MENDELU Working Papers
in Business and Economics

92/2023

𝕏 Bots and Earnings Announcements

Jan Hanousek, Jr., Jan Hanousek, and Konstantin Sokolov

MENDELU Working Papers in Business and Economics
Research Centre
Faculty of Business and Economics
Mendel University in Brno
Zemědělská 1, 613 00 Brno
Czech Republic
http://vyzc.pef.mendelu.cz/en
+420 545 132 605

Citation

Hanousek, J., Jr. et al. (2023). 𝕏 Bots and Earnings Announcements. *MENDELU Working Papers in Business and Economics* 92/2023. Mendel University in Brno. Cited from: http://ideas.repec.org/s/men/wpaper.html

Abstract

Jan Hanousek, Jr. et al.: **𝕏 Bots and Earnings Announcements**

This paper studies the rationale and effects of buying bots on 𝕏 (former Twitter). We observe that a large amount of attention to corporate 𝕏 accounts around earnings announcements is driven by bots. Bot activity is a significant predictor of investor disagreement, which is persistent in the long term. Moreover, bot activity increases analyst dispersion for the following quarterly earnings announcement. Consistent with managerial short-termism, bot activity often accompanies intense earnings management. Our results are robust to various specifications, including a matching approach indicating causal interpretation.

Keywords

Analysts' dispersion, artificial attention increase, earnings management, investors' disagreement, quarterly earnings, 𝕏 bots, Twitter bots.

JEL: G11, G14, G30, D22, C26

Contacts

Jan Hanousek, Jr., Fogelman College of Business Economics, University of Memphis, E-mail: jan.hanousek@memphis.edu. Corresponding author.

Jan Hanousek, Department of Economics, Faculty of Business and Economics, Mendel University in Brno, Zemědělská 1, 613 00 Brno, Czech Republic and CEPR London, jan.hanousek@mendelu.cz.

Konstantin Sokolov, Fogelman College of Business Economics, University of Memphis, E-mail: ksokolov@memphis.edu

Acknowledgments

We are grateful for comments from a brown bag seminar at the University of South Florida.

1. Introduction

𝕏 (formerly Twitter) emphasizes its efforts to bring about a better platform for healthy civic discourse.[1,2] However, these efforts are undermined by a class of software that imitates human communication, commonly referred to as “bots". Multiple studies show that bots can be weaponized to manipulate elections and misinform the public (Bessi and Ferrara, 2016; Ferrara, 2017; Broniatowski et al., 2018; Stella et al., 2018). The corporate sector is also becoming increasingly aware of the power of social media. In fact, companies are willing to purchase 𝕏 bots and pay as much as $20,000 for a single tweet to promote their image on social media (Confessore et al., 2018).

Our paper sheds light on the bot activity around earnings announcements. Specifically, we study the effects and rationale of this artificial increase in attention and reach of the announcements. We observe that companies tweet[3] more frequently during quarterly earnings periods and receive significantly more attention. The lion’s share of this attention is driven by bots.

We find that bot activity is a significant predictor of investor disagreement, both in terms of abnormal volume and stock price volatility. As such, bots play a divisive role among investors. In particular, an increase in bot activity by 1% typically leads to a 0.98% increase in the standard deviation of returns. Bot activity plays a role both in the week before earnings announcements and the week after, with the bot activity increasing volatility and volume in both periods. Moreover, we find that the effect of bot activity on volatility is persistent for longer time horizons, including up to the subsequent earnings announcement.

Since bot activity fuels investor disagreement, it may also undermine a company’s performance. We, therefore, proceed with exploring why companies buy bots. We find that managers are likely responsible for purchasing bots to attract attention to their company’s current financial performance. This activity is especially pronounced when earnings announcements are in line with the expectations,

[1] Available at https://blog.twitter.com/official/en_us/a/2012/shutting-down-spammers.html, last accessed January 18, 2023.

[2] Available at https://blog.twitter.com/official/en_us/topics/company/2018/2016-election-update.html, last accessed January 18, 2023.

[3] Following the renaming of Twitter into 𝕏, tweets became known as posts. However, due to the ambiguity associated with the term post, we continue to use the word tweets.

as such announcements carry little new information. We find that the management may buy bots for their tweets on 𝕏 to emphasize the small amount of information that is present in such announcements.

While it is not possible to identify who is responsible for the bot activity, only insiders are privy to the earnings information prior to the announcement. And we are observing the highest levels of bot activity during earnings, where such activity should have the strongest impact. Moreover, purchasing bots might be a high-risk strategy for non-insiders, who are risking putting a spotlight on underwhelming earnings. Given that corporations are becoming increasingly aware of the importance of social media presence, we believe that our results indicate that managers are most likely suspects behind the purchase of bot activity.

Moreover, consistent with managerial short-termism, we find that bot activity often accompanies intense earnings management. Specifically, bot activity has a significant negative relationship with discretionary spending, which is a common tool influencing short-term returns. We observe that an increase of bot activity by 1% is associated with a mean cut in discretionary expenses by 2.7%, with a more substantial impact in the case of the positive earnings surprise (5%). This finding suggests that the bot activity helps management underscore the company's meeting of short-term goals, which is often achieved through earnings management.

Finally, we explore the long-term effects of bot activity. We find that bot activity leads to higher volatility post-earnings and that these effects persist for up to three months following the quarterly earnings. We find that an increase in bot activity by 1% is associated with a rise in return volatility by 0.57% in the first week after the QE, by 0.76%, and by 0.80% in the second and third weeks after the QE, respectively. In the horizon of two to three months, the mean effect falls to 0.53% and 0.40%, respectively, which shows that the increased bot activity has persistent and economically significant results.

Moreover, we show that not only investors are affected by bot activity. The analyst dispersion, another indicator of investor disagreement (Diether et al., 2002), significantly increases in the next quarter following high bot activity. Market analysts might consider the popularity of firms on social media, consistent with Gerken and Painter (2023). This finding further outlines a bot activity's impact since analysts can only observe a firm's apparent popularity and cannot observe how much of it is driven by bots.

Our paper contributes to the literature in three ways. First, it contributes to the growing literature on social media (e.g., Chen et al., 2014; Farrel et al., 2022; Bradley et al., 2022). Social media is becoming increasingly important, and research suggests it contains value-relevant information (e.g., Chen et al., 2014). Most businesses have official social media accounts (Jung et al., 2015), which they use to communicate with shareholders, e.g., to dampen adverse price reactions to consumer product recalls (Lee et al., 2015). However, it is not clear how bot activity affects social media. Bots can increase investors' informativeness by increasing the social reach (Blankespoor et al., 2014). However, bots can also inflate the apparent popularity and sway the opinion of investors and analysts covering the company. We hypothesize that bots primarily serve to artificially increase the apparent popularity of a firm and expect the effect to be most substantial during earnings announcements. We find that corporations tweet significantly more frequently during quarterly earnings, and their posts receive more attention due to bot activity.

Second, the paper contributes to the growing literature on investor disagreement. Previous research shows that social media increases investors' disagreement (Cookson and Niessner, 2020; Antweiler and Frank, 2004). Traditional finance literature finds that disagreement induces trading (e.g., Hirschleifer, 1977; Diamond and Verrecchia, 1981). Empirically, Antweiler and Frank (2004) show that Yahoo!Finance message board activities and disagreement predict market volatility and trading volume. In a recent paper, Cookson and Niessner (2020) study disagreement on the investing platform StockTwits. They provide further evidence that disagreement increases trading volume. While bot activity effects have been overlooked in finance, current research shows that bot activity tends to be linked with conflict and polarization (e.g., Stella et al., 2018; Broniatowski et al., 2018). As a result, we hypothesize that bot activity will have a polarizing effect on investors translating into higher investor disagreement. We find that bot activity positively and significantly impacts both abnormal volume and volatility, and the effects persist for up to three months following the earnings.

