https://doi.org/10.1016/j.jacceco.2020.101334

Can social media distort price discovery? Evidence from merger rumors

Weishi Jia
Cleveland State University
w.jia1@csuohio.edu

Giulia Redigolo
ESADE Business School
giulia.redigolo@esade.edu

Susan Shu*
Carroll School of Management
Boston College
shus@bc.edu

Jingran Zhao
Hong Kong Polytechnic University
jingran.zhao@polyu.edu.hk

*Corresponding Author. Carroll School of Management, Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA 02467. Phone number: (617)552-1759. Email: susan.shu@bc.edu

We thank Mengyao Cheng, Ilia Dichev, Carlo Gallimberti, Amy Hutton, Lian Fen Lee, Alvis Lo, Michelle Lowry, Stan Markov, Jeff Ng, Grace Pownall, Sugata Roychowdhury, David Solomon and workshop participants at Boston College, CUHK, Emory University, University of Hong Kong and University of Rochester for valuable feedback and comments.

Can social media distort price discovery? Evidence from merger rumors

Abstract

We study whether social media can play a negative information role by impeding price discovery in the presence of highly speculative rumors. We focus on merger rumors, where most do not materialize. We find that merger rumors accompanied by greater Twitter activity elicit greater immediate market reaction even though rumor-related Twitter activity is unrelated to the probability of merger realization. The price distortion associated with tweet volume persists weeks after a rumor and reverses only after eight weeks. The price distortion is more pronounced for rumors tweeted by Twitter users with greater social influence, for target firms with low institutional ownership, and for rumors that supply more details. Our evidence suggests that social media can be a rumor mill that hinders the market's price discovery of potentially false information.

JEL classification: G14; M15; M40; M41

Keywords: social media; merger and acquisition; rumor; Twitter

1. Introduction

One of the fundamental questions in financial economics is how information affects the capital market. The rise of social media has drastically reshaped the ways information is created, disseminated, and consumed (Miller and Skinner, 2015).[1] Social media is distinguished from traditional media in its user-generated content, its speed, reach, and network effects, among other things. These particular features suggest that social media can play a distinct role from traditional media. Within the financial markets, social media has been shown to facilitate the aggregation of individual opinions and improve price efficiency via the wisdom of the crowd (Chen et al., 2014; Bartov et al., 2017; Tang, 2018), reduce information asymmetry among investors and mitigate adverse market reaction to product crises (Blankespoor et al., 2014; Lee et al., 2015). However, the very features that facilitate information discovery can also turn social media into a rumor mill whereby false information proliferates and distorts price discovery in the capital market. This negative information role of social media is largely unexplored in the literature. In this study, we shed light on the downside of social media as an information channel by examining whether social media activity impedes price discovery in the face of potentially false rumors.

Although social media's rumor mill role is unexplored in the academic literature, anecdotal evidence has long highlighted the negative effect of social media. Many have expressed the concern that rumors could "outrace the truth" (Nyhan, 2014). The Securities and Exchange Commission (SEC) has been so concerned about the role of social media in spreading false information that it has issued multiple investor alerts.[2] Such negative views of social media

[1] According to a 2017 survey from Pew Research Center, two-thirds of Americans now get at least some of their news on social media. http://www.journalism.org/2017/09/07/news-use-across-social-media-platforms-2017/
[2] https://www.sec.gov/oiea/investor-alerts-bulletins/ia_rumors.html

are widespread despite the lack of systematic evidence. To fill the gap in the literature, we explicitly explore whether individual actions in spreading potentially false rumors, via social media, can distort, at least temporarily, price discovery.

Social media has several key features for false information to spread and take hold, and at least temporarily, distort prices. First, social media widens the reach of false information, essentially lowering the cost of obtaining such information by capacity-constrained investors.[3] Investors' response to what they perceive as valid information can temporarily distort stock prices. Second, the network feature of social media exacerbates investors' *persuasion bias,* that is, the failure to adjust appropriately for the repetitions in the information one receives when assessing its validity (DeMarzo et al., 2003).[4] Persuasion bias applies "not only to information coming from one source over time but also to information coming from multiple sources connected through a social network." (DeMarzo et al., 2003). Simply put, repeated encounters of the same information in one's social media network can lead to an inflated assessment of its validity and temporarily distort price discovery.[5] Thus, investors with bounded rationality may incorrectly impound social media activity in their response to rumors.

However, the price distortion associated with the rumor mill effect may be difficult to detect. First, it can be dominated by the crowd wisdom effect. Prior literature shows that social

---

[3]Capacity constraint means processing capacity is a scarce resource that investors have to allocate across various activities. Because information processing is costly, investors have to decide what information to acquire, analyze, and incorporate into trading decisions (see Blankespoor et al. 2020).

[4] Repetition-induced bias has been shown to influence decision-making in various experiments (e.g., Hawkins and Hoch, 1992; Gilbert, 1991). DeMarzo et al. (2003) argue that persuasion bias can be "viewed as a simple, boundedly rational heuristic for dealing with a very complicated inference problem," as "agents cannot determine (or recall) the source of all the information that has played a role in forming their beliefs."

[5] We do not attempt to empirically distinguish the two different sources of price distortions because both can lead to the rumor mill effect.

media aggregates individual wisdom into the wisdom of the crowd, which predicts future sales and earnings and facilitates price discovery (Chen et al., 2014; Bartov et al., 2017; Tang, 2018). Second, investors may ignore rumors (especially false ones) on social media, in which case social media activity would likely have no bearing on price discovery in our empirical setting. Consistent with this view, Guess et al. (2018) show that false news on social media, while having a broad reach, only affected a small segment of the population during the 2016 presidential election. Ultimately, whether we will find evidence of the rumor mill effect of social media in the price discovery process is an empirical question.

We use Twitter, one of the most influential social media platforms, to measure social media activity. For financial rumors, we focus on a sample of 304 rumors related to mergers and acquisitions that surface in the media from 2009 to 2014. Prior literature has highlighted traditional media's incentives to publish sensational but potentially untrue stories, such as merger rumors, to attract readership (e.g., Ahern and Sosyura, 2015).[6] These rumors are highly speculative, with only 10%–30% of rumors materializing (Ahern and Sosyura, 2015; Ma and Zhang, 2016). Using detailed Twitter data, we can explore whether and how individual engagement in terms of sharing, discussing, and spreading the rumors (via social media) can influence price discovery.[7]

---

[6] Appendix B provides two examples of merger rumors. The first example, "Boeing explores purchase of Mercury Systems," appeared in Reuters, and was an "exclusive" story based on two anonymous sources. In the second example, "Should VF acquire Columbia Sportswear," the author is merely expressing his or her opinions.

[7] We focus on understanding the role that individuals play in stock price formation via social media by spreading potentially false rumors to networks of users, potentially distorting, at least temporarily, price discovery. Ahern and Sosyura (2015), on the other hand, focus on the incentives of newspapers to generate potentially untrue stories to attract readership. Although they show that investors underestimate newspapers' incentives to sensationalize and overreact to the rumors, they do not study *how* these rumors lead to market over-reaction. One such channel could be individuals spreading merger rumors through their circles, leading to temporary price distortions driven by, for instance, individuals' persuasion bias. Ahern and Sosyura (2015) are neither interested in nor equipped to address

To capture social media activity related to a merger rumor, we use the volume of rumor-related tweets in the rumor announcement window as our primary construct. The rumor mill effect is predicated on social media distorting investors' information processing through both its reach and its network features. The volume of rumor-related tweets captures the tendency for such distortive effect to be significant: highly tweeted merger rumors can spread to more investors, and they are more likely to be seen repeatedly through one's social network. Both circumstances can lead to an inflated assessment of rumor validity by investors and potentially distort price discovery.[8]

We start by examining whether rumor-day Twitter activity helps predict the accuracy of a merger rumor (i.e., whether it materializes). If more "valid" rumors attract more tweets, then tweet volume should be positively related to rumor accuracy. If social media posts are motivated not by rumor validity but by users' desire to share attention-grabbing stories, then rumor-day tweet volume would be unrelated to rumor accuracy. To ascertain the rumor date for each merger rumor, we identify the original "scoop" article that first reported a merger rumor. We find that rumor-day abnormal tweet volume does not predict rumor accuracy, meaning that highly tweeted rumors are not more likely to be accurate compared with rumors that attract low tweet volume. This finding is consistent with the interpretation that social media posts are motivated not by the validity of merger rumors but by users' desire to share attention-grabbing stories.

---

this question. By exploring Twitter activity around the rumors, we can examine the role of individual actions via social media in price discovery.

[8] We note that our use of tweet volume captures key aspects of the rumor mill effect. In our setting, the role of crowd wisdom is more limited because the vast majority of these merger rumor related tweets (over 70%) are neutral in sentiment (see Appendix C). Focusing on tweet volume increases the power of the tests on the rumor mill effect, which, unlike the crowd wisdom effect, has not been documented in the literature.

Next, we examine whether the target's stock price reaction to the merger rumors is related to the rumor-window Twitter activity. We first sort the rumored targets into high/low portfolios based on rumor-day abnormal user tweet volume. The abnormal return for the event period [0, +1] is 7.57% for the high portfolio, in contrast to 2.87% for the low portfolio. Regressions of rumor-window abnormal returns on abnormal tweet volume yield positive and significant coefficients, implying that highly tweeted rumors elicit a greater market reaction. In terms of economic significance, an increase in tweet volume by one standard deviation translates to an abnormal return of 0.29%. Despite the insignificant relation between tweet volume and rumor accuracy, investors behave as if a higher rumor-window tweet volume implies a higher likelihood of a deal.

Since rumor-window tweet volume is insignificantly related to rumor accuracy, the heightened market reaction to highly tweeted merger rumors is suggestive of the rumor mill effect of social media.[9] To solidify our inference, we further examine the price evolution around merger rumors that fail to materialize. A merger rumor fails to materialize because either (1) it is pure fabrication, or (2) the merger negotiations are initiated but fall apart. Because the eventual outcome is unknown at the time of the rumor, the price reaction reflects investors' collective assessment of the likelihood of deal realization. For inaccurate rumors, investors will eventually

---

[9] While tweet volume is unrelated to deal likelihood, we cannot rule out the possibility that it is related to perceived merger gains conditional on deal completion. To mitigate this concern, we perform several robustness tests. First, we include additional control variables shown to be related to potential merger gains, and our main inferences are unchanged. Second, we examine whether abnormal tweet volume is related to potential merger gains for the 63 realized rumors. To capture potential merger gains, we sum up the abnormal returns from two different windows: the actual takeover announcement and the merger rumor announcement. We find that rumor-day tweet volume is not related to this measure of merger gains. To the extent that perceived merger gains constitute a correlated omitted variable, our instrumental variables approach can hopefully further mitigate this concern.

recognize the low probability of a deal, at which point the stock prices should revert to pre-rumor levels. Price evolution around rumors, including post-rumor reversals, sheds light on the role of social media in the price discovery of (ex-post) false rumors.

Even among rumors that fail to materialize, the immediate market reaction is positively and significantly related to the rumor-day tweet volume, consistent with the overall sample results. In addition, the price distortion driven by tweet volume persists for weeks, and price correction occurs only after eight weeks. The heightened immediate response to highly tweeted rumors, combined with post-rumor reversals, is consistent with price distortion being driven by investors' tendency to attribute higher validity (than warranted) to highly tweeted merger rumors.

We explore three sources of cross-sectional variation in the rumor mill effect: the social influence of those who tweeted about the merger rumor, the investor base of the rumored target, and the apparent credibility of a rumor. First, we find that the distorted reaction to highly tweeted merger rumors is more pronounced for rumors tweeted by "influential" Twitter users. Because high-influence Twitter users have more followers and are more active, their tweets have greater reach and a higher chance of repeated exposure. Besides, information tweeted by influential posters is also likely to be viewed as more credible by capacity-constrained investors. Second, the rumor mill effect is more pronounced among firms with low institutional ownership. These firms have a larger base of retail investors who are perhaps more susceptible to social media hype. In addition, firms with low institutional ownership are more likely to have binding short-sale constraints, which limit the arbitrage by sophisticated investors (Nagel, 2005). Third, the rumor mill effect is more pronounced for rumors that come across as more definitive based on rumor details such as purported talk stage, potential bidder name, and deal price.

While we aim to establish a causal link between Twitter activity and abnormal market reaction, our inference that Twitter activity *affects* market reaction is confounded by two important endogeneity concerns: reverse causality (i.e., abnormal returns *causing* abnormal tweets) and unobservable correlated variables that are omitted (e.g., perceived rumor validity or anticipated merger gains driving both investor reaction and Twitter activity). The complicated relationship between tweet volume and abnormal returns is illustrated in Figure 1. We perform a battery of tests to strengthen our causal inferences.

First, we analyze two separate samples of rumors that surfaced during non-trading windows and examine the relationship between the abnormal tweets during the non-trading windows (as users can still tweet) and the *subsequent* market reaction. The first sample consists of rumors that surfaced during weekends or holidays. The second sample consists of rumors that surfaced either before the stock market opened or after the market closed. In both cases, we find a positive and significant relationship between the abnormal tweet volume, measured in the non-trading windows, and the abnormal returns in the next trading day. This analysis helps alleviate the concern of reverse causality because the tweets in the non-trading windows precede and thus are by construction unaffected by the subsequent price movements.[10]

Second, we employ an instrumental variable approach, which can mitigate both correlated omitted variables and reverse causality problems. We use two instruments, Twitter

[10] Although investors can still trade during extended market hours, this market is notoriously thin, with wide bid-ask spreads and low likelihood of order execution. The effect of the sporadic returns on Twitter activity is likely limited.

outage and the presence of "viral" tweets, in two-stage least square (2SLS) regressions. A Twitter outage is a plausible instrument because it captures exogenous variations in Twitter usage, yet it is driven entirely by technical issues unrelated to stock returns. Similarly, we argue that viral tweets likely distract some Twitter users from merger rumors and therefore reduce the circulation of these rumors, but they are otherwise unlikely to affect the market reaction to merger rumors.[11] As expected, when a merger rumor coincides with a Twitter outage or a viral tweet, the volume of rumor-related tweets decreases significantly. Either instrument is highly significant in the first stage and meets conventional criteria for instrument relevance. Importantly, the instrumented tweet volume continues to be positive and significant in the second stage, indicating that rumor-related tweet volume leads to higher abnormal returns. Collectively, these tests strengthen the causal link between tweets and returns. Although some of these tests are based on relatively small sample sizes, we are reassured by the consistent inferences across multiple approaches.

Our evidence of the rumor mill effect in the presence of highly speculative rumors does not contradict the crowd wisdom effect because we specifically design our tests to maximize the chances of detecting the rumor mill effect. Within the subsample of merger rumors that materialized, we find weak evidence that the rumor-day tweets accelerate the price discovery of merger rumors: rumors accompanied by a higher number of tweets experience higher rumor-day cumulative abnormal returns (CARs) yet lower takeover announcement CARs. We do not draw

[11] To identify viral tweets, we rely on a variety of publications that rank top tweets by year. These viral tweets are unrelated to financial or economic news. They are most frequently posted by celebrities (e.g., Justin Bieber and Kanye West) about their personal and professional life (see Appendix D for the list).

strong inferences because realized rumors only represent 21% of the rumor sample, and the crowd wisdom effect of social media has been well documented.

