003 (0.147)	0.004 (0.254)	0.003 (0.144)	0.003 (0.158)
CHG_BACKLOG _{i,q-1}	(?)	0.087 (0.728)	0.086 (0.707)	0.085 (0.702)	0.081 (0.592)
ADVERTISE _{i,q-1}	(?)	-0.105 (0.509)	-0.069 (0.277)	-0.104 (0.508)	1.107 (0.623)
Ln (AVG_PURCHASE _{i,q})*ADVERTISE _{i,q-1}					-0.024 (0.655)
AVG_POSITIVE _{i,q} *ADVERTISE _{i,q-1}					-1.334 (0.635)
Industry fixed effects		Included	Included	Included	Included
Number of observations		2,281	2,281	2,281	2,281
Likelihood ratio		109.4	109.5	111.4	115.3

Table 8
(continued)

Regression results from model (3):

$$\text{UNEXPECTED_SALESGROWTH}_{i,y+1,q} = \alpha + \beta_1 \text{AVG_PURCHASE (AVG_POSITIVE)}_{i,q} + \beta_2 \text{Ln (NUM_FORECAST}_{i,q}) + \beta_3 \text{ACTUAL_FORECAST_DAYS}_{i,q} + \beta_4 \text{Ln (SALES}_{i,q-1}) + \beta_5 \text{SAMEQUARTER_SALESGROWTH}_{i,q-4} + \beta_6 \text{CHG_BACKLOG}_{i,q-1} + \beta_7 \text{ADVERTISE}_{i,q-1} + \varepsilon_{it}$$

PURCHASE is the number of tweets on Twitter that express purchase intent for a company i 's products and brands on a daily basis. POSITIVE is the ratio of the number of positive tweets over the number of non-neutral (positive and negative) tweets about a company i 's products and brands on a daily basis. AVG_PURCHASE _{i,q} averages daily PURCHASE over quarter q . AVG_POSITIVE _{i,q} averages daily POSITIVE over quarter q . SAMEQUARTER_SALESGROWTH _{i,q} is measured as the percentage change in sales in quarter q relative to sales in quarter $q-4$. B2B is defined as 1 if a firm's major customers are businesses, and 0 if a firm's major customers are consumers. SALES _{$i,q-1$} is measured as sales in the prior quarter. SAMEQUARTER_SALESGROWTH _{$i,q-4$} is measured as the same-quarter sales growth in quarter $q-4$. ADVERTISE _{$i,q-1$} is measured as the ratio of pro-rated quarterly advertising expense over sales during the most recent fiscal quarter. CHG_BACKLOG _{$i,q-1$} is measured as the ratio of the change in deferred revenue over sales during the previous quarter. NUM_FORECAST _{i,q} is the number of forecasts included in the consensus forecast for sales in quarter q . ACTUAL_FORECAST_DAYS _{i,q} is measured as number of calendar days between the report date of realized sales in quarter q and the report date of the consensus forecast. UNEXPECTED_SALESGROWTH _{i,q} is measured as the difference between sales in quarter q and the consensus analyst forecast for sales in quarter q divided by sales in the same quarter of the previous fiscal year (SALES _{$i,q-4$}). Standard errors are clustered at the 6-digit NAICS level.

***Coefficients are significant at 0.01 level. **Coefficients are significant at 0.05 level. *Coefficients are significant at 0.10 level.

Table 9
Unexpected Sales and Product Information on Twitter by Tweet Handle

Panel A: The predictive power of PURCHASE by tweet handle with respect to unexpected sales