We also contribute to the literature on myopic behaviors and earnings management. Past research shows that CEOs might use advertising to achieve myopic goals (Lou, 2014) by spotlighting short-term results. Bot activity might play a similar role by increasing the reach of social media posts and signaling higher popularity to investors. Moreover, bot activity is very similar to earnings management, specifically to real earnings management[4], as defined by Roychowdhury (2006). Managers can even alleviate negative news by using the official corporate social media account (Lee et al., 2015) and indicate positive reception. Furthermore, unlike other types of earnings management, this approach is inexpensive and very difficult to detect. We hypothesize that bot activity is positively related to myopic behavior and earnings management and that it is positively and significantly associated with cutting discretionary expenses, which is frequently used in earnings management to indicate positive performance (Graham et al., 2005; Bhojraj et al., 2009; Acharya and Xu, 2017). Therefore, a firm's management can manipulate the expectation and reception of earnings reports through bot activity to achieve myopic goals.

This paper is organized as follows. Section 2 outlines the bot industry, the history of bot activity on 𝕏, and the impacts of bot activity. Section 3 covers data collection, sample creation, and summary statistics. We report our analysis results and robustness checks in Section 4. Lastly, Section 5 concludes the paper.

2. The Bot Industry

In this section, we discuss the scale of bot activity on 𝕏 (formerly Twitter). 𝕏 research has been primarily focused on establishing a link between public sentiment on 𝕏 and stock prices. Bollen et al. (2011) developed six mood dimensions using 𝕏 messages and showed they are significant predictors of DJIA. Similarly, Gu and Kurov (2020) use the company 𝕏 sentiment from Bloomberg and show that it predicts stock returns without subsequent reversals. They further show that 𝕏 sentiment can be used to provide new information about the recommendations by analysts, as well as price targets and quarterly earnings. This is further supported by Bartov et al. (2018), who use a large volume of tweets mentioning a specific company to show that the sentiment is a significant predictor of quarterly earnings and price reaction. However, most financial research on social media does not consider the sheer amount of bot activity on 𝕏.

[4] While the more common way to manage earnings is through earnings accruals, the CEO can also manage the earnings using real activities manipulations (falsification), which may draw less scrutiny (Graham et al., 2005).

Officially, 𝕏 has reported in its 10-Q form that 5% of the activity is attributable to bots[5] (e.g., the 10-Q report from March 31, 2022). However, this figure has been heavily scrutinized. One of the most publicized criticisms came from Elon Musk before his acquisition of 𝕏, where he referenced various estimations provided by third parties, which measured the bot level activity at around 11% (Duffy and Fung, 2022). Furthermore, the extent of bot activity on 𝕏 has been questioned in the past. Timber and Dwoskin (2018) report that 𝕏 suspended more than 1 million accounts per day in the third quarter of 2018 to escalate the battle against disinformation and bot activity. This finding was preceded by a press release[6] by 𝕏, in which it vowed to provide a better platform for healthy civic discourse. However, this was not the first of assurances from 𝕏, with the company pledging to fight spam in a 2012 press release.[7] For example, the co-founder of Twitter, Jack Dorsey, stated in 2018: "We aren't proud of how people have taken advantage of our service or our inability to address it fast enough."[8]

The problem of bot accounts on social media is well-established in the political literature (e.g., Stella et al., 2018). Most research focuses on election manipulation (e.g., Bessi and Ferrara, 2016; Ferrara, 2017) or discourse amplification (e.g., Broniatowski et al., 2018). However, the market for bot accounts is not limited to the political sphere. In an investigation into the multimillion-dollar company Devumi, the *New York Times* discovered that the company had provided customers with over 200 million 𝕏 bot followers (Confessore et al., 2018). Their clients include television stars, professional athletes, politicians, and reporters. As a result, significant research has also been done on creating tools to detect bot accounts. These tools mainly rely on language processing and machine learning (e.g., Cresci et al., 2015; Sayyadiharikandeh et al., 2020).

Companies are aware of the importance of social media presence and reach, which is why individuals with 100,000 followers might earn, on average, $2,000 for a single promotional tweet. This can drive companies to build up their social media presence, frequently leading to their marketing and public

---

[5] 𝕏 does not report the total number of users on its accounts. Instead, it reports the total number of monthly active users during its quarterly earnings. Moreover, 𝕏 does not report the number of suspended accounts. As a result, it is difficult to test and verify its estimates.

[6] Available at https://blog.twitter.com/official/en_us/topics/company/2018/2016-election-update.html, last accessed January 18, 2023.

[7] Available at https://blog.twitter.com/official/en_us/a/2012/shutting-down-spammers.html, last accessed January 18, 2023.

[8] https://www.theguardian.com/technology/2018/mar/01/twitter-jack-dorsey-pledge-harassment-fake-news, last accessed January 18, 2023.

relations agencies purchasing bots to meet their goals faster (Confessore et al., 2018). This decision has a limited downside. 𝕏 does not typically suspend users suspected of buying bots (Confessore et al., 2018). This is because while 𝕏 can observe and block possible bot activity, it does not know who is responsible for any purchase. Moreover, it is essential to note that 𝕏 directly benefits from bot accounts since the individuals controlling them need to purchase the commercial API that allows automation.

Due to COVID-19, many companies had to move most of their business online, which also has moved customer interactions online. This change opens the company to the risk of bots attacking the company's reputation[9] via unfavorable reviews or posting negative messages on social media.[10]

𝕏 has acknowledged the problems of bot accounts on several occasions (e.g., Timber and Dwoskin, 2018) and has developed several initiatives to increase transparency and offer tools to encourage research. Notably, it created a website focusing on Platform Manipulation, where it documented the number of spam reports and bot challenges.[11] However, to the best of our knowledge, these initiatives stopped during the rumors of acquisition by Elon Musk, and they have not been updated since the end of 2021. Moreover, in February 2023, 𝕏 announced the termination and removal of academic API access.[12] Before 2023, academic scholars had access to free API, which has been used in more than 17,500 academic papers since 2020.[13] Instead of free API, the company presented new API packages in March 2023, with the cheapest option giving researchers access to 50 million tweets for $42,000 monthly. In comparison, the previous academic API package gave access to 10 million tweets monthly for no charge. Under the new system, the most expensive option offers academics 200 million tweets at $210,000 a month; however, this only accounts for roughly 0.3% of 𝕏's monthly activity. As a result, future research into 𝕏, the bot activity, and its impact is limited.

Another serious concern is the advancements in artificial intelligence (AI). One of the main tools available to combat bot activity is a challenge-response test known as CAPTCHA, which requires the user to enter a sequence of letters or numbers from a distorted image. 𝕏 frequently uses this tool and issues

---

[9] https://www.forbes.com/sites/googlecloud/2021/04/01/bot-attacks-are-the-biggest-online-risk-you-havent-addressed/?sh=37275fe6dda0.

[10] https://www.nytimes.com/2017/04/14/business/united-airlines-passenger-doctor.html.

[11] Available at https://transparency.twitter.com/en/reports/platform-manipulation.html#2021-jul-dec.

[12] Available at https://twitter.com/XDevelopers/status/1621026986784337922.

[13] Available at https://www.wired.com/story/twitter-data-api-prices-out-nearly-everyone/.

it to accounts suspected of bot activity.[14] However, contemporary research shows that traditional and widely used CAPTCHAs have been insecure (e.g., Dinh et al., 2023; Wang et al., 2023). Specifically, new deep learning AI appears capable of completing some of these tests, which, given their availability, raises further concerns about the ability of 𝕏 and other social media platforms to combat bot accounts effectively. As a result, the impacts of bots on social media and public discussions might increase in severity and call into question any methods of obtaining public sentiment, which will hamper future research.

3. Data and Methodology

3.1 Data collection process

Our data comes from the official corporate 𝕏 accounts of the S&P 500 companies. Specifically, we rely on the API offered by 𝕏 to scrape all original tweets.[15] In April 2013, the SEC allowed US-listed firms to make public disclosures through social media (SEC, 2013). We, therefore, use April 1, 2013, as the beginning of the sample and conclude the data collection on December 31, 2022. The resulting sample contains 23,451 unique tweets from 294 S&P 500 constituents with official 𝕏 accounts. Figure 1 shows that the corporate 𝕏 accounts experienced activity spikes around the quarterly earnings announcements. Following Yang et al. (2020), we focus on a period of 15 days surrounding the quarterly earnings announcement date.