We provide the first systematic evidence on the rumor mill effect of social media. As an emerging and increasingly important information channel, social media has unique features that are distinct from traditional media. By providing a window into individual opinions and actions, social media platforms enable researchers to study the role of individuals in influencing price discovery. Prior literature documents various benefits of social media, both as a corporate disclosure channel (Blankespoor et al., 2014; Lee et al., 2015) and as an information platform that efficiently aggregates individual wisdom (Chen et al., 2014; Bartov et al., 2017; Tang, 2018). We highlight a potential downside of social media: in the face of highly speculative financial rumors, social media facilitates the spread of rumors and, at least temporarily, distorts price discovery. While we focus on merger rumors, our finding potentially extends to other corporate rumors such as those about management turnover, FDA and patent approvals, or new product launches. In an age when fake news can dominate the headlines, and an increasing number of readers get their news from social media, our evidence is of interest to academics, practitioners, and regulators.

Additionally, our paper extends the literature that examines media's incentives to provide biased coverage to cater to their readers, attract new readership, and generate advising revenue (e.g., Core et al., 2008; Gentzkow and Shapiro, 2010; Ahern and Sosyura, 2015; Gurun and Buter, 2012). For example, Ahern and Sosyura (2015) document that media outlets publish potentially false merger rumors to sell more newspapers, yet investors underestimate newspapers' incentives to sensationalize. When these potentially false stories surface in the media, individuals can play a role, via social media, in the price discovery process. Social media activity could help investors ascertain the validity of inaccurate media stories vis-à-vis the wisdom of the crowd. Or, it could also amplify the negatives of traditional media by increasing the exposure to false information. Our evidence highlights the interaction between social media and traditional media in generating and spreading false information, which furthers our understanding of the role of both information channels in the capital market.

2. Related literature and hypothesis development

Prior studies have documented various informational benefits associated with social media. One strand of the literature focuses on how social media helps disseminate corporate information. Blankespoor et al. (2014) find that the use of Twitter to disseminate corporate news reduces information asymmetry among investors and increases market liquidity, although Jung et al. (2018) document that firms strategically avoid disseminating bad news on social media. Lee et al. (2015) find that corporate social media attenuates the negative price reaction to product recalls. However, Lee et al. (2015) also allude to a potential downside of corporate use of social media, as social media accounts can facilitate the spread of negative sentiment about the firm and its products during product recalls.

Another strand of the literature focuses on the useful and predictive role of social media. Because social media provides a platform to efficiently aggregate individual wisdom, information on social media has the potential to predict future firm performance and improve stock price formation. Chen et al. (2014) find that views expressed on an investment-related website, *Seeking Alpha*, predict future stock returns and earnings surprises. Similarly, Bartov et al. (2017) find that the aggregate opinion from individual tweets can predict a firm's future earnings and announcement returns, after controlling for the role of media. Tang (2018) finds that customers' opinions posted on Twitter predict future firm sales growth.

Despite the many benefits documented by prior studies, social media can also pose an information challenge for market participants. Information on social media is unfiltered, and social media users are motivated to post sensational stories. Toubia and Stephen (2013) find that social media users are driven by image-related utility. They post content to increase their recognition from others in the form of the number of followers, likes, and retweets, regardless of the content's accuracy. To explore the negative information consequences of social media, we focus on the *rumor mill* effect of social media in the face of highly speculative merger rumors.

Our first hypothesis explores whether Twitter activity helps predict the accuracy of the rumor (i.e., whether it will materialize). If more "valid" rumors attract more tweets, then tweet volume should be positively related to rumor accuracy. On the other hand, if social media posts are motivated not by rumor validity but by users' desire to post and share attention-grabbing stories, then rumor-window tweet volume would have no bearing on rumor accuracy. Our first hypothesis, stated in its null form, is as follows:

H1: Twitter activity related to a merger rumor does not predict the accuracy of the rumor, that is, whether the rumor will materialize.

Our next hypothesis is related to how investors impound rumor-related information on Twitter. Twitter activity can distort information processing by capacity-constrained investors with bounded rationality and, at least temporarily, induce higher stock returns. First, social media widens the reach of information, even false information, essentially lowering investors' acquisition cost of these false rumors and inducing buying pressure that can temporarily increase stock prices. Second, the network feature of social media exacerbates investors' persuasion bias, that is, the tendency to overweigh information that one repeatedly encounters (DeMarzo et al., 2003; Hawkins and Hoch, 1992; Gilbert, 1991). DeMarzo et al. (2003) point out that the failure to adjust for repetition applies to many situations, including information "coming from multiple sources connected through a social network." This bias is particularly relevant for social media because users are connected through a network and are likely to repeatedly encounter the same (even false) information, leading to an inflated assessment of its validity. The association between market reaction and rumor-window tweeting activity can also be driven by rational expectations, for example, if tweet volume indicates a higher deal likelihood. Our second hypothesis, stated in the null form, is as follows:

H2: The market reaction to the announcement of a merger rumor is unrelated to the Twitter activity in the rumor window.

The first two hypotheses, H1 and H2, jointly shed light on the rumor mill role of social media. Specifically, if rumor period tweets are unrelated to deal realization, yet the market nevertheless reacts to heightened Twitter activity, the findings will point to the rumor mill effect of social media.

To solidify the rumor mill role of social media in the price discovery process, we turn to price evolution around inaccurate rumors (i.e., those that ultimately fail to materialize). Since the merger outcome is not observable at the time of a rumor, investors form expectations of deal likelihood based on their assessment of rumor validity. If, however, investors attribute higher validity (than warranted) to highly tweeted yet inaccurate merger rumors (due to persuasion bias or widened rumor reach), the immediate market reaction will be distorted, and the information discovery process prolonged. Our third hypothesis, stated in its null form, is as follows:

H3: The price evolution around inaccurate rumors, both during the event window and in the post periods, is unrelated to the Twitter activity in the rumor window.

3. Data and descriptive statistics

3.1. Sample selection

We follow a two-step process to identify rumors that surfaced in the media. First, we compile an initial list of merger rumors by searching RavenPack News Analytics (*Dow Jones Edition and Web Edition*)[12] for all stories tagged under the TOPIC: "Business"; GROUP: "Acquisition-mergers"; TYPE: "Acquisition"; SUB_TYPE: "Rumor".[13] We retain rumors related to target companies. This process yields 590 rumors from 2008 to 2014. To expand our sample, we also add 70 stories in Ahern and Sosyura (2015) but not in our initial rumor sample.[14] In RavenPack, stories covering the same "event" receive the same event key, and we only keep the earliest event in a sequence. When multiple rumor events about the same target surface within 60 days of each other, we only keep the earliest one to avoid confounding effects.[15] After

[12] RavenPack retrieves and analyzes news articles from three sources: (1) *Dow Jones Edition*: Dow Jones Newswires, Wall Street Journal, and Barron's; (2) *Web Edition*: business publishers, national and local news, blog sites, government, and regulatory updates; and (3) *PR Edition*: press releases, regulatory, corporate, and news services. We only use the *Dow Jones Edition* and *Web Edition*. RavenPack has been widely used in finance and accounting studies (e.g., Kolasinski et al., 2013; Dang et al., 2015; Dai et al., 2015; Massa et al., 2015).

[13] In RavenPack, relevant stories about entities are classified into a set of predefined event categories following the RavenPack taxonomy. The taxonomy is as follows: TOPIC is a subject or theme of events, and it is the highest level of the RavenPack Taxonomy; GROUP is a collection of related events; TYPE is a class of events, the constituents of which share similar characteristics; and SUB_TYPE is a subdivision of a particular class of events.

[14] We thank Kenneth Ahern for generously providing the rumor data used in Ahern and Sosyura (2015).

[15] We also ran robustness checks, excluding six rumors in our final sample that were followed by another merger rumor within 60 days. The main inferences are unchanged.

excluding rumors related to non-U.S. targets, foreign subsidiaries of U.S. firms, or private firms, we have a total of 474 rumors remaining.

Next, we identify scoop articles, defined as those that first reported a specific merger rumor. We carefully read each story and use Factiva (*All Sources*) to follow the citation trail until we identify the original story that does not cite another publication. This story is identified as the scoop article, and the reported date is identified as the rumor day. If there is any ambiguity regarding the source article, we supplement it with a Google search. We define a "rumor" as "unverified news." That is, a merger talk/deal is reported, but no official confirmation is found. We remove 28 misclassified rumors (e.g., a takeover bid is already confirmed).

Our use of the citation trail in Factiva to identify scoop articles closely follows Ahern and Sosyura (2015). However, our initial sample is based on RavenPack (*Dow Jones Edition* and *Web Edition*), while Ahern and Sosyura (2015) use Factiva (but limited to major news and business sources) to retrieve the initial rumor sample. One advantage of RavenPack is that the database covers a broader range of publications, including online blog sites and trade journals, compared with Factiva (major news and business sources). Our sampling choice likely results in more speculative rumors, which is well suited for studying the rumor mill effect.[16]

To collect Twitter data, we use Crimson Hexagon, a web-based platform that provides access to full Twitter Firehose (all public tweets). For each rumored target, we search the database and download all public Tweets in specified windows. Consistent with prior literature,

[16] In our final sample (with all available data), a total of 29 rumors are exclusively from the Ahern and Sosyura (2015) sample. When we compare the 29 rumors with the rest of our sample for the years 2009–2011 (to match their time period), we find that the rumors from Ahern and Sosyura (2015) are indeed more accurate, are published by more reputable outlets, and are significantly more likely to report bidder and price information. In robustness checks, we exclude these 29 rumors from the sample, and our main inferences remain unchanged.

we use “$” and the stock symbol when we search for tweets to minimize misclassification. For example, the search query for Apple Inc. is “$AAPL.” Appendix C presents examples of tweets related to merger rumors. After removing rumors in 2008 due to insufficient Crimson Hexagon coverage (33), rumors missing Compustat /CRSP data (61), and observations with insufficient data to compute the main Twitter activity measure (48), our final sample consists of 304 rumors from 2009 to 2014. We detail the sample selection in Table 1.[17]

Table 2, Panel A, presents the distribution of the merger rumors by year.[18] While there are more rumors in 2010 and 2014, there is no obvious time-series concentration. We follow Ahern and Sosyura (2015) in determining the rumor realization. Realized rumors are defined as cases in which a proposed takeover of the rumored target is publicly announced within one year of the rumor date, whether or not the deal was ultimately completed. Proposed takeovers are identified in the SDC Platinum database based on the following criteria: (1) the target is a public U.S. firm, (2) the status of the deal is “Completed,” “Pending,” or “Withdrawn,” and (3) the form of the deal excludes “Buyback,” “Exchange offers,” and “Recapitalization.” Unrealized rumors are those not receiving a takeover offer within a year. The accuracy rate of our sample is roughly 21%, suggesting that the vast majority of the rumors do not materialize.[19]

---

[17] Our main measure of rumor-window Twitter activity is standardized by the time-series standard deviation of “normal” tweet volume. If a target firm experiences zero Twitter activity throughout the control window (defined later), we will not be able to calculate the time-series standard deviation needed for our main measure of Twitter activity. This requirement resulted in a loss of 48 observations.

[18] Our rumor sample starts in the year 2009 to ensure reasonable Twitter coverage. Twitter data become available to download starting from May 2008, but the Twitter coverage is relatively poor for 2008.

[19] A few reasons likely account for the difference in our accuracy rate compared with previous papers. The first is the initial data source. Our initial rumor sample is from RavenPack (*Dow Jones Edition* and *Web Edition*), which covers a variety of publications. Using Factiva (and limiting to major publications) for their initial sample, Ahern and Sosyura (2015) report an accuracy rate of 33%. Using Capital IQ, Ma and Zhang (2016) report an accuracy rate of 10.6%. Second, the sample period varies across studies. The sample period is 2000–2011 for Ahern and Sosyura (2015) and 2005–2011 for Ma and Zhang (2016). In contrast, our sample period is 2009–2014.

3.2. Measurement of variables and descriptive statistics

3.2.1. Rumor article characteristics

From the scoop articles, we manually collect detailed rumor information. The rumors vary in reported details, and we present the rumor characteristics in Table 2, Panel B and part of Panel C. As shown in Panel B, more than half of the rumors are in the "speculation" stage (54.61%), without any mention of the status, and another 17.43% are pure opinion pieces (e.g., "Should GE buy Boeing?"). Among the remaining rumors, the purported merger talks range from preliminary to advanced.[20] We define a categorical variable, *SPECULATIVE*, that is equal to 2 if the merger stage is classified as "opinion pieces," 1 if the merger stage is speculation (i.e., no mention of merger talk stages), and 0 otherwise. In addition, some rumors report the rumored bidders or even deal prices, while others contain no such specific information. As shown in Panel C of article characteristic variables, potential bidders are mentioned in 73% of the rumors, and deal prices are mentioned in 16%.

Based on the various rumor details, we construct a binary variable *SPECIFIC*, which is equal to 1 for rumors that appear most definitive: both the takeover prices and the potential bidders are mentioned, and the merger talks are rumored to be in advanced stages (i.e., "made offer," "evaluating bid," or "in advanced talk"). It is equal to 0 otherwise.

The publications that first reported these rumors can be major newspapers or magazines, such as the *Wall Street Journal*, *New York Times*, or *Forbes*, or trade magazines in retail and

[20] Similar to Ahern and Sosyura (2015), the classification includes the following stages: (i) preliminary talk, (ii) in talk, (iii) made offer, (iv) preparing bid, (v) for sale, (vi) evaluating bid, and (vii) speculation. We also add two additional categories for rumors that are reportedly in advanced talk and rumors that appear in the form of individual opinions. We label them respectively (viii) "in advanced talk" and (ix) "opinion pieces."

technology. They also include blogs associated with print newspapers, such as *Dealbook* by the *New York Times*, and online investor sites such as *Seeking Alpha*, *The Street*, or *TechCrunch*. We track and code the outlet for each scoop article. We rely on the RavenPack ranking of the influence and trustworthiness of various publications. This ranking ranges from 1 (for the most trusted, reputable, and impartial sources) to 3 (for the least reliable sources). We add another rank category equal to 4 for sources not recorded and classified by RavenPack. We label this categorical variable as *OUTLET_RANK*. The mean of *OUTLET_RANK* is equal to 1.65, suggesting that many scoop articles come from highly reputable news outlets. Table 2, Panel D presents the frequency distribution of *OUTLET_RANK*.

3.2.2. Twitter activity

To empirically capture the social media activity related to a merger rumor, we extract from Crimson Hexagon the number of daily tweets involving the target firm. For the rumor window, we define a standardized abnormal Twitter activity as follows:

$$TWEET\_RUMOR_{Event} = \frac{(Tweet\ volume\_event\ window) - (Mean\ tweet\ volume\_control\ window)}{Standard\ deviation\ of\ daily\ \ tweet\ volume\ in\ the\ control\ window}$$

Where *Tweet volume_event window* is equal to the number of rumor-day tweets, and *Mean tweet volume_control window* is equal to the average number of daily control-window tweets. User tweet volume in the control window [-90, -21] captures the normal level of daily Twitter activity related to a firm.[21] We deflate the difference by the (time-series) standard deviation of the normal

[21] We assume abnormal tweet volume represents merger-rumor related tweets. This measure has pros and cons compared with a more precise keyword search based procedure, especially on Twitter, where users often use non-conventional wording. While it no doubt contains measurement error, we do not foresee any bias. To validate our assumption, we extract the number of tweets that contain the following keywords on the rumor date: "acquir*" or "acquisition*" or "merger*" or "deal*" or "takeover*" or "buyout*" or "bid*" or "buy*" or "rumor*" or "rumour*" or "speculat*" or "said to be" or "talk*" or "target*." We then conduct two validation tests. First, we correlate the

tweet volume. As can be seen in Panel C of Table 2, the mean of $TWEET\_RUMOR_{[0]}$ is 5.99, suggesting a significant spike in Twitter activity during the event window.