	DEPENDENT VARIABLE = UNEXPECTED_SALESGROWTH _{i,q}					
	Expected Sign	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)	Coefficient (<i>chi</i> -square)
Intercept		Included	Included	Included	Included	Included
Ln (MEDIA_AVG_PURCHASE _{i,q})	(?)	0.048*** (8.374)			0.017 (0.271)	0.073 (2.425)
Ln (EXPERT_AVG_PURCHASE _{i,q})	(?)		0.112 (0.822)		-0.083 (1.229)	-0.121 (0.639)
Ln (CROWD_AVG_PURCHASE)	(?)			0.004* (3.138)	0.003 (1.625)	0.002 (0.473)
Ln (NUM_FORECAST _{i,q})	(?)	0.002 (0.110)	0.003 (0.222)	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)
ACTUAL_FORECAST_DAYS _{i,q}	(?)	0.001* (2.720)	0.001* (2.704)	0.001* (2.807)	0.001* (2.771)	0.001* (3.007)
Ln (SALES _{i,q-1})	(?)	-0.011** (4.96)	-0.011** (4.857)	-0.011** (4.605)	-0.013** (4.436)	-0.011** (4.389)
SAMEQUARTER_SALESGROWTH _{i,q-1}	(?)	0.003 (0.219)	0.003 (0.246)	0.003 (0.145)	0.003 (0.147)	0.003 (0.139)
CHG_BACKLOG _{i,q-1}	(?)	0.087 (0.736)	0.087 (0.734)	0.087 (0.728)	0.086 (0.726)	0.086 (0.711)
ADVERTISE _{i,q1}	(?)	-0.071 (0.306)	-0.071 (0.281)	-0.106 (0.515)	-0.104 (0.487)	-0.200 (0.848)
ADVERTISE _{i,q-1} * Ln (MEDIA_AVG_PURCHASE _{i,q})	(?)					-1.648 (1.911)
ADVERTISE _{i,q-1} * Ln (EXPERT_AVG_PURCHASE _{i,q})	(?)					0.687 (0.105)
ADVERTISE _{i,q-1} * Ln (CROWD_AVG_PURCHASE _{i,q})	(?)					0.027 (0.148)
Industry fixed effects		Included	Included	Included	Included	Included
Number of observations		2,281	2,281	2,281	2,281	2,281
Likelihood ratio		111.5	109.4	109.4	113.4	119.4

Table 9 (continued)

Panel B: The predictive power of POSITIVE by tweet handle with respect to unexpected sales

	DEPENDENT VARIABLE = UNEXPECTED_SALESGROWTH _{i,q}					
	Expected Sign	Coefficient (<i>chi-square</i>)	Coefficient (<i>chi-square</i>)	Coefficient (<i>chi-square</i>)	Coefficient (<i>chi-square</i>)	Coefficient (<i>chi-square</i>)
Intercept		Included	Included	Included	Included	Included
MEDIA_AVG_POSITIVE _{i,q}	(?)	0.001 (0.611)			0.001 (0.208)	-0.003* (3.807)
EXPERT_AVG_POSITIVE _{i,q}	(?)		0.005*** (27.627)		0.004*** (30.394)	0.007*** (26.190)
CROWD_AVG_POSITIVE _{i,q}	(?)			-0.244 (1.617)	-0.217 (1.480)	-0.188 (1.345)
Ln (NUM_FORECAST _{i,q})	(?)	0.018* (3.620)	0.021** (4.033)	0.014* (3.403)	0.017** (4.027)	0.021** (4.179)
ACTUAL_FORECAST_DAYS _{i,q}	(?)	0.003 (0.928)	0.003 (1.015)	0.002 (0.581)	0.002 (0.934)	0.002 (0.795)
Ln (SALES _{i,q-1})	(?)	0.006 (1.017)	0.006 (0.800)	0.008 (0.996)	0.007 (0.794)	0.009 (0.819)
SAMEQUARTER_SALESGROWTH _{i,q-4}	(?)	0.126* (2.811)	0.137** (4.832)	0.117* (2.803)	0.130** (4.589)	0.126* (4.225)
CHG_BACKLOG _{i,q-1}	(?)	2.741** (4.290)	2.773** (4.585)	2.712** (4.288)	2.743** (4.589)	2.747** (4.778)
ADVERTISE _{i,q-1}	(?)	0.425 (1.048)	0.327 (0.772)	0.487 (1.234)	0.399 (0.781)	0.327 (0.772)
ADVERTISE _{i,q-1} * MEDIA_AVG_POSITIVE _{i,q}	(?)					0.079* (3.091)
ADVERTISE _{i,q-1} * EXPERT_AVG_POSITIVE _{i,q}	(?)					-0.098** (6.335)
ADVERTISE _{i,q-1} * CROWD_AVG_POSITIVE _{i,q}	(?)					-2.866 (0.425)
Industry fixed effects		Included	Included	Included	Included	Included
Number of observations		244	244	244	244	244
Likelihood ratio		44.9	44.6	44.8	48.5	54.3