[Figure 1]

We proceed with measuring 𝕏 bot activity around earnings announcements. The bot identification procedure involves the following steps. First, we collect data on all 𝕏 users who liked or reposted[16] the original S&P 500 companies' tweets. Then, we check whether 𝕏 has suspended these users. If a user's account is suspended, we assume it is a bot. We acknowledge that 𝕏 suspends accounts for reasons

[14] Available at https://help.twitter.com/en/managing-your-account/suspended-twitter-accounts.

[15] We restrict the sample to tweets with at least 100 reposts or 100 favorites. The breakpoint of 100 was chosen because it is the minimum amount one can buy from a typical website offering bot services (e.g., https://venium.com/twitter/buy-twitter-retweets/). If a company has several official accounts, e.g., @Amazon and @AmazonNews, we use the main official account that the company displays on its website.

[16] Prior to the name change from Twitter to X, likes were called favorites and reposts were called retweets.

other than the malicious use of automation. These reasons, however, primarily include engaging in illegal activities, such as calls to violence, terrorism, and child sexual exploitation. Thus, we believe that the accounts suspended for engaging in these illicit activities are unlikely to interact with S&P 500 companies' earnings disclosure tweets. Moreover, 𝕏 account suspension typically results from failing a bot challenge. In fact, the 𝕏 guidelines state, "Most of the accounts we suspend are suspended because they are spammy or just plain fake."[17] We also recognize that 𝕏 may fail to suspend some of the bot accounts. Accordingly, our identification procedure results in a conservative estimate of bot activity. We provide more detailed information and a description of the data collection process in Appendix A1.

The data reveals that the sample companies are subject to significant bot activity. In total, their official 𝕏 accounts have been liked 37,531,072 times. Of that, 7,641,950 (20.3%) "likes" were done by bot accounts.

3.2 Sample construction

Earning people's trust is the core value of 𝕏.[18] This trust, however, may be undermined if a large proportion of 𝕏 activity comes from inauthentic users. To capture this proportion, we calculate the *BotRatio.*

$$BotRatio_{i,t} = \frac{BotLikes_{i,t}}{TotalLikes_{i,t}},$$

where $BotLikes_{i,t}$ is the number of likes by suspended accounts on company *i* tweets that were tweeted on day *t*. Similarly, $TotalLikes_{i,t}$ are the total number of likes on the official tweets of company *i* that were tweeted on day *t*.

Aggregation on a daily level allows us to reduce the potential impact of outliers. For the tweets that occurred on weekends or holidays, we set the day as the nearest following business day. Similarly, for tweets that occurred outside of business hours, we set the day as the nearest following business day.

[17] For completeness, the account might be suspended for other reasons than being a bot (e.g., violence, terrorism, and child sexual exploitation). We argue, however, that accounts of individuals reacting to official corporate accounts will not be associated frequently with these other reasons for suspension.

[18]https://workat.tech/company/twitter#:~:text=Twitter's%20mission%20is%20to%20give,and%20information%20instantly%20without%20barriers.

We define the following dependent variables to analyze the investor disagreement and price effects of bot activity. First, *volatility* is the standard deviation of market returns since this measure has been used to measure investor disagreement (Antweiler and Frank, 2004). Secondly, we use the log of abnormal volume (*AbLogVol)*, a measure defined by Cookson and Niessner (2020), to quantify investor disagreement by analyzing the trade volume. *AbLogVol* is the difference between the share volume on day *t* and the average share volume between days *t*-120 and *t*-20. Lastly, we use the analyst dispersion, which has been shown to indicate both analyst and investor disagreement (Diether et al., 2002). Specifically, we use *absolute dispersion*, which is defined as the absolute value of the standard deviation of analysts' forecasts scaled by the consensus for the quarterly earnings per share issued just before the quarter-end date. We use the absolute value of the dispersion since we do not distinguish between positive and negative earnings.

We follow Bartov et al. (2018) and Cookson and Niessner (2020) for the choice of control variables. Specifically, we use *size*, defined as the market value of equity; *book-to-market ratio*; *illiquidity,* defined by Amihud (2002); past returns described by Brennan et al. (2012); the number of analysts covering the stock; the indicator for the fourth quarter; and analyst dispersion. Detailed variable definitions of these control variables are available in Appendix A2.

It is essential to note that the global level of bot activity varies over time. As the 𝕏 transparency center outlines, there are periods of increased efforts in detecting bots. Moreover, continuous technological improvements can make bot detection either easier or harder. Furthermore, many bots are used to spread misinformation during political events (e.g., Bessi and Ferrara, 2016). This might increase scrutiny during those periods and temporarily affect *BotRatio*. Therefore, our analysis requires additional control variables to account for the global bot activity level. In particular, for each firm *i* and day *t*, we define the global level of bot activity as a mean bot ratio for all firms apart from the given firm during the 3- and 7-day window surrounding day *t*. We can then define 3-day abnormal bot activity (7-day abnormal bot activity), a dummy variable equal to one if the *BotRatio* is greater than the average 3-day (7-day) global bot activity.

3.3 Summary statistics and determinants of bot activity

Table 1 presents the summary statistics. We observe considerable variation in the number of tweets by corporations and the number of likes on the given tweets. Some firms connect with their customers more frequently (e.g., Netflix, Amazon, or Apple), while others do not even have an official 𝕏 account (e.g., Berkshire Hathaway). The mean number of tweets for any given day is over 2; however, the distribution is skewed to the right. Moreover, there is a considerable variation in the number of likes. We only focus on tweets with more than 100 likes since most websites offering bot likes require purchasing at least 100 likes in an order. The median number of likes in a given day is 419; however, the mean is 3,646. There are several outliers with more than 100,000 likes (generally for new product announcements, such as new show announcements on Netflix).

[Table 1]

The average *BotRatio* is 17%. This figure is more prominent than the 5% level reported by 𝕏's 10-Q report and is greater than the estimate of 11% suggested by Elon Musk (Duffy and Fung, 2022). This number indicates that the 𝕏 accounts of S&P 500 companies experience a relatively higher level of attention by bots than other users' accounts. To further verify our methodology, we calculate bot activity using reposts of the original tweets instead of likes. The ratio of suspended accounts is almost identical to those based on tweet likes.[19] Overall, this suggests that the main drivers of account suspension are similar among the users who liked or reposted a post. This gives further credence to the idea that these users are mostly bots. We can also see that roughly one-fourth of the sample had bot activity higher than the global level of bot activity. This result suggests that firms regularly purchase some base-level bot activity for all their tweets and increase it for specific tweets.

4. Results

4.1 Bot activity and investor disagreement

We begin by testing how the bot activity impacts investor disagreement. This analysis requires addressing the potential endogeneity of the bot activity on the firm's 𝕏 account. We rely on the 2SLS approach, where we instrument the *BotRatio* by the following IV regression:

$$BotRatio_{i,t} = \alpha_i + \beta_i 3\ day\ global\ bot\ activity_{i,t} + \gamma_i 7\ day\ global\ bot\ activity_{i,t} +$$

$$quarter_t + yearFE + IndustryFE + \varepsilon_{i,t} \qquad (1)$$

[19] The results are available on request.

where $BotRatio_{i,t}$ is the number of suspended accounts that liked any tweet of the corporation on the given day divided by the total number of accounts that liked a tweet of the corporation on the given day; *3-day global bot activity* is the average bot ratio for all other corporate tweets, excluding the company *i*, that occurred during the 3-day window; and *7-day global bot activity* is calculated accordingly. This approach is used to reflect valid instruments of bot activities on the other firm's 𝕏 accounts with additional variables controlling for the industry and time fixed effects. Appendix A3 reports the resulting estimates of the first-stage regression.