3.2.3. Other variables

To examine the market reaction to merger rumors, we focus on $CAR\_RUMOR_{[0, +1]}$, the event window market-adjusted abnormal return around the rumor publication date. We measure abnormal returns over two days [0, +1] because the abnormal tweets measured on Day 0 include those posted in the after-market hours. The two-day return window ensures that the abnormal returns fully impound the rumor-related tweets. As can be seen in Panel C of Table 2, the average market reaction to a rumored target is 5.22% for the event window, which suggests that merger rumors elicit significant and positive market reactions for rumored targets.

After the original scoop article comes out, other traditional media outlets continue to provide details and updates. These articles can potentially influence the market reaction to the rumors. We use Factiva to collect the number of traditional media articles about a merger rumor, $\#MEDIA\_ARTICLES_{[0, +1]}$.[22] We measure traditional media articles over [0, +1] (instead of just

---

two measures of rumor-related tweets, one based on our keyword search and one based on the difference between event day tweets minus control window tweets. The correlation is 36% and highly significant. Next, we regress the unscaled abnormal tweet volume on the keyword-based measure, with size and industry and year fixed effects. The coefficient is positive and highly significant. These two validation tests provide some confidence for our measure of rumor-related tweets.

[22] We search the following string in the Factiva "free text form": Company name and (acquire or acquisition or merger* or deal or takeover or buyout or bid or buy) and (rumor or rumor or speculat* or "said to be" or talk*). We use the company name (instead of the Factiva identifier) as our first search term because it produces a higher number of merger-rumor related articles. This discrepancy is due to Factiva's indexing algorithms, which require a company to be discussed in sufficient detail for the Factiva identifier to be assigned. For publication sources, we select "Major News and Business Sources: U.S." to reduce potential measurement errors. Since our keyword combinations are designed to be sufficiently broad to maximize relevant output, the searches will invariably capture unrelated articles. Ahern and Sosyura (2015) give an example (p. 2055). This type of measurement error is exacerbated if we search using "all sources," when a large number of publications are non-business related (such as *The Hollywood Reporter*, *Advertising Age, etc.)*. Still, when we use traditional media coverage from "all sources," all the inferences remain unchanged.

Day 0) to allow for a potential publication lag in the traditional press. Because we control for the role of the traditional media, the effects of Twitter activity can then be interpreted as incremental to the effects of traditional media.[23]

Firms with widely recognizable brands are more likely to be rumor targets because they attract widespread interest. We follow Ahern and Sosyura (2015) and use the ranking data from *Interbrand* and *BrandZ*, two consulting firms that publish an annual list of 100 most valuable global brands.[24] Because of the selective nature of these two lists, we define a binary variable, *VALUABLE_BRAND*, for any target firm that appeared on either list for any of the years from 2006 to 2014. We expect rumors involving household names to elicit greater reactions from both social media users and investors. As summarized by *VALUABLE_BRAND*, roughly 10% of our sample target firms have high brand recognition, which is a fairly high percentage considering that the brand measure is constructed from a very exclusive list of global brands.

Finally, several standard firm characteristics are included in the main tests, such as firm size (*LN_ASSET*) of a rumored target. To control for the growth prospect and "glamour" factor of the target firm, we include Tobin's Q (*TOBINQ)* and the ratio of research and development expenditure to total assets (*R&D*).[25] We also include the advertising expenditures to total assets ratio (*ADV*) as an additional measure of the breadth and prominence of the target firm. Because the ex-ante probability of being a takeover target varies across firms, we also calculate the

---

[23] We also collect news articles published by traditional media during and after the rumor announcement for the period up to Day +10. We use this variable as a measure of ongoing media attention. For the sake of robustness, we rerun all our specifications, including this variable as a control. Controlling for ongoing media coverage after the rumor announcement leaves all our results unchanged.

[24] http://interbrand.com/best-brands/best-global-brands/ and http://www.millwardbrown.com/brandz/top-global-brands/

[25] Tobin's Q is the market value of equity, plus the difference between assets and common equity, scaled by assets.

likelihood of being acquired based on Cremers et al. (2009).[26] Detailed descriptions of the variables are provided in Appendix A.

Table 3 presents the Pearson and Spearman correlation coefficients for the main variables. The correlation between *TWEET_RUMOR* $_{[0]}$ and *CAR_RUMOR* $_{[0, +1]}$ is positive and significant, suggesting a positive association between abnormal tweet volume and a target firm's abnormal return on the rumor day.

4. Does user tweet volume predict rumor accuracy?

In this section, we examine Hypothesis 1, that is, whether rumor-day tweet volume helps *predict* rumor accuracy (i.e., whether a merger rumor will be realized within a year). As discussed earlier, the relationship between abnormal tweet volume and rumor accuracy is not clear a priori. Rumor-day tweet volume can be a significant predictor of rumor accuracy if more "valid" rumors attract greater Twitter activity. If Twitter posts are not information-driven but motivated by users' desire to share attention-grabbing stories, then tweet volume should be unrelated to rumor accuracy.

We run a logistic regression to examine whether rumor-day abnormal tweet volume can predict rumor accuracy. Our logistic regression model is specified as follows:

$$REALIZED = \beta_0 + \beta_1\, TWEET\_RUMOR_{[0]} + Controls + \varepsilon. \qquad (1)$$

The dependent variable is *REALIZED*, which is equal to 1 if the rumor materializes within one year and 0 otherwise. The variable of interest is *TWEET_RUMOR* $_{[0]}$. The coefficient

[26] The predictive variables include Tobin's Q, PPE (property, plant, and equipment scaled by assets), cash, the presence of block holders (13F), firm size, leverage, return on assets, and whether there were merger events in the industry in the prior year.

β_1 captures whether the rumor-day Twitter activity predicts rumor realization. We present the results in Table 4. In Model 1, we include the variable of interest as well as target firm control variables. In Models 2 and 3, we include additional control variables for article characteristics, such as whether a bidder or the price is mentioned in the rumor article (*BIDDER_MENTIONED* and *PRICE_MENTIONED*), and the specificity of the rumor article (*SPECIFIC*).

In all specifications, the coefficient on $TWEET\_RUMOR_{[0]}$ is insignificant, suggesting that user tweet volume does not predict rumor realization. In other words, merger rumors accompanied by greater user tweet volume are no more likely to be accurate than those with lower tweet volume.

5. Twitter activity and market reaction to merger rumors

5.1. Market reaction to merger rumors and Twitter activity

In this section, we examine Hypothesis 2, that is, whether investors' reaction to merger rumors varies with user Twitter activity. We first sort the rumored targets into two portfolios based on the abnormal user tweet volume on the rumor day. The high portfolio consists of target firms who experience above-median abnormal rumor-day user tweet volume, while the low portfolio consists of target firms with below-median Twitter attention. We conduct standard event studies for various event windows for the two portfolios. Abnormal returns are defined as raw returns in excess of CRSP value-weighted market returns. The results for the cumulative abnormal returns (*CARs*) for various windows are presented in Table 5.

Panel A presents the *CARs* for the pre-event window [-20, -1]. In the run-up period, there is some apparent leakage for the realized rumors, but not for the unrealized rumors before the rumor announcements.

In Panel B, we examine the stock market reaction in the event windows and find a positive reaction to merger rumors. For the overall sample, the abnormal return in the [0, +1] window is 5.22%. For firms in the high (low) portfolio split based on rumor-window Twitter activity, the *CAR* is 7.57% (2.87%). This finding suggests that rumored targets with higher tweet volume elicit greater investor reaction when the rumors surface. In addition, realized rumors experience a greater market reaction, as the *CAR* in the [0, +1] window for the realized rumor sample is 9.12%, compared with 4.16% for the unrealized rumors. Importantly, even for the unrealized rumors, user Twitter activity seems to affect abnormal returns. Specifically, the rumored targets with high tweet volume experience an average *CAR* of 5.74%, compared with 2.59% for those with low tweet volume.

In Panel C, we examine how investors react to the merger rumors after the rumor dates. We examine returns in windows Day +2 to Day +20, Day +21 to Day +40, and Day +41 to Day +60. Overall, we find that for unrealized rumors, significant return reversals occur after the initial positive response to merger rumors. As late as eight weeks after the rumor date, the high portfolio still experiences a reversal for the unrealized sample. As expected, no reversal occurs for realized rumors.

In Table 6, we examine the link between the market reaction to rumors during the event window [0, +1] and the abnormal user tweet volume using a regression framework. We run OLS regressions using the following specification:

$$CAR\_RUMOR_{[0,+1]} = \beta_0 + \beta_1 TWEET\_RUMOR_{[0]} + Controls + \varepsilon \qquad (2)$$

The dependent variable is the Day [0, +1] *CAR* for the rumored targets. The main variable of interest, $TWEET\_RUMOR_{[0]}$, is the standardized abnormal user tweet volume. The coefficient β_1 captures how investors react to abnormal tweet volume. We also include the *CAR* in the pre-rumor period, $CAR\_PRE_{[-5,-1]}$, to control for prior leakage. Industry, year, and outlet rank fixed effects are included in all specifications.

Consistent with earlier results, we find that the market reaction to merger rumors is positively related to abnormal user Twitter activity. The coefficients on $TWEET\_RUMOR_{[0]}$ are positive and significant across all specifications, including both realized and unrealized rumors. In terms of economic significance, an increase in the number of tweets by one standard deviation translates to an abnormal return of 0.29%. When interpreted together with the evidence that tweet volume is unrelated to rumor realization, the heightened market reaction to rumors accompanied by greater user tweeting activity is consistent with the rumor mill effect.

5.2. Twitter activity and price discovery for unrealized merger rumors

Our baseline results suggest that social media activity distorts market reaction to merger rumors. To solidify our inference of the rumor mill effect of social media, we turn to the sample of rumors that are likely inaccurate, i.e., those rumors that fail to materialize within a year. The price evolution around these rumors, including immediate reaction and post-rumor reversals, provides a powerful setting to study the role of social media in the price discovery of (ex-post) false rumors. Since the eventual merger outcome is unknown at the time of a rumor, the market reaction to merger rumors reflects investors' collective assessment of the validity of the rumor. If investors overestimate the validity of highly tweeted rumors, the immediate market response should be positively related to tweet volume, but the price distortion should reverse over time as investors recognize the low deal likelihood.

Results in Table 6 show that investor reaction to unrealized merger rumors is positively related to user tweet volume. The coefficient is positive and highly significant across all specifications, even after we control for firm and rumor article characteristics. This result suggests that when merger rumors surface, investors behave as if tweet volume conveys higher deal likelihood, even among the unrealized rumors.

If the immediate market reaction to highly tweeted rumors documented above represents temporary price distortion, the positive coefficient on abnormal tweet volume should reverse over time. Table 7 examines the post-rumor price reversals for the unrealized rumors. Because the exact timing of the reversal is ex-ante unclear, we follow an approach used in prior literature (e.g., Hartzmark and Shue, 2018) and examine the price reversals over different intervals. Specifically, we study the price movement over roughly 60 trading days after a merger rumor, divided into three 20-day intervals. In the first two post-rumor reversal windows, [+2, +20] and [+21, +40], the coefficients on abnormal tweet volume are negative but insignificant, suggesting limited price reversal in the first eight weeks. However, the coefficient on abnormal tweet volume in the [+41, +60] window is significantly negative, consistent with price reversal. The magnitude of the coefficient is substantially larger than that in the previous two windows.[27] Together, the evidence suggests that the price distortion associated with tweet volume persists

[27] In un-tabulated tests, we explore the effects of later-period media coverage and Twitter activity over three different windows (matching the return windows). Post-rumor CARs are unrelated to either later-period media coverage or Twitter activity (possibly due to measurement error). Our main inferences are largely unaltered.

weeks after the rumors, and meaningful price correction takes place more than eight weeks post rumor announcement.[28] The post-rumor reversal results are consistent with our main inference; that is, user tweets distort investors' immediate response and lead to protracted price discovery for the highly tweeted rumors.

5.3. Cross-sectional analysis of price distortions

We explore three sources of potential cross-sectional differences of the rumor mill effect: author influence, investor base, and rumor specifics. We present the results in Table 8.

First, we investigate whether tweets posted by users with higher social media influence generate greater stock market impact. We collect the influence score from Crimson Hexagon for each Twitter user that tweeted in the rumor period. Klout score is a number between 1 and 100, with a higher score indicating a higher ranking of the breadth and strength of one's online social influence.[29] We define a binary variable, *HIGH_INFLUENCE*, which is equal to 1 if the daily mean Klout score is above the sample median, 0 otherwise, and use this variable as a cross-sectional cut. For those rumors tweeted by Twitter users with greater influence, the price distortion of abnormal tweets should be more pronounced, as social influence likely exacerbates both distortive effects of social media (persuasion bias and heightened awareness). Besides, information tweeted by influential social media users can also be viewed as more credible by capacity-constrained investors. When we examine the moderating effect of high influence, the

[28] When we modify the dependent variables to reflect the overall return windows (i.e., [0, +40] and [0, +60]), we find the coefficient is positive and significant in the former but insignificant in the latter, which confirms that the price distortion persists up till Day +40 but dissipates by Day +60.
[29] Klout scores are computed by a company named Klout. Klout assigns a score to all users based on their influence on nine major social media networks, including Twitter and Facebook. Klout scores consider both long-lasting features (e.g., followers) and dynamic features (e.g., retweets) (Rao et al., 2015).

coefficient on *TWEET_RUMOR*HIGH_INFLUENCE* is positive and highly significant, suggesting that the price distortion associated with the rumor mill effect is more pronounced for rumors accompanied by more influential tweets.[30]

Second, we study whether the investor base of a rumored target plays a role in the market reaction. We define a binary variable, *LOW_INST*, which is equal to one if the percentage of total shares owned by institutional investors (based on Thomson Reuters 13(f) filings) is in the lowest decile, zero otherwise. Because retail investors are less sophisticated and face greater information-processing constraints, the rumor mill effect is likely more pronounced among firms with a larger base of retail investors (i.e., lower institutional ownership).[31] In addition, prior literature documents that firms with lower institutional ownership are more likely to have binding short-sale constraints, which limit effective arbitraging by sophisticated investors (e.g., Nagel, 2005). Indeed, the coefficient on *TWEET_RUMOR*LOW_INST* is positive and significant, suggesting that the price distortion associated with the rumor mill effect is exacerbated for targets with lower institutional ownership.

Lastly, we examine whether the effect of tweet volume on abnormal returns is more pronounced among rumors that appear more specific based on deal information from the scoop articles. The coefficient on *TWEET_RUMOR*SPECIFIC* is positive and significant. This outcome suggests that investors are more influenced by social media activity if the rumors appear more credible at the time, even for rumors that eventually fail to materialize. Because

[30] In un-tabulated results, we find that social media influence, alone or interacted with tweet volume, is unrelated to either the likelihood of merger realization or perceived merger gains. These findings confirm our interpretation that highly influential Twitter users exacerbate price distortions driven by tweet volume.