Table 9
(continued)

Regression results from model (3):

$$\text{UNEXPECTED_SALESGROWTH}_{i,q} = \alpha + \beta_1 \text{AVG_PURCHASE (AVG_POSITIVE)}_{i,q} + \beta_2 \text{Ln (NUM_FORECAST}_{i,q}) + \beta_3 \text{ACTUAL_FORECAST_DAYS}_{i,q} + \beta_4 \text{Ln (SALES}_{i,q-1}) + \beta_5 \text{SAMEQUARTER_SALESGROWTH}_{i,q-4} + \beta_6 \text{CHG_BACKLOG}_{i,q-1} + \beta_7 \text{ADVERTISE}_{i,y-1} + \epsilon_{it}$$

where MEDIA_AVG_PURCHASE(POSITIVE), EXPERT_AVG_PURCHASE(POSITIVE), and CROWD_AVG_PURCHASE(POSITIVE) replace AVG_PURCHASE(POSITIVE).

PURCHASE is the number of tweets on Twitter that express purchase intent for a company i 's products and brands on a daily basis. POSITIVE is the ratio of the number of positive tweets over the number of non-neutral (positive and negative) tweets about a company i 's products and brands on a daily basis. AVG_PURCHASE $_{i,q}$ averages daily PURCHASE over quarter q . AVG_POSITIVE $_{i,q}$ averages daily POSITIVE over quarter q . MEDIA_AVG_PURCHASE $_{i,q}$ is AVG_PURCHASE initiated by the media over quarter q . MEDIA_AVG_POSITIVE $_{i,q}$ is AVG_POSITIVE initiated by the media over quarter q . EXPERT_AVG_PURCHASE $_{i,q}$ is AVG_PURCHASE initiated by product experts over quarter q . EXPERT_AVG_POSITIVE $_{i,q}$ is AVG_POSITIVE initiated by product experts over quarter q . CROWD_AVG_PURCHASE $_{i,q}$ is AVG_PURCHASE initiated by the crowd over quarter q . CROWD_AVG_POSITIVE $_{i,q}$ is AVG_POSITIVE initiated by the crowd over quarter q . SALES $_{i,q-1}$ is measured as sales in the prior quarter. SAMEQUARTER_SALESGROWTH $_{i,q-4}$ is measured as the same-quarter sales growth in quarter $q-4$. ADVERTISE $_{i,q-1}$ is measured as the ratio of prorated quarterly advertising expense over sales during the most recent fiscal quarter. CHG_BACKLOG $_{i,q-1}$ is measured as the ratio of the change in deferred revenue over sales during the previous quarter. NUM_FORECAST $_{i,q}$ is the number of forecasts included in the consensus forecast for sales in quarter q . ACTUAL_FORECAST_DAYS $_{i,q}$ is measured as number of calendar days between the report date of realized sales in quarter q and the report date of the consensus forecast. UNEXPECTED_SALESGROWTH $_{i,q}$ is measured as the difference between sales in quarter q and the consensus analyst forecast for sales in quarter q divided by sales in the same quarter of the previous fiscal year (SALES $_{i,q-4}$). Standard errors are clustered at the 6-digit NAICS level.

***Coefficients are significant at 0.01 level. **Coefficients are significant at 0.05 level. *Coefficients are significant at 0.10 level.