The second stage regression specification is as follows.

$$DepVar_{i,t} = \alpha_i + \beta_1 \widehat{BotRatio}_{i,t} + Controls_{i,t} + \varepsilon_{i,t}, \qquad (2)$$

where $DepVar_{i,t}$ is one of the dependent variables measuring investor disagreement and $\widehat{BotRatio}_{i,t}$ is the predicted value of 𝕏 bot activity from the first-stage regression. $Controls_{i,t}$ include size, book-to-market value ratio, fourth-quarter fixed effect, illiquidity, analyst dispersion, number of analysts, and the company's past returns. Table A2 provides definitions of the control variables. We also include year and industry fixed effects in the regressions.

Following the literature on investor disagreement, we test the impact of bot activity on return and volume volatility (e.g., Hirschleifer, 1977). We begin by analyzing the impact of bot activity on return volatility, which we define as the standard deviation of daily returns. Table 2 shows that the bot ratio is a positive and significant predictor of volatility. In particular, estimated elasticities in mean[20] shows that an increase in bot activity by 1% typically leads to a 0.98% increase in the standard deviation of returns. This means the posts for which bot accounts were purchased to magnify the apparent popularity and social reach will lead to higher disagreement.

[Table 2]

The above result holds for both periods - before quarterly earnings as well as in the post-earnings period. This finding suggests that bots can be used even after the quarterly earnings announcements are made public to limit negative responses, as Lee et al. (2015) outlined. Overall, these results are

[20] In computation we used the Delta method where the elasticity in means is estimated as $\beta \frac{\bar{y}}{\bar{x}}$.

consistent with past research on social media bots (Ferrara et al., 2016) that shows that they increase polarization and disagreement. Therefore, the use of bots could negatively impact market stability.

However, the previous literature (e.g., Cookson and Niessner, 2020) argues that volatility is an indirect measure of disagreement that relies on observed trading patterns. Instead, instead recommend using an abnormal volume variable (*AbLogVol*), which should directly capture the dispersion of investor opinions. To address this concern, Table 3 complements the findings on return volatility with the analysis of the impact of bot activity on abnormal trading volume. Similar to return volatility, we can see that the bot activity significantly increases the volume during the quarterly earnings period. Specifically, an increase in bot activity by 1% typically increases the abnormal volume by 2.2%.

[Table 3]

Overall, these results suggest that bot activity significantly impacts both the return volatility and the abnormal volume. This result is consistent with the hypothesis that the primary motivation for bot activity is to indicate positive expectations or the positive reception of earnings. Furthermore, the results contradict the notion that bots are used to increase the post's social reach and increase the investor's informativeness. In that case, we should have observed lower volatility and volume since higher informativeness decreases disagreement (e.g., Milgrom and Stokey, 1982; Karpoff, 1986). Therefore, our results indicate that bots are used to help improve the reception of the QE announcements or their expectations at the expense of more considerable volatility and investor disagreement. This strategy is similar to real earnings management (Roychowdhury, 2006). However, to truly understand the impact of bots, we need to better understand the motivations for bot activity.

4.2 Rationale behind bot activity

In this section, we explore the motivation behind buying 𝕏 bots. One possible rationale is that the management might buy bots to increase attention and indicate a positive reaction to the company's earnings. For example, large firms enjoy more extensive analyst coverage than smaller firms (O'Brien and Bhushan, 1990). As a result, investors might not be as motivated to search for additional sources of information for larger firms. Conversely, smaller firms might benefit more from bot activity, which may help indicate agreement among investors or the positive reception of their earnings. We show the relationship between firm size and bot activity in Figure 2. The results show that firms with the smallest market capitalization have much more significant levels of bot activity, which is consistent with our intuition.

[Figure 2]

The strategy of buying bots to increase attention to the company's earnings announcements can be especially effective when the earnings announcements are not highly informative. For example, an earnings announcement that aligns with the expectations carries little information. This situation may incentivize the management to purchase 𝕏 bots to emphasize the little information that exists.

We, therefore, proceed with exploring the relationship between bot ratio and earnings surprise. Figure 3 shows that, indeed, the bot ratio is the largest when the QE results are in line with analyst expectations. This result is consistent with the suggestion that bots are used to increase attention when earnings announcements carry little additional information.

[Figure 3]

When the announced earnings align with the analysts' assessments, investors might be looking for other sources of information to gain an informational advantage, and as a result, purchasing bot activity can indicate positive earnings reception.[21] However, Bhojraj et al. (2009) outline that a company in line with analyst expectations might indicate short-termism. This issue arises when managers are motivated to meet earnings expectations at the expense of future economic growth. Therefore, if a manager decides to pursue myopic goals, bot activity might be a novel tool to spotlight the company's current short-term economic outlook.

While the users cannot observe the bot ratio, they can watch the user's reaction to social media posts. Since bots inflate the number of likes (by roughly 17% on average in our sample), it may lead users to believe that the company has broader support and further increase investor attention. This result is supported by the findings of Lou (2014). He finds that managers adjust firm advertising to attract attention and influence short-term returns further. In this way, bot activity might serve as another type of advertising.[22]

[21] In this sense, bot activity can be taken as a potential real earnings management outlined by Roychowdhury (2006). While the more common way to manage earnings is through earnings accruals, the CEO can also manage the earnings using real activities manipulations (falsification), which may draw less scrutiny (Graham et al., 2005).

[22] Let us again note that compared to other methods of advertising, the purchase of bots is relatively cheap, with many websites offering 100 likes for around 3$.

These results indicate that managers are likely responsible for purchasing 𝕏 bots to bring attention to a company's current financial results, even at the expense of future performance. To further test this hypothesis, we focus on an indicator of short-termism. Graham et al. (2005) report that over 80% of CEOs polled suggested that they would decrease discretionary expenses in order to meet the current earnings target. This result is further supported by Chen et al. (2015), who find that when the monitoring of CEOs decreases, CEOs' compensation increases. They focus on short-term returns, even at the cost of value-destroying acquisitions and actions such as reducing discretionary expenses.

To calculate the abnormal discretionary expenses (XSGA), we follow the model by Roychowdhury (2006) and estimate the following cross-sectional regression for each industry and quarter:

$$\frac{XSGA_{it}}{Assets_{i,t-4}} = \beta_1 * \frac{1}{Assets_{i,t-4}} + \beta_2 * \frac{Sales_{i,t-4}}{Assets_{i,t-4}} + \epsilon_{it} , \tag{3}$$

where $XSGA_{it}$ is the SG&A costs, including marketing and R&D for firm *i* in quarter[23] *t*. The residuals then account for the levels of abnormal discretionary expenses. Negative values should, therefore, imply reductions in discretionary expenses, indicate short-termism, and lead to an increase in the current period earnings. We analyze the relationship between the bot ratio and abnormal discretionary expenses in Table 4.

[Table 4]

Table 4 reveals that the bot ratio has a significant and negative relationship with discretionary spending. Specifically, an increase of bot activity by 1% is associated with a mean cut in discretionary expenses by 2.7%. High mean elasticity suggests that high bot activity is related to reductions in spending, most likely in an attempt to increase earnings for the period, as outlined by Bhojraj et al. (2009) and Acharya and Xu (2017). The bot activity can then serve as a tool to attract attention and influence short-term returns further (Lou, 2014). These results indicate that the management or insiders may be the ones who

[23] Roychowdhury (2006) analyzes annual statements in his paper, since CEOs have higher incentive to meet annual earnings expectations. However, the measure can also be constructed at quarterly level, given that the assets and sales are lagged by one year, to account for the seasonality of earnings (Roychowdhury, 2006).

purchased the bots. We further split the sample based on earnings surprise in columns (3) and (4).[24] We find that the relationship is the strongest for the positive earnings surprise, where a 1% increase in bot activity is associated with a 5% decrease in discretionary spending. This result is intuitive since higher cuts in discretionary spending will lead to higher earnings, which may be related to higher advertising through bots. Moreover, managers might be further motivated to beat analyst expectations to increase equity issuances and insider selling. This decision would encourage them to use bots to advertise the financial results.