[31] A crucial assumption underlying this cross-sectional analysis is that institutional ownership captures the proportion of sophisticated trades occurring, which may not always be true.

more specific merger rumors are more likely to materialize (Table 4), capacity-constrained investors can rationally look to rumor specificity as an indication of rumor accuracy when they respond to highly tweeted rumors at the time of the merger rumors.

6. Endogeneity and the Causal Link between Tweets and Abnormal Returns

While we aim to establish a causal link between Twitter activity and abnormal market reaction, the inter-relationship between the two is complex, as depicted in Figure 1. Our inference that Twitter activity affects market reaction (causal link 1) is confounded by two important endogeneity concerns: reverse causality (i.e., abnormal returns causing abnormal tweets, shown as causal link 2) and unobservable omitted correlated variables (e.g., perceived rumor validity or anticipated merger gains driving both investor reaction and Twitter activity). Below, we explore a few settings to provide evidence that strengthens our interpretation of the causal link between tweets and abnormal market reactions.

6.1. Rumors in non-trading windows

Because abnormal returns and tweets are measured contemporaneously, it is challenging to isolate the effect of tweets on returns. To mitigate concerns of reverse causality, we analyze the market reaction to rumors that surface during non-trading windows. Even though the stock market is closed, Twitter users can still tweet about merger rumors that surface during these windows. By construction, these tweets are unlikely driven by the market reaction, which helps to mitigate concerns of reverse causality. We exploit two different non-trading settings.

First, we identify 27 rumors that surfaced on weekends or holidays. We regress the abnormal return of the *first available* trading day on the rumor-day tweet volume.[32] The results are presented in Table 9, Panel A. Due to the small sample size, we include limited control variables in addition to industry and year fixed effects. Interestingly, despite the small sample size, we find that the subsequent available stock market reaction to a merger rumor is positively and significantly related to the rumor-related tweet volume in the non-trading window.

Next, we identify 36 rumors that surfaced during non-trading hours by using article time stamps (whenever available). If we measure tweets in the non-trading hours, either before the market opens (from 0:00 to 9:29 am) or after the market closes (from 4:01 to 11:59 pm), the tweets should not be affected by the returns.[33] Any association between the tweet volume measured in the non-trading hours and the subsequent return most likely reflects the effect of tweets on returns, not vice versa. The results, presented in Panel B, show that when we regress the abnormal return of the subsequent trading day on abnormal tweet volume measured over the prior non-trading window, the coefficients are positive and highly significant.[34] Both the weekend/holiday test and the non-trading-hour test offer similar insight; that is, tweet volume likely has an effect on returns, and our main results are unlikely driven by reverse causality.

---

[32] If a rumor surfaced on Friday after the close of the market, $TWEET\_RUMOR_{[0]}$ is measured for Friday aftermarket, Saturday, and Sunday. If a rumor surfaced on Saturday, $TWEET\_RUMOR_{[0]}$ is measured for Saturday and Sunday. If a rumor surfaced on Sunday, $TWEET\_RUMOR_{[0]}$ is measured for Sunday.

[33] Although trades can happen during extended market hours through ECNs (electronic communications networks), most investors use limit orders, and the market is notoriously thin, resulting in wide bid-ask spreads and low likelihood of order execution. Given the sparse trading during these hours, Twitter activity is unlikely to be driven by the sporadic return observations, compared with regular trading hours.

[34] If a rumor surfaced on Day t before the market opens (0:00–9:29 am), $TWEET\_RUMOR_{[0]}$ is measured for the pre-market window. If a rumor surfaced on Day t after the market closes (4:01–11:59 pm), abnormal Twitter activity is measured for the post-market window on Day t and for the pre-market window on Day $t + 1$ separately. $TWEET\_RUMOR_{[0]}$ is then obtained by summing up the two.

6.2. Instrumental variable approach

6.2.1. Twitter outages

To address the two endogeneity concerns discussed above, reverse causality and unobservable correlated omitted variables, we explore possible exogenous shocks to Twitter activity caused by systematic and prolonged Twitter outages using an instrumental variable approach. We collect data on Twitter interruptions during our sample period from Twitter's official support account (@TwitterSupport), which regularly posts tweets informing users of issues on Twitter. Using Twitter API, we retrieve the time stamp and content of all tweets posted by @TwitterSupport from 2009 to 2014. We then manually go through tweets that are posted around our rumor days to identify potentially relevant tweets describing Twitter outages.[35]

For each outage-related tweet, we record the nature of the outage as well as the duration. We estimate outage duration by examining the timestamps of the initial tweet announcing an issue and a subsequent tweet announcing the resolution.[36] After deleting insignificant interruptions (i.e., those lasting less than 45 minutes), a total of 42 episodes of Twitter interruption overlap with our sample rumor dates. We create a binary variable *OUTAGE* to denote rumors that coincide with an outage in the sample.

In the first two columns of Table 10, Panel A, we report the estimation results of the 2SLS using the instrument *OUTAGE*. In the first-stage regression, the coefficient on *OUTAGE* is

[35] Besides tweets on Twitter malfunctions, @TwitterSupport also frequently posts tweets to promote new functions of Twitter and provide answers to frequently asked questions about Twitter.
[36] For example, at 8:34 pm on June 3, 2013, @TwitterSupport posted, "Some users may be experiencing issues accessing Twitter. Our engineers are working to resolve it. Find updates at status.twitter.com." At 9:23 pm on the same day, it subsequently posted, "The site access issues have now been resolved. Thanks for your patience, everyone!" We calculate the duration of this outage episode as 49 minutes, the time difference between the initial tweet at 8:34 pm and the subsequent tweet at 9:23 pm.

negative, as expected, and significant at the 10% level, indicating fewer rumor-related tweets when Twitter is experiencing outages. However, the partial F statistic on the excluded instrument (3.28) indicates a weak instrument (Stock et al., 2002). In the second stage, the coefficient on the *instrumented* tweet volume is positive, as predicted by the rumor mill effect, but insignificant.

We conjecture that the weak significance in the initial outage analysis is attributable to a lack of power because many of these interruptions involve minor or local issues that only affect a small subset of users or functions. To focus on more severe and widespread outages, we classify the 42 Twitter outages into the following four categories: (1) site crashes, or high/elevated error rates on Twitter, (2) sign-in/access/connectivity/stability issues, (3) isolated local issues for specific groups (e.g., AT&T users) or related to specific functions (e.g., profile picture change), and (4) other miscellaneous issues. We conjecture that outages pertaining to general accessibility issues (i.e., the first two categories) are more "severe" in nature because they lead to system-wide interruptions that negatively affect users' ability to view and post tweets. As expected, severe outages are relatively rare, with only ten such episodes in our rumor sample. However, we expect the impact of these outages to be greater than that of minor outages, and as a result, they are likely to generate meaningful variations in expected tweet volume.

In Columns 3 and 4, we repeat the 2SLS analysis but focus on more severe outages. As can be seen in Column 3, the coefficient on the instrument, *OUTAGE_SEVERE* is negative and highly significant, indicating lower rumor-related tweet volume when the rumor coincides with a severe Twitter outage. The partial F statistic on the excluded instrument is 14.28, which easily satisfies the conventional criteria for instrument relevance. Column 4 presents the second-stage estimation results, where we regress the rumor-window abnormal return on the instrumented tweet volume from the first stage. The coefficient on the instrumented tweet volume is positive and significant. Overall, the 2SLS results suggest a causal link between rumor-related tweets and the market reaction to rumors.

6.2.2 Viral tweets

In this section, we use the presence of a viral tweet (unrelated to merger rumors) on a rumor date to instrument for the volume of rumor-related tweets. The presence of a viral tweet is a plausible instrument because it potentially diverts the attention of Twitter users from merger rumors and reduces the circulation of these rumors on Twitter, and because these viral tweets are otherwise unlikely to be correlated with the market reaction to merger rumors.[37]

To identify dates with viral tweets, we compile a list of viral tweets from 2009 to 2014 from a variety of sources. *Time* magazine publishes an annual article series called *Top 10 of Everything*, including a list of Top 10 Tweets for 2009 and 2010. From 2011 to 2013, the series expanded to include lists for Top 10 Best Tweets and Top 10 Worst Tweets. In 2014, the magazine stopped publishing the list of top tweets. As a result, we use a set of noteworthy tweets in 2014, summarized by Twitter.[38] Several lists of popular tweets from other media sources are also collected and used to supplement the lists from *Time*.[39] Our final viral tweet sample consists

[37] Since exclusion restrictions cannot be tested directly, we take several steps to ensure that the criteria are reasonably met. First, we provide the list of viral tweets that overlap with the event window of a sample rumor (Appendix D). As can be seen, these tweets were mostly posted by celebrities (e.g., Kanye West and Justin Bieber), and based on the content, most of the viral tweets are unlikely to affect the stock market directly. Some of the tweets, however, are more ambiguous, such as one related to Twitter IPO and six others posted by prominent businesspeople or politicians. When we delete these seven viral tweets, our 2SLS results remain unchanged.

[38] Available at https://twitter.com/twitter/timelines/540615025182773248?lang=en

[39] These supplemental lists include: (1) List of Most Retweeted Tweets, https://en.wikipedia.org/wiki/List_of_most-retweeted_tweets; (2) Golden Tweets of 2012, https://2012.twitter.com/en/golden-tweets.html; and (3) Most Retweeted Tweets of 2013, http://time.com/11777/this-is-the-most-retweeted-tweet-of-2013/

of more than 120 tweets that occurred from 2009 to 2014. For each viral tweet, we collect the date the tweet was published, the content, and the author of the tweet.[40] We then identify merger rumors that occur within the event window of a viral tweet and create a binary variable *VIRAL_TWEET* to denote these rumors. A total of 35 viral tweets coincide with a rumor.

Panel B of Table 10 reports the 2SLS results using only *VIRAL_TWEET* as an instrument (Columns 1 and 2) or using both *VIRAL_TWEET* and *OUTAGE_SEVERE* as instruments (Columns 3 and 4). As expected, the coefficient on *VIRAL_TWEET* is negative and highly significant in both specifications, suggesting that when Twitter is abuzz with other viral tweets, fewer Twitter users tweet about a rumored takeover target. The partial F statistic on the excluded instrument is 9.56 when only one instrument is used and 10.96 when two instruments are used, both satisfying the conventional criteria for instrument relevance (Stock et al., 2002). In the second stage, the coefficient on the *predicted* rumor-related tweet volume is positive and significant in both specifications (Columns 2 and 4). Taken together, the 2SLS evidence using Twitter outages and/or viral tweets provides consistent support for the causal inference of the effects of abnormal tweets on abnormal returns.

7. Additional analyses

7.1. Analysis of rumors that were eventually realized

So far, we provide evidence of the rumor mill role of social media in a setting where the underlying news is likely false. Our results do not speak to the crowd wisdom role, especially

[40] In some cases, it is difficult to identify the dates of these tweets precisely. The lists compiled by *Time* do not always provide the specific date of the viral tweet, so in these cases, we have to use Google to identify the date. Moreover, some original tweets, especially tweets on the Top 10 Worst Tweets list, were deleted after the initial posting by the authors. In these cases, we rely on related media reports to approximate the tweet date.

when the underlying news is likely true. Nevertheless, it is interesting to examine whether and how rumor-day tweet volume affects the price discovery of those rumors that eventually materialized. The results are presented in Table 11. A few interesting findings emerge. First, the immediate market reaction to the 63 "realized" rumors is more pronounced for rumor targets accompanied by more tweets. Second, the tweet-volume-driven market reaction does not reverse. Third, the takeover announcement is negatively associated with rumor-day tweets. Viewed collectively, the evidence alludes to the crowd wisdom effect for realized rumors, as Twitter posts appear to accelerate price discovery. However, we hesitate to draw strong inferences because only a small minority of merger rumors (21%) eventually materialize.

7.2. Detailed analysis of journalist and article characteristics

Data limitations prevent us from including detailed article characteristics and journalist characteristics in our primary analyses, even though they could be correlated omitted variables (Ahern and Sosyura, 2015). To address this concern, we hand collect additional journalist and article characteristics from various public sources. To gather journalist characteristics, we extract the journalist's name for each rumor article and locate the journalist's biographical webpage or LinkedIn profile. We record the following information: (1) education history, (2) age, (3) gender, (4) location, (5) awards, (6) industry expertise, (7) graduate degree, and (8) columnist. We define these variables in Appendix A. We are able to obtain detailed journalist characteristics for only 118 rumors. We rerun our main analyses with these additional journalist variables. In Table 12, Panel A, we present the logit model predicting merger realizations (Columns 1 and 2) and the OLS results of CARs on abnormal tweet volume (Columns 3–6). Adding journalist characteristics does not alter our main inferences. As shown in Columns 3–6, the coefficient on our main variable of interest, abnormal tweet volume, continues to be positive and significant in all four specifications, despite the substantially reduced sample size.

Additionally, we collect and control for detailed article-level characteristics, including (1) weak modal words, (2) anonymous source, (3) target response, (4) rumor in headline, and (5) # of numerical values in the article. As in the above analysis, we rerun our main analyses with these additional article-level variables, and the results are presented in Panel B. The first two columns present the logit regression predicting merger realizations. Abnormal tweet volume continues to be insignificant, consistent with our main results. In Columns 3–6, we present the OLS regression results of rumor-window CARs on abnormal tweet volume. Our main variable of interest, abnormal tweet volume, continues to be positive and significant after we control for detailed article-level characteristics.[41]

7.3. Target response on social media

In the main analysis, we have ignored possible responses by rumored targets on social media. If rumored targets use Twitter to provide clarifications and updates on rumored mergers, these firm-initiated tweets can change the information dynamics and thus the price evolution. To examine responses by rumored targets or bidders (if mentioned in the scoop article) via Twitter, we start by randomly selecting 120 rumors, corresponding to slightly more than one-third of our final sample. We then manually collect the official Twitter handle of each rumored target and the

[41]CONFIRMED RUMOR and DENIED RUMOR cases are rare, with only four such cases in our sample, two realized and two unrealized. In our logistic regressions (columns 1 and 2), the two confirmed rumors predict success (deal realization) perfectly, so the coefficient on CONFIRMED RUMOR could not be estimated. Similarly, DENIED RUMOR predicts failure (unrealized deals) perfectly. In columns 5 and 6, the coefficients on CONFIRMED RUMOR cannot be estimated because only *unrealized* rumors are used in these regressions.

corresponding bidder (if available) from Google and Twitter. Next, we search for any tweet posted by a rumored target or bidder during a five-day window starting from the rumor date. Out of the 120 rumors, only two firms provided a response on Twitter, one by a target firm and the other by a bidding firm. In both cases, the merger rumor was denied. Deleting these two observations does not change any of our inferences. It appears that a firm's response on Twitter has no material impact on our results.

8. Conclusion

We provide the first systematic evidence on the rumor mill effect of social media, using a sample of merger rumors in which the reported news is highly speculative, and the vast majority of the rumors do not materialize. While social media has the potential to facilitate the aggregation of information and help stock price formation, it also provides fertile ground for disinformation to spread and hinder price discovery. We find evidence of the rumor mill effect in the presence of highly speculative rumors. Specifically, user Twitter activity around merger rumors distorts, rather than facilitates, price discovery.

We first document that rumor-window user tweet volume is insignificantly related to rumor realization (i.e., whether a merger rumor materializes within a year). Despite the insignificant relationship between tweet volume and rumor accuracy, merger rumors with greater Twitter activity experience *higher* rumor-window abnormal returns. Investors behave as if tweet volume conveys valid information about deal likelihood. Together, the findings are suggestive of the rumor mill role of social media.

When we specifically focus on the price evolution around unrealized merger rumors, we find similar price distortions driven by tweet volume, as evidenced by immediate overreaction and prolonged price discovery. The price distortion associated with rumor-day tweet volume persists even 40 trading days after a rumor surfaces. Taken together, the evidence highlights a potential downside of social media; that is, it can facilitate the spread of potentially false information and distort the price discovery process in the capital market. Additional tests suggest that our main inferences are robust and are not likely affected by reverse causality or correlated omitted variables such as the perceived validity of these merger rumors or anticipated merger gains (upon completion). Although some of these causality tests are based on relatively small sample sizes, we are reassured by the consistent inferences across multiple approaches to address endogeneity.

In addition, we find that author influence, investor base, and rumor specificity affects the relationship between social media and price discovery. For the unrealized rumor targets, the documented relation between high social media activity and distorted price discovery is more pronounced for rumors followed by more influential Twitter users, for target firms with low institutional holdings, and for rumors that come across as more authentic.

Although we find evidence of the rumor mill effect in the presence of highly speculative rumors, we are not dismissing the crowd wisdom role of social media in other settings. Our research design, while seeking to maximize the power of the tests related to the rumor mill effects, is not necessarily suited to testing the crowd wisdom effect.

Our evidence that social media is a double-edged sword is of interest to academics, practitioners, and regulators. In an age when fake news can take on a life of its own and people are increasingly turning to social media as an information channel, our paper also contributes to the broader debate regarding the roles and responsibilities of media and social media outlets.

It is important to note that we study how social media affects the spreading of financial rumors while being silent on how social media could have changed the origination of rumors. It would be fruitful to study whether the proliferation of social media use has affected the origination of financial rumors. The evolution of social media has most likely changed how and when the news is reported by traditional media (Alejandro, 2010). The interaction between social media and traditional press presents an important and interesting area for future research.

References

Ahern, K.R., Sosyura, D., 2015. Rumor has it: Sensationalism in financial media. Rev. Financial Stud. 28(7), 2050–2093.

Alejandro, J., 2010. Journalism in the age of social media. Reuters Institute Fellowship Paper, University of Oxford, 2009–2010.

Bartov, E., Faurel, L., Mohanram, P.S., 2017. Can Twitter help predict firm-level earnings and stock returns? Account. Rev. 93(3), 25–57.

Blankespoor, E., deHaan, E., Marinovic, I. 2020. Disclosure processing costs, investors' information choice, and equity market outcomes: a review. J. Account. Econ. Forthcoming.

Blankespoor, E., Miller, G.S., White, H.D., 2014. The role of dissemination in market liquidity: Evidence from firms' use of Twitter. Account. Rev. 89(1), 79–112.

Chen, H., De, P., Hu, Y.J., Hwang, B.H., 2014. Wisdom of crowds: The value of stock opinions transmitted through social media. Rev. Financial Stud. 27(5), 1367–1403.

Core, J. E., Guay, W., Larcker, D. F., 2008. The power of the pen and executive compensation. J. Financial Econ. 88(1), 1-25.

Cremers, K.M., Nair, V.B., John, K., 2009. Takeovers and the cross-section of returns. Rev. Financial Stud. 22(4), 1409–1445.

Dai, L., Parwada, J.T., Zhang, B., 2015. The governance effect of the media's news dissemination role: Evidence from insider trading. J. Account. Res. 53(2), 331–366.

Dang, T.L., Moshirian, F., Zhang, B., 2015. Commonality in news around the world. J. Financial Econ. 116(1), 82–110.

DeMarzo, P., Vayanos, D., Zwiebel, J., 2003. Persuasion bias, social influence, and unidimensional opinions. Q. J. Econ. 118(3), 909–968.

Gentzkow, M., Shapiro, J. M., 2010. What drives media slant? Evidence from U.S. daily newspapers. Econometrica, 78(1), 35-71.

Gilbert, D.T., 1991. How mental systems believe. Am. Psychologist 46(2), 107–119.

Guess, A., Nyhan, B., Reifler, J., 2018. Selective exposure to misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign. European Research Council.

Gurun, U.G., Butler, A.W., 2012. Don't believe the hype: Local media slant, local advertising and firm value. J. Finance 67(2), 561–598.

Hartzmark, S.M., Shue, K., 2018. A tough act to follow: Contrast effects in financial markets. J. Finance 73(4), 1567–1613.

Hawkins, S., Hoch, S., (1992). Low-involvement learning: Memory without evaluation. J. Consumer Res. 19(2), 212–225.

Jung, M.J., Naughton, J.P., Tahoun, A., Wang, C., 2018. Do firms strategically disseminate? Evidence from corporate use of social media. Account. Rev. 93(4), 225–252.

Kolasinski, A.C., Reed, A.V., Ringgenberg, M.C., 2013. A multiple lender approach to understanding supply and search in the equity lending market. J. Finance 68(2), 559–595.

Lee, L.F., Hutton, A.P., Shu, S., 2015. The role of social media in the capital market: Evidence from consumer product recalls. J. Account. Res. 53(2), 367–404.

Loughran, T. and McDonald, B., 2011. When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks. J. Finance, 66, 35-65.

Ma, M., Zhang, F., 2016. Investor reaction to merger and acquisition rumors. SSRN Working Paper. http://dx.doi.org/10.2139/ssrn.2813401.

Massa, M., Qian, W., Xu, W., Zhang, H., 2015. Competition of the informed: Does the presence of short sellers affect insider selling? J. Financial Econ. 118(2), 268–288.

Miller, G.S., Skinner, D.J., 2015. The evolving disclosure landscape: How changes in technology, the media, and capital markets are affecting disclosure. J. Account. Res. 53(2), 221–239.

Nagel, S., 2005. Short sales, institutional investors and the cross-section of stock returns. J. Financial Econ. 78(2), 277–309.

Nyhan, B., 2014. Why rumors outrace the truth online. New York Times Upshot Blog, September 29, 2014. https://www.nytimes.com/2014/09/30/upshot/its-so-much-more-fun-to-spread-rumors-than-the-truth.html (accessed 20 March 2020).

Rao, A., Spasojevic, N., Li, Z., Dsouza, T., 2015. Klout score: Measuring influence across multiple social networks, in: 2015 IEEE International Conference on Big Data (Big Data). IEEE, pp. 2282–2289.

Stock, J., Wright, J. Yogo, M., 2002. A survey of weak instruments and weak identification in generalized method of moments. J. Bus. Econ. Statistics 20(4), 518–529.

Tang, V.W., 2018. Wisdom of crowds: Cross-sectional variation in the informativeness of third-party-generated product information on Twitter. J. Account. Res. 56(3), 989–1034.

Toubia, O., Stephen, A.T., 2013. Intrinsic vs. image-related utility in social media: Why do people contribute content to Twitter? Marketing Sci. 32(3), 368–392.

Appendix A. Variable Definitions

Variable	**Variable Definition**	**Notes**
Tweet variables		
TWEET_RUMOR $_{[0]}$	Standardized abnormal Twitter volume for rumor day [0]: abnormal daily tweets relative to control window, scaled by standard deviation of daily tweets in the control window [-90, -21].	Source: Crimson Hexagon
HIGH_INFLUENCE	Binary variable equal to 1 if the daily mean Twitter user Klout score (between 1 and 100, with a higher score indicating greater social media influence) is above the sample median, 0 otherwise.	
VIRAL_TWEET	Binary variable equal to 1 if a merger rumor surfaces within [-1, +1] days of a viral tweet, 0 otherwise.	Manually collected
OUTAGE	Binary variable equal to 1 if a merger rumor coincides with a generic outage that lasts for longer than 45 minutes when duration is available, 0 otherwise.	
OUTAGE_SEVERE	Binary variable equal to 1 if a merger rumor coincides with an OUTAGE (described above) where the nature of the issue pertains to (1) site crashes, or high/elevated error rates on Twitter, or (2) sign-in/access/connectivity/stability issues, 0 otherwise.	
Stock return variables		
CAR_RUMOR $_{[0]}$	Cumulative abnormal returns in percentages at rumor announcement dates.	Source: CRSP
CAR_RUMOR $_{[0, +1]}$	Cumulative abnormal returns in percentages for window [0, +1] around rumor announcement dates.	
CAR_PRE $_{[-5, -1]}$	Cumulative abnormal returns in percentages for window [-5, -1] before rumor announcement dates.	
CAR $_{[+2, +20]}$	Cumulative abnormal returns in percentages for window beginning on day +2 and ending on day +20, after rumor announcement dates.	
CAR $_{[+21, +40]}$	Cumulative abnormal returns in percentages for window beginning on day +21 and ending on day +40, after rumor announcement dates.	
CAR $_{[+41, +60]}$	Cumulative abnormal returns in percentages for window beginning on day +41 and ending on day +60, after rumor announcement dates.	
CAR_TAKEOVER_ANN $_{[0, +1]}$	Cumulative abnormal returns in percentages for window [0, +1] around takeover announcement dates.	
Article characteristic variables		
REALIZED	Binary variable equal to 1 if the rumor materializes within one year after the rumor announcement date, 0 otherwise.	Source: SDC
SPECULATIVE	Categorical variable equal to 2 if merger stage is “opinion pieces”, equal to 1 if merger stage is “speculation”, 0 otherwise.	
BIDDER_MENTIONED	Binary variable equal to 1 if potential bidder/s is/are mentioned in the scoop article, 0 otherwise.	Manually

PRICE_MENTIONED	Binary variable equal to 1 if takeover price is mentioned in the scoop article, 0 otherwise.	collected
SPECIFIC	Binary variable equal to 1 if takeover price and bidder/s are mentioned in the scoop article, and merger stage is equal to "made offer", or "evaluating bid", or "in advanced talk", 0 otherwise.	
OUTLET_RANK	Categorical variable ranging from 1 to 4, where a rank of 1 is the highest value (i.e., most trusted, reputable, and impartial source), and a rank of 4 is the lowest one (i.e., unknown identity and seemingly unreliable).	Source: RavenPack
#MEDIA_ARTICLES $_{[0, +1]}$	Number of traditional press articles related to a merger rumor (based on keyword search) during the rumor period [0, +1].	Source: Factiva
WEAK MODAL WORDS (%)	Ratio of weak modal words in the text of a scoop article on total number of words. Weak modal words are defined in Loughran and McDonald (2011) and include the following: apparently, appeared, appearing, appears, conceivable, could, depend, depended, depending, depends, may, maybe, might, nearly, occasionally, perhaps, possible, possibly, seldom, seldomly, sometimes, somewhat, suggest, suggests, uncertain, and uncertainly.	Manually Collected
ANONYMOUS SOURCE	Binary variable equal to 1 if an article does not identify a specific source of the rumor, 0 otherwise.	
TARGET RESPONSE	Categorical variable that records the target firm's response to the rumor, according to the text of the newspaper article: No comment, Has conversations, Confirmed rumor, Denied rumor, Couldn't be reached, or Wasn't asked.	
MERGER STAGE	Categorical variable that records the stage of the rumored talks, according to the text of the scoop article: Preliminary talks, In talks, Made offer, Preparing a bid, For sale, Evaluating bid, In Advanced Talk, Speculation, and Opinion Piece.	
RUMOR IN HEADLINE	Binary variable equal to 1 if the scoop article refers to the rumor in the headline of the article., 0 otherwise.	
# OF NUMBERS	Number of numerical values reported in the scoop article.	

Target characteristic variables

VALUABLE_BRAND	Binary variable equal to 1 if the target firm appears on the list of the 100 most valuable global brands, 0 otherwise.	Manually collected
LN_ASSET	Natural log of total assets	Source: Compustat
TOBINQ	Market value of equity plus the difference between assets and common equity, scaled by total assets.	
R&D	R&D expenses scaled by total assets.	
ADV	Advertising expense scaled by total assets.	
TAKEOVER_PROBABILITY	Probability of being a takeover target for a firm (in percentages). It is predicted by a set of variables including Tobin's Q, Property, Plant and Equipment (PPE), cash, the presence of block holders, market capitalization, leverage, Return on Assets (ROA), and the presence of merger events in the industry in the prior year.	Calculated based on Cremers et al. (2009)

LOW_INST	Binary variable equal to 1 if the total amount of shares owned by institutional investors is in the lowest decile, 0 otherwise.	Source: Thomson Reuters 13(f) database

Journalist characteristic variables

LOG (JOURNALIST AGE)	Log of the average age of all journalists listed as authors of a scoop article.	Manually collected
UNDERGRADUATE DEGREE	Binary variable equal to 1 if an article is written by a journalist who graduated with a major in one of the following categories: Business & Economics, English, Political Science, History, Other, 0 otherwise. Degree categories are as defined in Ahern and Sosyura (2015).	
EXPERT in TARGET INDUSTRY	Binary variable equal to 1 if any journalist listed as author of a scoop article is an expert in the same industry as the primary industry of the rumor target, using Fama-French 17 industry codes, 0 otherwise.	
NY BASED	Binary variable equal to 1 if at least one of the authors of a scoop article is based in New York City, 0 otherwise.	
AWARD WINNER	Binary variable equal to 1 if at least one of the authors of a scoop article has been nominated for or received journalism related award, such as the Pulitzer Prize or the Gerald Loeb Award, 0 otherwise.	
GENDER	Binary variable equal to 1 if an article has at least one female coauthor, 0 otherwise.	
GRADUATE DEGREE	Binary variable equal to 1 if an article is written by a journalist who has a graduate degree, 0 otherwise.	
COLUMNIST	Binary variable equal to 1 if at least one of the authors of a scoop article is a columnist, 0 otherwise.	

Appendix B. Rumor article examples

Exclusive: Boeing explores purchase of Mercury Systems – sources

Reuters
Nicola Leske and Mike Stone
3 April 2014

Boeing Inc (BA.N) is considering buying Mercury Systems Inc (MRCY.O), a supplier of digital signal and image processing systems to the aerospace and defense industry, according to two people familiar with the matter.

……

A move by Boeing to buy one of its suppliers would allow the company to save costs and increase its footprint in commercial aerospace components. It would also allow the Chicago-based company to gain access to Mercury's microprocessor business, which can be used on unmanned aerial vehicles, according to one of the people.

……

A Boeing spokesman said that the company did not comment on rumors of potential mergers, acquisitions, divestitures or joint ventures. Mercury Systems was not immediately available for comment

Should VF Acquire Columbia Sportswear?

Motley Fool
10 March 2014

Certain brands always seem to hang around despite no longer facing hugely popular demand. One such brand is Columbia Sportswear (NASDAQ: COLM).

However, recent results indicate that the Columbia Sportswear brand is gaining popularity. In many ways, the outdoors brand seems like an ideal candidate for a takeover. The company would fit perfectly into the large and diverse brand portfolio of VF (NYSE: VFC).

……

Columbia Sportswear just reported its best quarter ever and the best news is, unlike most other retailers of late, management is not blaming the weather! The company's overall brand seems to be in higher demand than usual. This makes the relatively small Columbia Sportswear a prime buyout candidate for a growth-hungry and acquisition-friendly retail giant like VF.