Our results indicate that managers who are focused on myopic goals can use bot accounts to achieve such goals. By employing bots on the official 𝕏 accounts of corporations, they inflate the number of likes and the apparent popularity of the company, which can lead investors to believe that the company is more popular and, as a result, more valuable. The results further suggest that bot activity can be used to increase the effects of positive earnings, which the managers achieve by cutting discretionary expenses. However, given the data limitations, there is no formal test to determine who purchased or is responsible for the bot activity.

4.3 Long-term impact of bot activity

To better disentangle the effects of bots on firm financials, this section analyzes the long-term impact of bot activity. First, we focus on its effect on volatility and consider various time windows from one week after earnings until three months after earnings.[25] We report the results of the analysis of bot activity on the volatility of returns in Table 5.

[Table 5]

We use daily data to analyze the long-term impact consistently with previous models. We observe that bot activity has a long-lasting effect for up to 3 months following the earnings. While the coefficients suggest that bot activity around the QE are gradually losing their impact and significance, the mean elasticities show that the effect remains steady over time. Specifically, the increase in bot activity by 1% is associated with an increase in return volatility by 0.57% in the first week after the QE, 0.76% in the

[24] In our partitioning, we are grouping together earnings that are in line with analyst expectations and negative earnings surprises. Negative earnings surprises constitute a very small sample and their exclusion or separate evaluation does not change our results and conclusions.

[25] While analysis of even longer time periods is possible, it would then cross another earnings period, which might bias the results.

two weeks after it, and 0.80% in the three weeks after it. In the horizon of two to three months, the mean effect falls to 0.53% and 0.40%, respectively, which indicates persistent and highly economically significant results. Moreover, there is a difference in the longevity of the impact depending on whether the bot activity occurred pre- or post-earnings announcement. While the bot activity before earnings remains significant at 1%, the bot activity post-earnings loses significance after one month and is substantially smaller.

Cookson and Niessner (2020) suggest that investor disagreement can be decomposed into two parts: 1) disagreement caused by different information sets and 2) disagreement about the interpretation of financial information. Therefore, the disagreement caused by bot activity before earnings, which might be linked to disagreement about the information sets, can have long-lasting effects and affect market stability. This result further supports our hypothesis that bot activity motivation might differ based on whether it is pre- or post-earnings. Since we observe that bot activity rises with the decrease in firm size, the observed bot activity can be the tip of the iceberg and be much higher comparatively for smaller firms.

However, the results so far could be attributed to retail investors, who might be more influenced by social media and the artificial increase in popularity caused by bot activity. Diether et al. (2002) suggest that analyst dispersion can be a proxy for investor disagreement. Since analysts could be considered more professional investors than retail investors, disagreement among them allows us to analyze whether bot activity can affect more sophisticated investors. Furthermore, analysts could incorporate social media popularity in their predictions, as Gerken and Painter (2023) outline. We explore how bot activity impacts future analyst dispersion following earnings in Figure 4.

[Figure 4]

In Figure 4 we analyze the analyst dispersion as analysts approach earnings announcement day, taking into account whether there was large or low bot activity in the previous quarter. We observe that while dispersion increases as the announcement day approaches, the dispersion is much larger for the firm's quarterly earnings if analysts observed high bot activity in the previous quarter. This result suggests that bot activity might have a polarizing effect, consistent with political literature (e.g., Broniatowski et al., 2018), and amplify the discourse among analysts. We further test the relationship between future analyst dispersion and bot activity in Table 6.

[Table 6]

We can again see that the bot ratio significantly affects future analyst dispersion. Namely, mean estimated elasticities show that a 1% increase in bot ratio leads to a 0.58% increase in analyst dispersion. This result taken together with the previous figure also indicates that analysts likely consider social media popularity as well as the reaction of investors on social media to previous earnings. Since the bot activity is not observable, only the number of likes or reposts, it can lead to an overestimation of the popularity and biased estimates. To further test how analysts are affected by the bot ratio, we again split the sample based on the earnings expectation, as we did in Section 4.2. Intuitively, earnings that are in line with analyst expectations carry little additional information. Therefore, investors might be motivated to search for other sources of information. In Table 6, we can see that when earnings are in line with expectations, the bot activity has a much higher effect on the analyst dispersion for the following earnings announcement.

4.4 Causality interpretation

While the results in the previous section indicate a correlation between bot activity and the subsequent increase in disagreement, they do not necessarily provide causality interpretation (e.g., Heckman, 2010; LaLonde, 1986; Blundell and Costa-Dias, 2000; DiNardo and Lee, 2010). To address this concern, we will use the Rubin Causal Model (Holland, 1986). The model is based on two possible outcomes: one with treatment and one without treatment.

$$y_{0i} = \mu_0 + \varepsilon_{0i} \text{ and } y_{1i} = \mu_1 + \varepsilon_{1i}.$$

Formally, it can be written as $y_{Ti} = \mu_T + \varepsilon_{Ti}$, where subscript T=1 denotes the treatment and T=0 represents the control group. We will use established randomized control trial (RCT) techniques to estimate the Average Treatment Effect on the Treated (ATET),[26] where y_{0i} is estimated using the nearest-neighbor approach with an extensive set of controls. In our models, we exactly match the given quarter and year and use the firm-level controls (size) as the coordinates for the approximate matching. We define the treatment group (T=1) as the companies with bot activity above the 75 percentile of the

[26] The treatment effect $E[y_{1i} - y_{0i}]$ is under random assignment equal to $\mu_1 - \mu_0$, motivating our choice of a randomized control trial (RCT).

global level. In comparison, the control group is represented by those with bot activity below the 45 percentile.[27] Table 7 reports the results of this analysis.

[Table 7]

To facilitate better matching, we aggregate the bot activity data by calculating the maximum[28] ratio for both the pre- and post-earning period. We can see that our results confirm our previous regression model results. The treatment effect is statistically and economically significant, and it shows that bot activity substantially affects the volatility and abnormal volume. Matching quality indicators are available in Table 7, and detailed balancing tests are available upon request. Overall, the results suggest that bot activity has a causal effect on higher volatility and volume, which is also economically significant. The ATET computes changes in the corresponding variable in the treated group compared to the similar firms in the control sample. Table 7 also presents the means of each variable for the entire sample and the treated group to evaluate the economic significance. For example, the effect of high versus low bot activity (before the QE, see columns (1)-(3)) is between 17 and almost 74% of the mean of the treated variables. When compared with the total average, we observe a higher impact. This means that bot activity has a strong causal impact on investor disagreement, in line with both previous theories on bot effects on social media (e.g., Bessi and Ferrara, 2016; Ferrara et al., 2016) as well as the social media disagreement literature (Antweiler and Frank, 2004; Cookson and Niessner, 2020). Consequently, it is crucial to understand the further role social media bots could have on financial markets.

5. Conclusion

Bot activity is an affordable and quick way for companies to attain a broad social reach, providing them with a virtual megaphone to reach investors. What consequences does 𝕏 bot activity entail? The analysis of bot activity in finance research is limited, as past research primarily focuses on issues relating

[27] Let us note that the 45-percentile level for the control group was chosen to guarantee more than one observation for the match sample, since we would not be able to use the robust option otherwise.

[28] We choose the maximum instead of the mean, since not all tweets are identical, and some tweets might not have high bot activity as a choice. Moreover, using the mean of AbLogVol is not recommended given the way the variable is constructed. Nevertheless, using mean does lower the estimates, but the conclusions do not change. These results are available upon request.

to election interference and the spread of misinformation. This gap in the literature exists despite corporations being aware of the positive impact that social media presence can have.

In this paper, we analyze the impact of bot activity during quarterly earnings announcements on the investors' reactions. Bot activity artificially increases the number of likes and reposts a company gets on its 𝕏 posts. This feature allows it to reach users who do not follow the company by making it recommended to other users. However, this may come at a cost because bot activity may lead to disagreement.