Appendix C. Rumor-period tweet examples

Tweet	Rumored Target	Date
Halftime: Is Office Depot the Next Takeover Target? http://bit.ly/fb51XJ $ARMH $INTC $LLNW $MSFT $ODP $OWW $VIX #WirelessCommunications	Office Depot Inc.	12/22/2010
Any thoughts on rumors that $VZ is looking to acquire $NFLX? - http://t.co/5FWX9n6i	Netflix Inc.	12/12/2011
Talbots May Be Acquired http://t.co/964NloCs $TLB	Talbots Inc.	1/20/2012
$SWY Safeway eyed by Cerberus, other PE firms, Reuters says http://t.co/ym9aVg7TVq	Safeway Inc.	10/22/2013
Should Pfizer Buy Eli Lilly?.. http://t.co/cHwfHucE9p $LLY #biotech #stocks	Eli Lilly & Co.	6/13/2014
Oracle: Said Near Deal to Buy Micros Systems -Bloomberg http://t.co/RIdxgv8JSi $ORCL $MCRS	Micros Systems Inc.	6/17/2014
Will Yahoo Or Facebook Acquire Yelp? http://t.co/UJIIIBC651 $BABA $FB $YHOO $YELP	Yelp Inc.	9/22/2014
Too bad if $MCRS gets gobbled up by $ORCL. It will be one Purple Chip swallowed by another and one less top co on our list of Purple Chips.	Micros Systems Inc.	6/17/2014
Sterne Agee out right now saying they believe $EMC will need to acquire $FIO!!!	Fusion-io Inc.	8/16/2012

This appendix presents examples of tweets related to merger rumors.

Appendix D. Viral tweet list

	Date	Tweet	Author	Source
1	10/8/2009	"The wait is over! The black Svengali has arrived! I'm on the street turning good girls bad and getting them pregnant!!!"	Tracy Morgan	Top 10 Tweets of 2009 - *Time*
2	10/8/2009	"FYI Liam doesn't have a Twitter and he wants ME to delete mine with good reason."	Miley Cyrus	Top 10 Tweets of 2009 - *Time*
3	3/24/2010	"3-D is a distracting, annoying, anti-realistic, juvenile abomination to use as an excuse for higher prices."	Roger Ebert	Top 10 Tweets of 2010 - *Time*
4	6/9/2010	"@Sn00ki u r right, I would never tax your tanning bed! Pres Obama's tax/spend policy is quite The Situation. but I do rec wearing sunscreen!"	John McCain	Top 10 Tweets of 2010 - *Time*
5	7/18/2010	"Ground Zero Mosque supporters: doesn't it stab you in the heart, as it does ours throughout the heartland? Peaceful Muslims, pls refudiate"	Sarah Palin	Top 10 Tweets of 2010 - *Time*
6	6/22/2010	"I am the real Liam Payne"	Liam Payne	List of Most Retweeted Tweets - Wikipedia
7	9/4/2010	"I'm sorry Taylor."	Kanye West	Top 10 Tweets of 2010 - *Time*
8	9/6/2010	"At any moment, Justin Bieber uses 3% of our infrastructure. Racks of servers are dedicated to him. - A guy who works at Twitter"	Dustin Curtis	Top 10 Tweets of 2010 - *Time*
9	2/9/2011	"Mission accomplished. Thanks to all the brave young Egyptians."	Wael Ghonim	Top 10 Best Tweets of 2011 - *Time*
10	7/27/2011	"@wiggsd Sorry to hear that. Fiscal policy is important, but can be dry sometimes. Here's something more fun: http://tinyurl.com/y8ufsnp #WHChat"	The White House	Top 10 Best Tweets of 2011 - *Time*
11	8/1/2011	"The #Capitol looks beautiful and I am honored to be at work tonight"	Gabrielle Giffords	Top 10 Best Tweets of 2011 - *Time*
12	10/5/2011	"For those of us lucky enough to get to work with Steve, it's been an insanely great honor. I will miss Steve immensely. http://b-gat.es/qHXDsU"	Bill Gates	Top 10 Best Tweets of 2011 - *Time*
13	5/24/2012	"People rarely look the way you expect them to, even when you've seen pictures."	New York Fiction	Top 10 Best Tweets of 2012 - *Time*
14	8/6/2012	"I'm safely on the surface of Mars. GALE CRATER I AM IN YOU!!! #MSL"	Curiosity Rover	Top 10 Best Tweets of 2012 - *Time*
15	9/26/2012	"RIP Avalanna. i love you"	Justin Bieber	Top 10 Best Tweets of 2012 - *Time*
16	10/22/2012	"I highly approve of Romney's decision to be kind and gentle to the retard."	Ann Coulter	Top 10 Worst Tweets of 2012 - *Time*
17	11/7/2012	"Four more years."	Barack Obama	Top 10 Best Tweets of 2012 - *Time*
18	3/21/2013	"I want @drake to …"	Amanda Bynes	Top 10 Worst Tweets of 2013 - *Time*
19	5/2/2013	"Warren is in the house."	Warren Buffet	Top 10 Best Tweets of 2013 - *Time*

20	6/25/2013	"Something special is happening in Austin tonight: http://OFA.BO/CBZ6c7 #StandWithWendy"	Barack Obama	Top 10 Best Tweets of 2013 - *Time*
21	9/12/2013	"Yesss ! I'm 20 ! Wohooo ! No more teens!"	Niall Horan	Most Retweeted Tweets of 2013 - *Time*
22	9/12/2013	"We've confidentially submitted an S-1 to the SEC for a planned IPO. This Tweet does not constitute an offer of any securities for sale."	Twitter	Top 10 Best Tweets of 2013 - *Time*
23	10/2/2013	"Listen, we're all *possibly* Frank Sinatra's son."	Ronan Farrow	Top 10 Best Tweets of 2013 - *Time*
24	1/24/2014	"YOU ARE ALL WORTHY NO MATTER WHAT ANYONE SAYS >> BE STRONG GOD IS WITH US ALL> MY BELIEBERS CHANGED MY LIFE> I WILL FOREVER BE GRATEFUL"	Justin Bieber	List of Most Retweeted Tweets - Wikipedia
25	2/20/2014	ありがとうございました！お疲れ様！！	Akiko Suzuki	Noteworthy tweets of 2014 - Twitter
26	4/17/2014	"#PRAYFORSOUTHKOREA"	g-Dragon	Noteworthy tweets of 2014 - Twitter
27	5/12/2014	"Hi this is Thierry Henry. This is my only official Twitter account. Everything posted on here will be my views. Welcome."	Thierry Henry	Noteworthy tweets of 2014 - Twitter
28	5/15/2014	"India has won! भारत की विजय। अच्छे दिन आने वाले हैं।"	Narendra Modi	Noteworthy tweets of 2014 - Twitter
29	5/15/2014	"Yes, that last tweet from 'Harry' really was by none other than #InvictusGames President Prince Harry!"	Invictus Games Foundation	Noteworthy tweets of 2014 - Twitter
30	6/4/2014	"Denny JA: Dengan RT ini, anda ikut memenangkan Jokowi-JK. Pilih pemimpin yg bisa dipercaya (Jokowi) dan pengalaman (JK). #DJoJK"	Denny Januar Ali	List of Most Retweeted Tweets - Wikipedia
31	7/8/2014	"Brasil, "levanta, sacode a poeira e dá a volta por cima"	Dilma Rousseff	Noteworthy tweets of 2014 - Twitter
32	7/8/2014	"Wow wow wow this is crazy. Whens #BRA going to turn up. I know predict a few red cards now @FIFAWorldCup"	Tim Cahill	Noteworthy tweets of 2014 - Twitter
33	8/13/2014	"It's time! The first ever premiere of "A Place With No Name" right now on Twitter"	Michael Jackson	Noteworthy tweets of 2014 - Twitter
34	9/8/2014	"The Duke and Duchess of Cambridge are very pleased to announce that The Duchess of Cambridge is expecting their second child"	Clarence House	Noteworthy tweets of 2014 - Twitter
35	9/21/2014	"Wonderful men out there. I'm launching a campaign - #heforshe. Support the women in ur lives and sign up here now!"	Emma Watson	Noteworthy tweets of 2014 - Twitter

This appendix presents the list of viral tweets that overlap with the event window of a rumor in our sample.

Figure 1.
Relationship between rumor-day tweets and market reaction

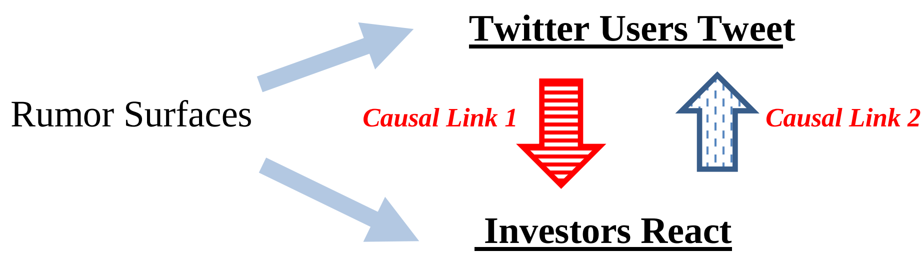

This figure plots the relationship between rumor-related tweet volume and abnormal returns.

Table 1. Sample selection

Rumor events identified in RavenPack since 2008	590
Plus: Rumors from Ahern and Sosyura (2015) sample	70
Less: Duplicates in 60-day window / Events related to foreign subsidiaries and non-US Firms	−186
Less: Misclassified rumors (e.g., already done deal)	−28
Less: Rumors in 2008 (Crimson Hexagon coverage limitations)	−33
Less: Rumors with missing financial data (Compustat/CRSP)	−61
Less: Rumors with missing standardized abnormal Twitter measure	−48
Final rumor sample	304

This table presents the sample selection procedure.

Table 2. Sample descriptive statistics

Panel A: Distribution of rumor sample by year

Year	Number of Rumors	Percentage (%) of Rumors	Number of Realized Rumors	Percentage (%) of Realized Rumors
2009	19	6.26	3	0.98
2010	68	22.37	18	5.92
2011	53	17.43	12	3.95
2012	53	17.43	11	3.62
2013	44	14.47	11	3.62
2014	67	22.04	10	3.29
Total	304	100.00	65	21.38

Panel B: Distribution of rumor sample by merger stage

Merger Stage	Number of Rumors	Percentage (%)
Preliminary Talks	8	2.63
In talk	25	8.22
Made offer	5	1.64
Preparing bid	5	1.64
For sale	11	3.62
Evaluating bid	12	3.95
In advance talk	19	6.26
Speculation	166	54.61
Opinion Pieces	53	17.43
Total	304	100.00

Panel C: Summary statistics of main variables

Variable	N	Mean	Std Dev	25th Pctl	Median	75th Pctl
Tweet and stock return variables						
TWEET_RUMOR $_{[0]}$	304	5.99	11.88	0.22	1.96	6.39
CAR_RUMOR $_{[0]}$	304	4.63	7.29	0.14	2.45	6.93
CAR_RUMOR $_{[0, +1]}$	304	5.22	7.94	0.24	3.05	7.30
Target characteristic variables						
CAR_PRE $_{[-5, -1]}$	304	1.32	6.98	-2.31	0.50	3.72
VALUABLE_BRAND	304	0.10	0.30	0.00	0.00	0.00
LN_ASSET	304	8.09	1.79	6.72	8.17	9.24
TOBINQ	304	2.40	1.85	1.25	1.70	2.76
R&D	304	0.06	0.10	0.00	0.01	0.08
ADV	304	0.02	0.05	0.00	0.00	0.03
TAKEOVER_PROBABILITY	291	4.98	1.71	3.60	4.82	6.25
LOW_INST	304	0.10	0.30	0.00	0.00	0.00

Article characteristic variables						
SPECULATIVE	304	0.89	0.67	0.00	1.00	1.00
BIDDER_MENTIONED	304	0.73	0.44	0.00	1.00	1.00
PRICE_MENTIONED	304	0.16	0.37	0.00	0.00	0.00
SPECIFIC	304	0.09	0.28	0.00	0.00	0.00
OULET_RANK	304	1.65	0.92	1.00	1.00	2.00
#MEDIA_ARTICLES $_{[0, +1]}$	304	4.80	7.25	0.00	2.00	6.50

Panel D: Distribution of rumor sample by outlet rank

	OUTLET_RANK		
	Rank	Number of rumors	Percentage (%)
More reputable (e.g., "Wall Street Journal")	1	175	57.57
↓	2	86	28.29
	3	18	5.92
Less reputable (e.g., "www.fool.com")	4	25	8.22
		304	100

This table presents descriptive statistics. All variables are defined in Appendix A.

Table 3. Correlation table

	Variable	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*	*12*
1	TWEET_RUMOR $_{[0]}$	1.00	0.39*	0.37*	0.06	−0.20*	−0.13*	−0.09	−0.06	−0.04	0.09	0.30*	0.11*
2	CAR_RUMOR $_{[0]}$	0.53*	1.00	0.83*	−0.14*	−0.20*	−0.27*	−0.02	0.12*	−0.04	0.18*	0.15*	0.04
3	CAR_RUMOR $_{[0, +1]}$	0.41*	0.84*	1.00	−0.06	−0.22*	−0.26*	−0.02	0.05	−0.06	0.19*	0.23*	0.06
4	CAR_PRE $_{[-5, -1]}$	−0.04	−0.12*	−0.12*	1.00	−0.00	−0.07	0.05	−0.02	0.06	−0.04	0.01	0.09
5	VALUABLE_BRAND	−0.13*	−0.15*	−0.14*	−0.02	1.00	0.36*	0.04	0.06	0.26*	−0.18*	0.14*	−0.03
6	LN_ASSET	−0.11*	−0.24*	−0.20*	−0.06	0.40*	1.00	−0.46*	−0.35*	−0.02	−0.04	0.25*	−0.03
7	TOBINQ	−0.00	0.01	−0.01	0.10*	−0.02	−0.44*	1.00	0.38*	0.00	−0.27*	−0.14*	0.08
8	R&D	−0.06	0.10*	0.06	−0.02	−0.02	−0.41*	0.34*	1.00	−0.04	0.14*	−0.08	0.02
9	ADV	0.02	0.05	0.02	0.07	0.04	−0.13*	0.07	−0.05	1.00	−0.11*	0.15*	−0.06
10	TAKEOVER_PROBABILITY	0.09	0.21*	0.19*	−0.05	−0.17*	−0.07	−0.30*	0.18*	−0.08	1.00	0.01	−0.05
11	#MEDIA_ARTICLES $_{[0, +1]}$	0.20*	0.13*	0.22*	0.07	0.14*	0.23*	−0.13*	−0.14*	0.04	−0.01	1.00	0.14*
12	SPECIFIC	0.05	0.10*	0.14*	0.11*	−0.03	−0.03	0.01	−0.06	−0.02	−0.07	0.19*	1.00

This table presents Pearson and Spearman correlation coefficients of all variables used in the regression analyses. Pearson correlations are reported on the left bottom corner and Spearman correlations are reported on the right top corner. * denotes significance level at less than 10%. All variables are defined in Appendix A.

Table 4. Twitter activities and merger realization: logistic regressions

DV: Probability of Merger Realization			
	(1)	(2)	(3)
TWEET_RUMOR $_{[0]}$	**0.016**	**−0.002**	**0.007**
	(1.269)	**(−0.148)**	**(0.505)**
VALUABLE_BRAND	−0.257	−0.445	−0.278
	(−0.337)	(−0.429)	(−0.288)
LN_ASSET	−0.326**	−0.362**	−0.401***
	(−2.487)	(−2.334)	(−2.692)
TOBINQ	−0.086	−0.036	−0.066
	(−0.734)	(−0.242)	(−0.423)
R&D	−0.591	−0.936	−0.773
	(−0.255)	(−0.357)	(−0.318)
ADV	−0.109	−3.393	−3.268
	(−0.023)	(−0.516)	(−0.531)
CAR_PRE $_{[-5, -1]}$	0.065***	0.069***	0.066**
	(2.693)	(2.606)	(2.542)
TAKEOVER_PROBABILITY		0.055	0.046
		(0.311)	(0.255)
#MEDIA_ARTICLES $_{[0, +1]}$		0.041	0.043*
		(1.550)	(1.796)
SPECIFIC			1.529***
			(3.228)
SPECULATIVE		−0.369	
		(−1.395)	
BIDDER_MENTIONED		−0.388	
		(−0.905)	
PRICE_MENTIONED		1.092**	
		(2.446)	
Intercept	−0.654	−0.081	−0.326
	(−0.410)	(−0.040)	(−0.156)
Industry / Year / Outlet FE	YES	YES	YES
Observations	304	291	291
R-squared	0.141	0.199	0.190

This table presents logistic regression results for the following equation:
$REALIZED = \beta_0 + \beta_1\, TWEET\_RUMOR_{[0]} + Controls + \varepsilon$. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Please see Appendix A for variable definitions. Year, industry, and outlet fixed effects are included. Standard errors are clustered at the firm level.

Table 5. CARs around rumor announcement dates

		(1)	(2)	(3)
		All	**Realized**	**Unrealized**
Panel A: Run-up period				
Days [-20, -1]	All	0.82	5.62**	−0.50
	High Twitter Portfolio	2.26*	6.30**	−0.12
	Low Twitter Portfolio	−0.61	4.97*	−0.86
Panel B: Event period				
Days [0, 1]	All	5.22***	9.12***	4.16***
	High Twitter Portfolio	7.57***	13.14***	5.74***
	Low Twitter Portfolio	2.87***	5.21***	2.59***
Day [0]	All	4.63***	6.64***	4.08***
	High Twitter Portfolio	7.00***	10.59***	6.07***
	Low Twitter Portfolio	2.26***	2.81***	2.10***
Panel C: Post-event period				
Days [2, 20]	All	−2.05***	0.18	−2.66***
	High Twitter Portfolio	−2.84***	0.23	−3.24***
	Low Twitter Portfolio	−1.27	0.13	−2.10**
Days [21, 40]	All	−0.03	0.67	−0.23
	High Twitter Portfolio	−1.00	0.02	−1.29
	Low Twitter Portfolio	0.93	1.31	0.82
Days [41, 60]	All	−0.73	−1.05	−0.65
	High Twitter Portfolio	−1.58*	−0.19	−1.85*
	Low Twitter Portfolio	0.09	−1.82	0.51

This table reports average cumulative abnormal returns (*CARs*) in percentages for target firms. High/Low Twitter portfolios are formed by a median split on abnormal Twitter volume on the rumor day [0]. Abnormal returns are defined as raw returns in excess of CRSP value-weighted market returns. Abnormal returns are winsorized at the 1% and 99% levels. Realized rumors are defined as rumors in which an official takeover announcement was made within one year of the rumor publication date. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test.

Table 6. Event-window CARs and Twitter activity

DV: CARs for target firms at rumor announcement dates [0, +1]						
	(1)	(2)	(3)	(4)	(5)	(6)
	All Sample		**Realized Rumors**		**Unrealized Rumors**	
TWEET_RUMOR $_{[0]}$	**0.290*** **	**0.249*** **	**0.493*** **	**0.462*** **	**0.258*** **	**0.237*** **
	(4.322)	**(3.784)**	**(3.439)**	**(3.011)**	**(3.304)**	**(2.901)**
VALUABLE_BRAND	−0.371	0.003	−15.009**	−9.023	0.979	1.184
	(−0.321)	(0.003)	(−2.668)	(−1.554)	(0.881)	(0.967)
CAR_PRE $_{[-5, -1]}$	−0.036	−0.033	−0.496***	−0.377**	0.033	0.062
	(−0.473)	(−0.416)	(−2.708)	(−2.328)	(0.338)	(0.624)
LN_ASSET	−0.974***	−1.234***	−0.748	−1.165	−1.092***	−1.247***
	(−2.832)	(−3.544)	(−0.890)	(−1.218)	(−3.380)	(−3.420)
TOBINQ	−0.647**	−0.477	−1.095	−0.401	−0.798***	−0.687**
	(−2.226)	(−1.552)	(−1.007)	(−0.382)	(−2.785)	(−2.109)
R&D	1.000	−1.823	−7.830	−15.879	−2.596	−3.216
	(0.165)	(−0.321)	(−0.457)	(−0.837)	(−0.381)	(−0.470)
ADV	7.186	−5.369	25.757	10.778	−3.313	−10.890
	(0.621)	(−0.420)	(1.281)	(0.471)	(−0.335)	(−0.934)
TAKEOVER_PROBABILITY		0.569		1.283		0.477
		(1.470)		(0.807)		(1.238)
#MEDIA_ARTICLES $_{[0, +1]}$		0.266***		0.375*		0.125
		(2.948)		(1.929)		(1.205)
SPECIFIC		2.480		−0.891		1.301
		(1.361)		(−0.202)		(0.606)
Intercept	10.558***	9.572**	12.605	1.874	11.333***	10.381**
	(2.838)	(2.103)	(1.189)	(0.128)	(3.218)	(2.220)
Industry / Year / Outlet FE	YES	YES	YES	YES	YES	YES
Observations	304	291	65	63	239	228
R-squared	0.315	0.389	0.629	0.682	0.290	0.318

This table presents OLS regression results for the following equation:
$CAR\_RUMOR_{[0, +1]} = \beta_0 + \beta_1 \, TWEET\_RUMOR_{[0]} + Controls + \varepsilon$. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Please see Appendix A for variable definitions. Year, industry, and outlet fixed effects are included. Standard errors are clustered at the firm level.

Table 7. Price reversals and Twitter activity: unrealized rumors

DV: CARs for target firms after rumor announcement dates			
	(1)	(2)	(3)
	CAR $_{[+2, +20]}$	**CAR $_{[+21, +40]}$**	**CAR $_{[+41, +60]}$**
TWEET_RUMOR $_{[0]}$	**−0.009**	**−0.038**	**−0.110****
	(−0.191)	**(−0.520)**	**(−2.013)**
Intercept	0.880	1.186	7.199
	(0.115)	(0.141)	(0.820)
Target characteristic controls	YES	YES	YES
Rumor characteristic controls	YES	YES	YES
Industry / Year / Outlet FE	YES	YES	YES
Observations	228	226	224
R-squared	0.114	0.074	0.130

This table presents OLS regression results for the following equation:
$CAR_{[+2, +20]}$ (or $CAR_{[+21, +40]}$, or $CAR_{[+41, +60]}$) $= \beta_0 + \beta_1$ *TWEET_RUMOR* $_{[0]}$ + *Controls* + ε.

"Target characteristic controls" includes: *VALUABLE_BRAND*, *CAR_PRE* $_{[-5, -1]}$, *LN_ASSET*, *TOBINQ*, *R&D*, *ADV*, and *TAKEOVER_PROBABILITY*. "Rumor characteristic controls" includes: *#MEDIA_ARTICLES* $_{[0, +1]}$ and *SPECIFIC*. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Please see Appendix A for variable definitions. Year, industry, and outlet fixed effects are included. Standard errors are clustered at the firm level.

Table 8. Cross-sectional tests

Panel A: Effects of high influence score– unrealized rumors

DV: CARs for target firms at rumor announcement dates [0, +1]		
	(1)	(2)
TWEET_RUMOR $_{[0]}$	0.160*	0.134
	(1.855)	(1.620)
TWEET_RUMOR $_{[0]}$ * HIGH_INFLUENCE	**0.201***	**0.210****
	(1.935)	**(2.045)**
HIGH_INFLUENCE	0.218	0.423
	(0.217)	(0.396)
Intercept	11.198***	9.983**
	(3.388)	(2.275)
Target characteristic controls	YES	YES
Rumor characteristic controls	NO	YES
Industry / Year / Outlet FE	YES	YES
Observations	239	228
R-squared	0.320	0.352

Panel B: Effects of institutional ownership– unrealized rumors

DV: CARs for target firms at rumor announcement dates [0, +1]		
	(1)	(2)
TWEET_RUMOR $_{[0]}$	0.199***	0.175**
	(2.809)	(2.505)
TWEET_RUMOR $_{[0]}$ * LOW_INST	**0.273*****	**0.257*****
	(3.450)	**(3.247)**
LOW_INST	-1.157	1.519
	(-0.760)	(0.636)
Intercept	11.564***	8.660*
	(3.305)	(1.857)
Target characteristic controls	YES	YES
Rumor characteristic controls	NO	YES
Industry / Year / Outlet FE	YES	YES
Observations	239	228
R-squared	0.318	0.354

Panel C: Effects of more specific rumors– unrealized rumors

DV: CARs for target firms at rumor announcement dates [0, +1]		
	(1)	(2)
TWEET_RUMOR $_{[0]}$	0.239***	0.214***
	(3.093)	(2.673)
TWEET_RUMOR $_{[0]}$ * SPECIFIC	**0.525*** **	**0.563*** **
	(3.257)	**(3.489)**
SPECIFIC	−1.106	−1.389
	(−0.466)	(−0.626)
Intercept	11.780***	10.783**
	(3.338)	(2.320)
Target characteristic controls	YES	YES
Rumor characteristic controls	NO	YES
Industry / Year / Outlet FE	YES	YES
Observations	239	228
R-squared	0.314	0.343

These tables present OLS regression results for the following equation:
$CAR\_RUMOR_{[0, +1]} = \beta_0 + \beta_1\ TWEET\_RUMOR_{[0]} + Controls + \varepsilon$, with moderating variables and their interactions with *TWEET_RUMOR* $_{[0]}$ included. "Target characteristic controls" includes: *VALUABLE_BRAND*, *CAR_PRE* $_{[-5, -1]}$, *LN_ASSET*, *TOBINQ*, *R&D*, and *ADV* in Column 1. "Target characteristic controls" includes: *VALUABLE_BRAND*, *CAR_PRE* $_{[-5, -1]}$, *LN_ASSET*, *TOBINQ*, *R&D*, *ADV* and *TAKEOVER_PROBABILITY* in Column 2. "Rumor characteristic controls" includes: *#MEDIA_ARTICLES* $_{[0, +1]}$ and *SPECIFIC*. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Please see Appendix A for variable definitions. Year, industry, and outlet fixed effects are included. Standard errors are clustered at the firm level.

Table 9. Rumors in non-trading windows and subsequent abnormal returns

Panel A: Rumors on non-trading days

Rumors on Non-Trading Days: Friday after market, Saturday, and Sunday			
DV: CARs for target firms at day +1 (first available trading day after rumor day)			
	(1)	(2)	(3)
TWEET_RUMOR $_{[0]}$	**1.727*** **	**1.904*** **	**1.717*** **
	(3.502)	**(3.331)**	**(2.818)**
LN_ASSET		−0.628	−3.384
		(−1.046)	(−1.698)
Intercept	−10.796*	−5.904	29.728
	(−1.722)	(−0.926)	(1.461)
Target characteristic controls	NO	NO	YES
Industry / Year FE	YES	YES	YES
Observations	27	27	27
R-squared	0.746	0.758	0.871

This table presents OLS regression results for the following equation:
$CAR\_RUMOR_{[0]} = \beta_0 + \beta_1 TWEET\_RUMOR_{[0]} + LN\_ASSET + \varepsilon$, for a sample of rumors that surfaced during non-trading days (i.e., Friday after market, Saturday, and Sunday). If a rumor surfaced on Friday after market, *TWEET_RUMOR* $_{[0]}$ is measured for Friday after market, Saturday, and Sunday. If a rumor surfaced on Saturday, *TWEET_RUMOR* $_{[0]}$ is measured for Saturday and Sunday. If a rumor surfaced on Sunday, *TWEET_RUMOR* $_{[0]}$ is measured for Sunday. "Target characteristic controls" includes: *VALUABLE_BRAND*, *CAR_PRE* $_{[-5,-1]}$, *LN_ASSET*, *TOBINQ*, *R&D*, and *ADV*. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Year and industry fixed effects are included. Standard errors are clustered at the firm level.

Panel B: Rumors during non-trading hours

Rumors after-market hours: 4:31-11:59 pm and 0:00-9:29 am			
DV: CARs for target firms at day +1 (next available trading day after market closure)			
	(1)	(2)	(3)
TWEET_RUMOR $_{[0]}$	**0.369*** **	**0.379*** **	**0.486*** **
	(3.455)	**(4.425)**	**(3.734)**
LN_ASSET		1.183	2.075
		(1.465)	(1.669)
Intercept	6.118	−2.600	−12.396
	(1.593)	(−0.341)	(−1.035)
Target characteristic controls	NO	NO	YES
Industry / Year FE	YES	YES	YES
Observations	36	36	36
R-squared	0.742	0.783	0.843

This table presents OLS regression results for the following equation:
$CAR\_RUMOR_{[0]} = \beta_0 + \beta_1 TWEET\_RUMOR_{[0]} + LN\_ASSET + \varepsilon$, for a sample of rumors that surfaced during non-market hours (i.e., pre-market from 0:00 to 9:29 am, and post-market from 4:01 to 11:59 pm). If a rumor surfaced on Day *t* during pre-market hours (0:00-9:29 am), *TWEET_RUMOR* $_{[0]}$ is measured for the pre-market window. If a rumor surfaced on Day *t* during post-market hours (4:01-11:59 pm), abnormal Twitter activity is measured for the post-market window on Day *t* and for the pre-market window on Day *t* + 1 separately. *TWEET_RUMOR* $_{[0]}$ is then obtained by summing up the two. "Target characteristic controls" includes: *VALUABLE_BRAND*, *CAR_PRE* $_{[-5,-1]}$, *LN_ASSET*, *TOBINQ*, *R&D*, and *ADV*. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Year and industry fixed effects are included. Standard errors are clustered at the firm level.

Table 10. 2SLS: Twitter outages and viral tweets

Panel A: 2SLS analysis with Twitter outage

	First Stage: DV= TWEET_RUMOR $_{[0]}$	**Second Stage: DV= CAR_RUMOR $_{[0, +1]}$**	**First Stage: DV= TWEET_RUMOR $_{[0]}$**	**Second Stage: DV= CAR_RUMOR $_{[0, +1]}$**
	Generic Twitter Outages		Severe Twitter Outages	
	(1)	(2)	(3)	(4)
Predicted TWEET_RUMOR $_{[0]}$		**0.292**		**0.357***
		(1.082)		**(1.695)**
OUTAGE	**−3.751***			
	(−1.810)			
OUTAGE_SEVERE			**−8.283*****	
			(−3.779)	
Intercept	2.495	10.244**	1.609	10.084**
	(0.322)	(2.280)	(0.210)	(2.248)
Target characteristic controls	YES	YES	YES	YES
Rumor characteristic controls	YES	YES	YES	YES
Industry / Year / Outlet FE	YES	YES	YES	YES
Observations	228	228	228	228
R-squared	0.190	0.312	0.196	0.288

F-test of excluded instruments for first stage regression in Column 1: F (1, 177) = 3.28, Prob>F = 0.0720
F-test of excluded instruments for first stage regression in Column 3: F (1, 177) = 14.28, Prob>F = 0.0002

This table presents 2SLS regression results for the following equation:
First Stage: TWEET RUMOR $_{[0]} = \beta_0 + \beta_1$ *OUTAGE (OUTAGE_SEVERE) + Controls +* ε.
Second Stage: CAR_RUMOR $_{[0, +1]} = \beta_0 + \beta_1$ *predicted TWEET_RUMOR* $_{[0]}$ *+ Controls +* ε.