We document that bot activity leads to a significant increase in investor disagreement. This results in higher volatility and trading volume. Furthermore, the increase appears to be long-term, as we observe that the effect persists for up to three months following the earnings. Moreover, we also detect that analysts could be affected by the bot activity since bot activity has a significant effect on the analyst dispersion for the following quarter. Our results are robust to various specifications, including a matching ATET approach indicating causal interpretation.

Social media users receive recommended accounts and tweets based on their interests and the perceived popularity of the given accounts and tweets. As a result, purchasing bot accounts increases the number of times the corporation's tweets are recommended to users. This may facilitate managers in pursuing myopic goals. In line with this, we analyze the relationship between bot accounts and short-termism by analyzing cutting discretionary expenses. We find that bot activity is high during quarters with cuts in discretionary expenses. The cuts to discretionary expenses are frequently motivated by the desire to increase earnings artificially to meet the short-term myopic goals of CEOs. Bot activity can be an essential tool to achieve this goal since it is inexpensive and difficult to detect.

In total, the impact of bots is substantial. While our research is limited to quarterly earnings due to data limitations, both corporations and individuals might employ them over more extended periods due to their inexpensiveness. Moreover, our sample is limited to constituents of S&P 500 companies. As a result, the effect can be more pronounced for smaller companies, which are less covered by analysts.

Following the acquisition of 𝕏 by Elon Musk in 2022, many tools available for studying or analyzing bot activity are no longer available. For instance, the academic API for 𝕏 data has been discontinued, and consequently, academics have been priced out of 𝕏 and bot research. Finally, the main tools for bot detection rest on bot challenges, such as CAPTCHA. Given the recent developments in AI technology, many of these challenges appear unreliable. As a result, the impact of bots on social media discussions can be expected to increase. Similarly, the severity of the effects might increase as well. Overall, our results and the worrying trend in bot activity underline the importance of future research and the discussion on addressing the issue.

References

Acharya, V., & Xu, Z. (2017). Financial dependence and innovation: The case of public versus private firms. *Journal of Financial Economics*, *124*(2), 223-243.

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, 5(1), 31-56.

Antweiler, W., & Frank, M. Z. (2004). Is all that talk just noise? The information content of internet stock message boards. *The Journal of Finance*, *59*(3), 1259-1294.

Bartov, E., Faurel, L., & Mohanram, P. S. (2018). Can Twitter help predict firm-level earnings and stock returns?. *The Accounting Review*, *93*(3), 25-57.

Bessi, A., & Ferrara, E. (2016). Social bots distort the 2016 US Presidential election online discussion. *First Monday*, *21*(11-7).

Bhojraj, S., Hribar, P., Picconi, M., & McInnis, J. (2009). Making sense of cents: An examination of firms that marginally miss or beat analyst forecasts. *The Journal of Finance*, *64*(5), 2361-2388.

Blankespoor, E., Miller, G. S., & White, H. D. (2014). The role of dissemination in market liquidity: Evidence from firms' use of Twitter™. *The Accounting Review*, 89(1), 79-112.

Blundell, R., & Costa Dias, M. (2000). Evaluation methods for non-experimental data. *Fiscal Studies*, 21(4), 427-468.

Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, *2*(1), 1-8.

Bradley, D., Hanousek Jr, J., Jame, R., & Xiao, Z. (2022). Place Your Bets? The Market Consequences of Investment Research on Reddit's Wallstreetbets. *SSRN.*

Brennan, M.J., Chordia, T., Subrahmanyam, A., Tong, Q., (2012). Sell-order liquidity and the cross-section of expected stock returns. *Journal of Financial Economics*, 105, 523-541

Broniatowski, D. A., Jamison, A. M., Qi, S., AlKulaib, L., Chen, T., Benton, A., ... & Dredze, M. (2018). Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *American Journal of Public Health*, *108*(10), 1378-1384.

Confessore, N.,Dance, G. J.X., Harris, R., and Hansen, M. (2018). The Follower Factory. *The New York Times*. Available at https://www.nytimes.com/interactive/2018/01/27/technology/social-media-bots.html

Cookson, J. A., & Niessner, M. (2020). Why don't we agree? Evidence from a social network of investors. *The Journal of Finance*, *75*(1), 173-228.

Chen, H., De, P., Hu, Y. J., & Hwang, B. H. (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. *The Review of Financial Studies*, *27*(5), 1367-1403.

Chen, T., Harford, J., & Lin, C. (2015). Do analysts matter for governance? Evidence from natural experiments. *Journal of Financial Economics*, *115*(2), 383-410.

Cresci, S., Di Pietro, R., Petrocchi, M., Spognardi, A., & Tesconi, M. (2015). Fame for sale: Efficient detection of fake Twitter followers. *Decision Support Systems*, *80*, 56-71.

Diamond, D. W., & Verrecchia, R. E. (1981). Information aggregation in a noisy rational expectations economy. *Journal of Financial Economics*, 9(3), 221-235.

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

DiNardo, J., & Lee, D. S. (2011). Program evaluation and research designs. *In Handbook of Labor Economics* (Vol. 4, pp. 463-536). Elsevier.

Dinh, N. T., & Hoang, V. T. (2023). Recent advances of Captcha security analysis: a short literature review. *Procedia Computer Science*, *218*, 2550-2562.

Duffy, C., Fung, B. (2022). Elon Musk commissioned this bot analysis in his fight with Twitter. Now it shows what he could face if he takes over the platform, *CNN*, available at https://edition.cnn.com/2022/10/10/tech/elon-musk-twitter-bot-analysis-cyabra/index.html

Farrell, M., Green, T. C., Jame, R., & Markov, S. (2022). The democratization of investment research and the informativeness of retail investor trading. *Journal of Financial Economics*, *145*(2), 616-641.

Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The rise of social bots. *Communications of the ACM*, *59*(7), 96-104.

Ferrara, E. (2017). Disinformation and social bot operations in the run up to the 2017 French presidential election. *arXiv preprint arXiv:1707.00086*.

Gerken, W. C., & Painter, M. O. (2023). The value of differing points of view: Evidence from Financial Analysts' Geographic Diversity. *The Review of Financial Studies*, 36(2), 409-449.

Graham, J.R., Harvey, C.R., Rajgopal, S., (2005). The economic implications of corporate financial reporting. *Journal of Accounting and Economics* 40, 3–73.

Gu, C., Kurov, A., 2020. Informational role of social media: Evidence from twitter sentiment. *Journal of Banking & Finance* 121, 105969. doi:https://doi.org/10.1016/j.jbankfin.2020.105969

Heckman, J. J. (2010). Building bridges between structural and program evaluation approaches to evaluating policy. *Journal of Economic literature*, *48*(2), 356-398.

Hirshleifer, J. (1977). The theory of speculation under alternative regimes of markets. The *Journal of Finance*, 32(4), 975-999.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, 81(396), 945-960.

Jung, M. J., Naughton, J. P., Tahoun, A., & Wang, C. (2015). Corporate use of social media. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2588081

Karpoff, Jonathan M., (1986), A theory of trading volume, *Journal of Finance* 41, 1069–1087.

LaLonde, Robert J. (1986). “Evaluating the econometric evaluations of training programs with experimental data.” *The American Economic Review*, 604-620.

Lee, L. F., Hutton, A. P., & Shu, S. (2015). The role of social media in the capital market: Evidence from consumer product recalls. *Journal of Accounting Research*, *53*(2), 367-404.

Lou, D. (2014). Attracting investor attention through advertising. *The Review of Financial Studies*, *27*(6), 1797-1829.

Milgrom, P., & Stokey, N. (1982), Information, trade, and common knowledge, *Journal of Economic Theory* 26, 17–27.

O’Brien, P. C., & Bhushan, R. (1990). Analyst following and institutional ownership. *Journal of Accounting Research*, 28, 55-76.

Roychowdhury, S., (2006). Earnings management through real activities manipulation. *Journal of Accounting and Economics*, 42, 335–370.