In Columns 1-2, *OUTAGE* is defined as one for all Twitter outage episodes with duration (when available) longer than 45 minutes; zero otherwise. In Columns 3-4, *OUTAGE_SEVERE* is defined as one for *OUTAGE* episodes involving (1) site crashes, or high/elevated error rates on Twitter, or (2) sign-in/access/connectivity/stability issues. It is zero otherwise.

"Target characteristic controls" includes: *VALUABLE_BRAND, CAR_PRE* $_{[-5, -1]}$, *LN_ASSET, TOBINQ, R&D, ADV,* and *TAKEOVER_PROBABILITY.* "Rumor characteristic controls" includes: *#MEDIA_ARTICLES* $_{[0, +1]}$ and *SPECIFIC.* ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Please see Appendix A for variable definitions. Year, industry and outlet fixed effects are included. Standard errors are clustered at the firm level.

Panel B: 2SLS analysis with viral tweets

	First Stage: DV= TWEET_RUMOR $_{[0]}$	Second Stage: DV= CAR_RUMOR $_{[0, +1]}$	First Stage: DV= TWEET_RUMOR $_{[0]}$	Second Stage: DV= CAR_RUMOR $_{[0, +1]}$
	Viral Tweets		Severe Twitter Outages and Viral Tweets	
	(1)	(2)	(3)	(4)
Predicted TWEET_RUMOR $_{[0]}$		**0.329***		**0.342****
		(1.881)		**(2.479)**
VIRAL_TWEET	**−4.865*****		**−4.715*****	
	(−3.092)		**(−3.122)**	
OUTAGE_SEVERE			**−7.986*****	
			(−4.221)	
Intercept	1.837	10.154**	1.004	10.121**
	(0.243)	(2.335)	(0.133)	(2.301)
Target characteristic controls	YES	YES	YES	YES
Rumor characteristic controls	YES	YES	YES	YES
Industry / Year / Outlet FE	YES	YES	YES	YES
Observations	228	228	228	228
R-squared	0.199	0.300	0.218	0.295

F-test of excluded instruments for first stage regression in Column 1: F (1, 177) = 9.56, Prob>F = 0.0023
F-test of excluded instruments for first stage regression in Column 3: F (2, 177) = 10.96, Prob>F = 0.0000

This table presents 2SLS regression results for the following equations:
First Stage: TWEET RUMOR $_{[0]} = \beta_0 + \beta_1$*VIRAL_TWEET + Controls* + ε.
Second Stage: CAR_RUMOR $_{[0, +1]} = \beta_0 + \beta_1$ *predicted TWEET_RUMOR* $_{[0]}$ *+ Controls* + ε.

First Stage: TWEET RUMOR $_{[0]} = \beta_0 + \beta_1$ *OUTAGE_SEVERE* + β_2*VIRAL_TWEET + Controls* + ε.
Second Stage: CAR_RUMOR $_{[0, +1]} = \beta_0 + \beta_1$ *predicted TWEET_RUMOR* $_{[0]}$ *+ Controls* + ε.

"Target characteristic controls" includes: *VALUABLE_BRAND*, *CAR_PRE* $_{[-5, -1]}$, *LN_ASSET*, *TOBINQ*, *R&D*, *ADV* and *TAKEOVER_PROBABILITY*. "Rumor characteristic controls" includes: *#MEDIA_ARTICLES* $_{[0, +1]}$ and *SPECIFIC*. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Please see Appendix A for variable definitions. Year, industry, and outlet fixed effects are included. Standard errors are clustered at the firm level.

Table 11. CAR and Twitter activities - realized rumors

DV: CARs for target firms at and after rumor announcement dates / at takeover announcement dates					
	(1)	(2)	(3)	(4)	(5)
	CAR_RUMOR $_{[0, +1]}$	**CAR $_{[+2, +20]}$**	**CAR $_{[+21, +40]}$**	**CAR $_{[+41, +60]}$**	**CAR_TAKEOVER_ANN $_{[0, +1]}$**
TWEET_RUMOR $_{[0]}$	**0.462***	**−0.129**	**0.043**	**0.158**	**−0.838***
	(3.011)	**(−0.563)**	**(0.244)**	**(1.072)**	**(−1.748)**
Intercept	1.874	2.016	−15.146	−18.411	93.773*
	(0.128)	(0.067)	(−0.736)	(−1.039)	(1.805)
Target characteristic controls	YES	YES	YES	YES	YES
Rumor characteristic controls	YES	YES	YES	YES	YES
Industry / Year / Outlet FE	YES	YES	YES	YES	YES
Observations	63	63	63	61	63
R-squared	0.682	0.261	0.409	0.596	0.445

This table presents OLS regression results for the following equation:
CAR_RUMOR $_{[0, +1]}$ (*CAR* $_{[+2, +20]}$, or *CAR* $_{[+21, +40]}$, *or CAR* $_{[+41, +60]}$, *or CAR_TAKEOVER_ANN* $_{[0, +1]}$) $= \beta_0 + \beta_1$ *TWEET_RUMOR* $_{[0]}$ + *Controls* + ε.

"Target characteristic controls" includes: *VALUABLE_BRAND*, *CAR_PRE* $_{[-5, -1]}$, *LN_ASSET*, *TOBINQ*, *R&D*, *ADV*, and *TAKEOVER_PROBABILITY*. "Rumor characteristic controls" includes: *#MEDIA_ARTICLES* $_{[0, +1]}$ and *SPECIFIC*. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Please see Appendix A for variable definitions. Year, industry, and outlet fixed effects are included. Standard errors are clustered at the firm level.

Table 12. Journalist and article characteristics

Panel A: Journalist Characteristics

	DV: Probability of Merger Realization		DV: CARs for target firms at rumor announcement dates [0, +1]			
			All Sample		Unrealized Rumors	
	(1)	(2)	(3)	(4)	(5)	(6)
TWEET_RUMOR $_{[0]}$	**0.005**	**0.015**	**0.456*** **	**0.464*** **	**0.399*** **	**0.421*** **
	(0.241)	**(0.609)**	**(7.256)**	**(5.631)**	**(6.882)**	**(5.680)**
LOG (JOURNALIST AGE)	2.390	1.414	−3.479	−5.090	0.293	−1.754
	(1.176)	(0.629)	(−0.964)	(−1.438)	(0.081)	(−0.474)
Undergraduate Degree						
BUSINESS & ECONOMICS	−0.679	0.197	0.320	−0.614	2.424	1.117
	(−0.605)	(0.212)	(0.094)	(−0.192)	(0.678)	(0.323)
JOURNALISM	1.253	2.217**	−1.419	−1.736	−1.155	−1.857
	(1.196)	(2.024)	(−0.676)	(−0.877)	(−0.440)	(−0.789)
ENGLISH	1.159	1.649	3.362	2.058	2.900	1.016
	(1.620)	(1.574)	(1.459)	(1.012)	(1.154)	(0.444)
POLITICAL SCIENCE	1.457	1.942	1.625	−0.181	1.523	−0.440
	(1.606)	(1.593)	(0.664)	(−0.083)	(0.671)	(−0.196)
HISTORY	1.173	1.441	−2.301	−3.597*	−2.698	−3.587
	(1.465)	(1.586)	(−1.124)	(−1.959)	(−0.991)	(−1.466)
OTHER	−0.613	−0.771	−0.620	−0.366	1.067	0.330
	(−0.747)	(−0.812)	(−0.251)	(−0.168)	(0.383)	(0.129)
EXPERT in TARGET INDUSTRY	−2.219	−2.838	1.757	2.215	1.772	2.741
	(−1.352)	(−1.572)	(0.721)	(0.963)	(0.822)	(1.358)
NY BASED	0.955	1.008	−0.017	−0.873	−0.103	−0.751
	(1.004)	(0.948)	(−0.011)	(−0.590)	(−0.072)	(−0.574)
AWARD WINNER	1.082	0.729	4.882*	4.754*	7.371**	6.176*
	(1.421)	(0.979)	(1.928)	(1.893)	(2.383)	(1.913)
GENDER	−0.935	−1.302*	1.471	0.618	1.961	1.162
	(−1.567)	(−1.778)	(0.870)	(0.393)	(1.046)	(0.713)
CAR_PRE $_{[-5,\,-1]}$	0.074*	0.087***	0.047	0.058	0.134	0.138
	(1.905)	(2.695)	(0.420)	(0.574)	(1.037)	(1.218)
LN_ASSET	−0.137	0.095	0.067	−0.407	0.508	−0.462
	(−0.650)	(0.327)	(0.148)	(−0.679)	(1.040)	(−0.730)
GRADUATE DEGREE		0.318		1.704		1.748
		(0.395)		(1.000)		(1.114)
COLUMNIST		1.910**		1.339		−0.896
		(2.558)		(0.827)		(−0.546)
Intercept	−13.510	−15.433	11.059	19.708	−8.673	9.715
	(−1.552)	(−1.350)	(0.733)	(1.236)	(−0.571)	(0.602)
Target characteristic controls	NO	YES	NO	YES	NO	YES
Industry / Year FE	YES	YES	YES	YES	YES	YES
Outlet FE	NO	NO	YES	YES	YES	YES
Observations	118	118	118	118	90	90
R-squared	0.339	0.423	0.553	0.633	0.615	0.688

This table presents logistic and OLS regressions results for the following equations:
$REALIZED = \beta_0 + \beta_1\ TWEET\_RUMOR_{[0]} + Controls + \varepsilon$
$CAR\_RUMOR_{[0,+1]} = \beta_0 + \beta_1\ TWEET\_RUMOR_{[0]} + Controls + \varepsilon,$
with journalist characteristic variables included. “Target characteristic controls” includes: *VALUABLE_BRAND*, *TOBINQ*, *R&D*, and *ADV*. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Please see Appendix A for variable definitions. Standard errors are clustered at the firm level.

Panel B: Article Characteristics

	DV: Probability of Merger Realization		DV: CARs for target firms at rumor announcement dates [0, +1]			
			All Sample		Unrealized Rumors	
	(1)	(2)	(3)	(4)	(5)	(6)
TWEET_RUMOR $_{[0]}$	**0.000**	**0.003**	**0.234*** **	**0.236*** **	**0.214*** **	**0.218** **
	(0.004)	**(0.221)**	**(3.518)**	**(3.299)**	**(2.668)**	**(2.442)**
WEAK MODAL WORDS (%)	0.005	0.018	−0.224	−0.159	−0.279	−0.247
	(0.030)	(0.118)	(−0.857)	(−0.645)	(−0.732)	(−0.714)
ANONYMOUS SOURCE	0.817*	0.896*	1.111	1.181	0.506	0.407
	(1.730)	(1.846)	(1.280)	(1.338)	(0.530)	(0.402)
Target Response						
HAS CONVERSATION	0.175	0.273	−6.545**	−5.721**	−7.232***	−6.775**
	(0.157)	(0.242)	(−2.575)	(−1.983)	(−2.880)	(−2.501)
CONFIRMED RUMOR	—	—	2.732	1.979	–	–
			(0.747)	(0.571)		
DENIED RUMOR	—	—	−7.376***	−7.107**	−7.285***	−7.572***
			(−3.496)	(−2.151)	(−3.148)	(−2.729)
COULDN'T BE REACHED	0.062	0.048	−1.028	−1.182	−0.841	−0.971
	(0.100)	(0.075)	(−0.831)	(−0.888)	(−0.560)	(−0.580)
WASN'T ASKED	−0.499	−0.506	−0.987	−0.998	−1.695	−1.596
	(−1.040)	(−0.988)	(−0.853)	(−0.874)	(−1.161)	(−1.094)
Merger Stage						
PRELIMINARY TALK	0.852	0.905	−3.598***	−3.227**	−3.087*	−2.001
	(1.046)	(1.128)	(−2.615)	(−2.275)	(−1.858)	(−0.766)
IN TALKS	−0.314	−0.995	3.010*	2.321	3.071	2.872
	(−0.570)	(−1.262)	(1.680)	(1.145)	(1.630)	(1.531)
MADE OFFER	−0.081	−1.020	3.643	3.217	−1.304	−0.724
	(−0.070)	(−0.775)	(0.874)	(0.688)	(−0.537)	(−0.215)
PREPARING BID	−1.800	−1.362	2.471	2.751	2.693	3.448
	(−1.329)	(−1.133)	(0.598)	(0.678)	(0.493)	(0.648)
FOR SALE	−0.155	−0.106	3.619	3.372	5.609	5.475
	(−0.144)	(−0.095)	(1.301)	(1.180)	(1.629)	(1.570)
EVALUATING BIDS	−0.303	−0.680	0.514	0.111	0.594	−0.071
	(−0.413)	(−0.850)	(0.321)	(0.062)	(0.371)	(−0.038)
#MEDIA_ARTICLES $_{[0, +1]}$	0.059**	0.061**	0.230***	0.239***	0.128	0.120
	(2.265)	(2.380)	(2.746)	(2.680)	(1.376)	(1.207)
RUMOR IN HEADLINE	0.179	0.226	1.204	1.040	0.988	0.634
	(0.318)	(0.413)	(1.173)	(0.929)	(0.855)	(0.534)
BIDDER MENTION	−0.516	−0.549	0.724	1.463	0.072	0.529
	(−1.182)	(−1.245)	(0.895)	(1.609)	(0.095)	(0.632)
PRICE MENTION	1.193**	0.616	−0.185	−0.063	−0.693	−0.803
	(2.317)	(0.920)	(−0.131)	(−0.032)	(−0.466)	(−0.378)
CAR_PRE $_{[-5, -1]}$	0.074***	0.076***	−0.080	−0.041	−0.007	0.057
	(2.772)	(2.870)	(−1.103)	(−0.527)	(−0.079)	(0.636)
LN_ASSET	−0.373***	−0.389**	−0.998***	−1.192***	−0.755***	−1.118***
	(−3.028)	(−2.406)	(−3.990)	(−3.352)	(−3.103)	(−3.082)
# OF NUMBERS		0.007		0.042		0.031
		(0.379)		(0.891)		(0.520)
SPECIFIC		1.579		−0.206		−0.015
		(1.570)		(−0.065)		(−0.004)
Intercept	−1.359	−1.481	8.908***	7.083	8.547**	10.080*
	(−0.661)	(−0.627)	(3.260)	(1.430)	(2.587)	(1.835)
Target characteristic controls	NO	YES	NO	YES	NO	YES
Industry / Year FE	YES	YES	YES	YES	YES	YES
Outlet FE	NO	NO	YES	YES	YES	YES
Observations	300	300	304	291	239	228
R-squared	0.205	0.216	0.407	0.436	0.352	0.386

This table presents logistic and OLS regressions results for the following equations:
REALIZED= $\beta_0 + \beta_1$ *TWEET_RUMOR* $_{[0]}$ + *Controls* + ε;
CAR RUMOR $_{[0, +1]}$ = $\beta_0 + \beta_1$ *TWEET_RUMOR* $_{[0]}$ + *Controls* + ε,

with article characteristic variables included. "Target characteristic controls" includes: *VALUABLE_BRAND*, *TOBINQ*, *R&D*, and *ADV*. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, using a two-tailed test. T-statistics are reported in parenthesis. Please see Appendix A for variable definitions. Standard errors are clustered at the firm level.