Sayyadiharikandeh, M., Varol, O., Yang, K. C., Flammini, A., & Menczer, F. (2020). Detection of novel social bots by ensembles of specialized classifiers. In Proceedings of the 29th ACM international conference on information & knowledge management (pp. 2725-2732).

SEC, 2013, Report of Investigation Pursuant to Section 21(a) of the Securities Exchange Act of 1934: Netflix, Inc., and Reed Hastings, Press Release No. 69279 www.sec.gov/litigation/investreport/34-69279.pdf

Stella, M., Ferrara, E., & De Domenico, M. (2018). Bots increase exposure to negative and inflammatory content in online social systems. *Proceedings of the National Academy of Sciences*, *115*(49), 12435-12440.

Timber C., and Dwoskin E. (2018) Twitter is sweeping out fake accounts like never before, putting user growth at risk. *The Washington Post*. Available at https://www.washingtonpost.com/technology/2018/07/06/twitter-is-sweeping-out-fake-accounts-like-never-before-putting-user-growth-risk/

Wang, P., Gao, H., Xiao, C., Guo, X., Gao, Y., & Zi, Y. (2023). Extended research on the security of visual reasoning captcha. *IEEE Transactions on Dependable and Secure Computing*.

Yang, Y. C., Zhang, B., & Zhang, C. (2020). Is information risk priced? Evidence from abnormal diosyncratic volatility. *Journal of Financial Economics*, *135*(2), 528-554.

Figure 1. The number of company tweets (weekly) around the quarterly earnings.

The figure contains the total number of weekly company tweets around the quarterly earnings. The numbers in brackets specify the week by the interval in days. Days with a minus sign denote the days before the QE, and days with a plus sign correspond to the day after the QE. For example, (-7; 0) marks the week before the QE, (0; 7) represents the week after the earnings, etc.

The figure shows the marginal effects associated with each week around the earnings. We controlled for company time-invariant heterogeneity (i.e., company fixed effects included).

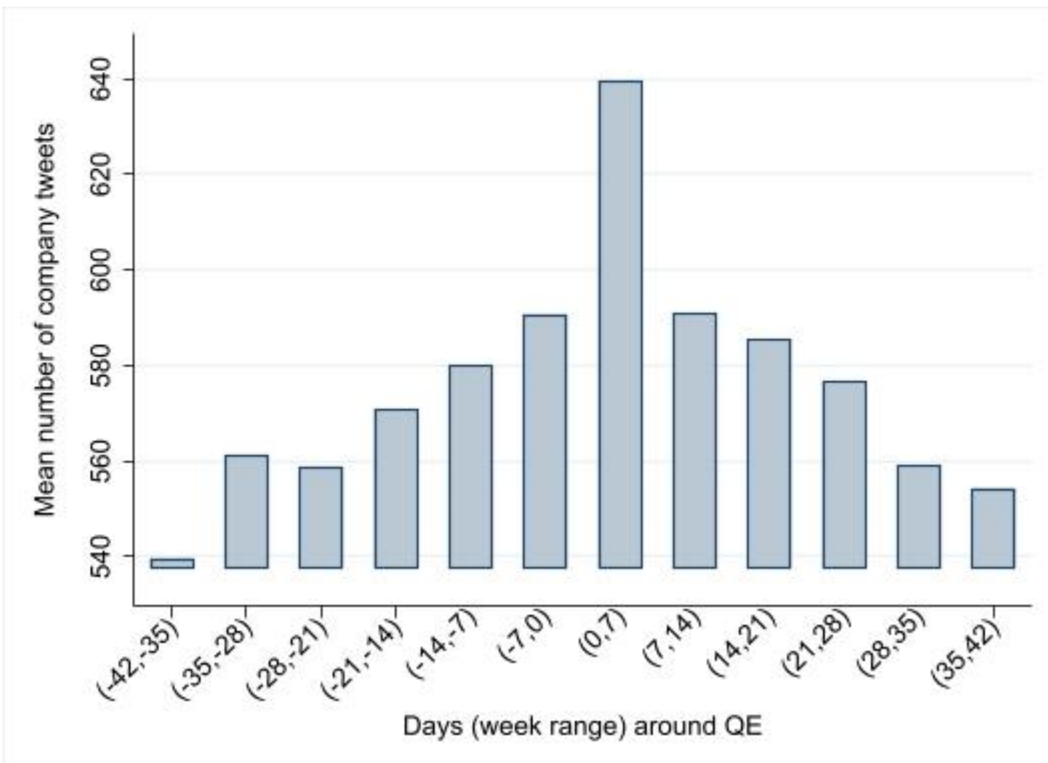

Figure 2. Bot activity around the QE by company size.

The figure shows the marginal effects of the bot activity associated with the company size. The first pentile (20th percentile) contains the smallest firms, the third pentile (60th percentile) represents the medium size firms, and the fifth pentile (i.e., 100th percentile) corresponds to the largest firms. The bot activity is measured as the ratio of suspended user accounts that put a like on the company tweets in one week before and one week after the QE.

The figure shows the mean share of suspended accounts (marginal effects) associated with each pentile of firm size. We controlled for the company time-invariant heterogeneity (i.e., company fixed effects included). Marginal effects were computed using the delta method.

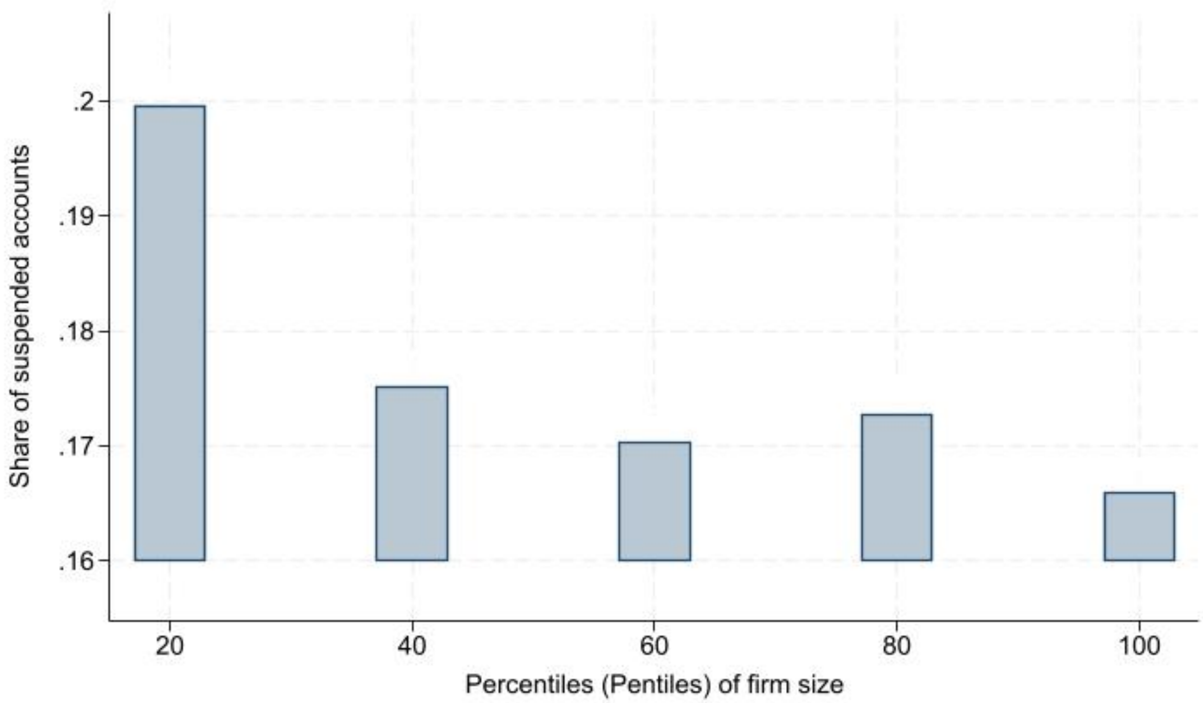

Figure 3. Bot activity around the QE by percentiles of earning surprise

The figure shows the marginal effects of the bot activity associated with the earning surprise. The earning surprise is measured as the absolute value of the analyst surprise (SUE Analyst). The first pentile (20th percentile) contains the firms with the slightest earning surprise, i.e., the firms align with the analyst expectation. The higher the percentile of the earning surprise, the higher the distance from the analyst's expectation. The bot activity is measured as the ratio of suspended user accounts that put a like on the company tweets in one week before and one week after the QE.

The figure shows the mean share of suspended accounts (marginal effects) associated with each pentile of firm size. We controlled for the company time-invariant heterogeneity (i.e., company fixed effects included). Marginal effects were computed using the delta method.

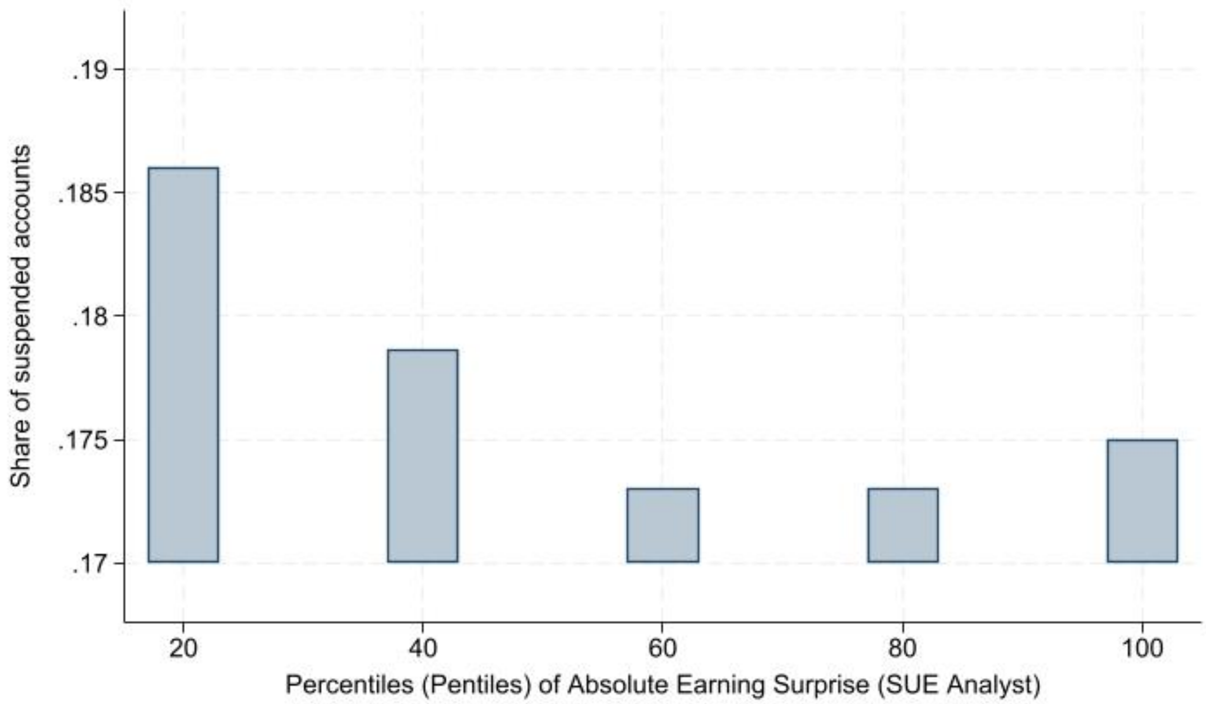

Figure 4. Bot activity effect on analyst dispersion

The figure shows the dispersion of analyst forecasts leading up to the announcement date. Analyst dispersion is calculated as the standard deviation of analyst forecasts issued during the given period for the given stock divided by the average stock price during the earnings period. The analyst dispersion is reported in percentages. We define the previous quarter as earnings with high bot activity if the bot activity was above the 75 percentile of the global level. Conversely, low bot activity is defined as low if the bot activity was below the 25 percentile of the global level in the previous quarter.

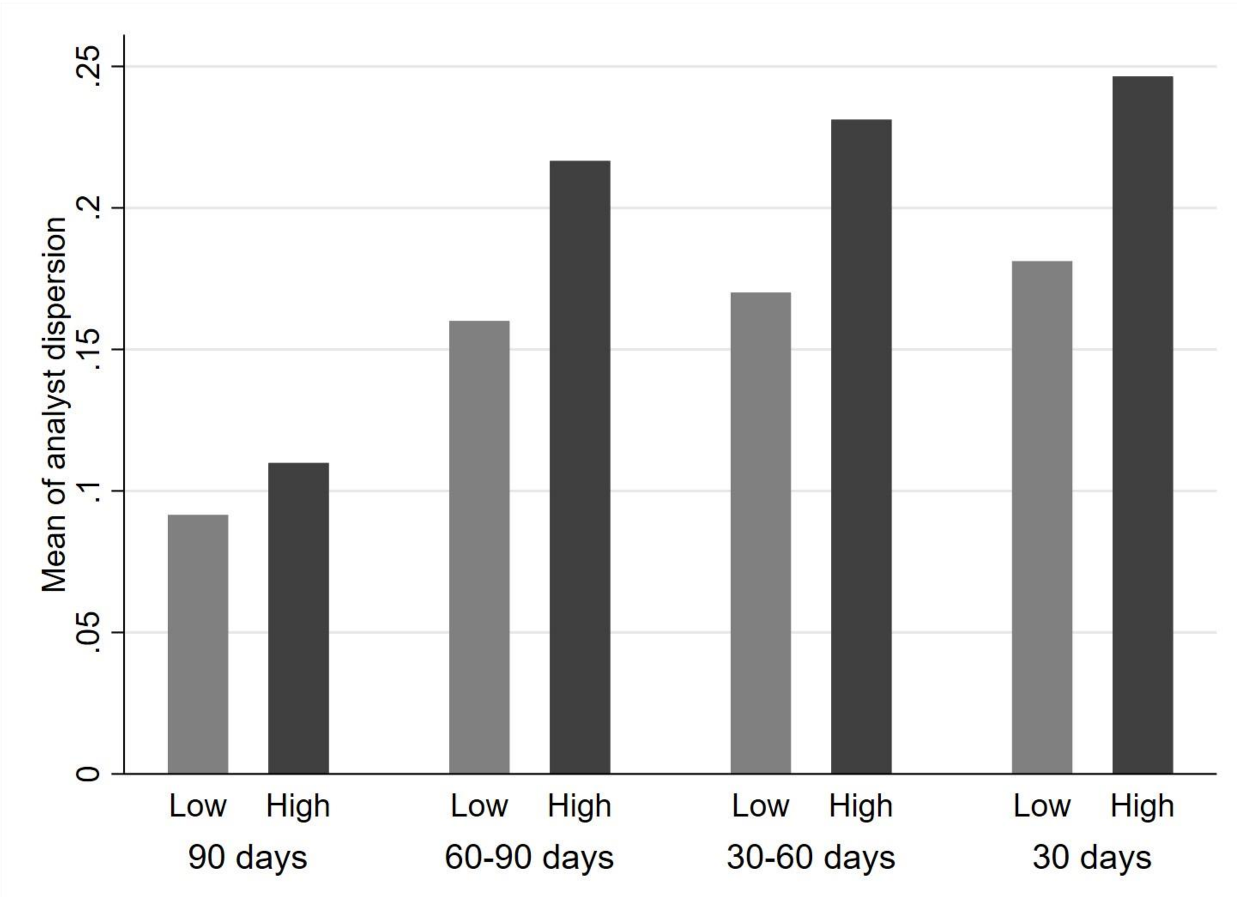

Table 1 – Sample descriptive statistics

This table provides the descriptive statistics for measures of bot activity, the dependent variables, and the control variables used in the analysis.

	N	Mean	SD	P25	Median	P75
The Social Reach of Tweets						
Number of company tweets on the given day	10,273	2.254	4.102	1	1	2
Number of tweets' likes on the given day	10,273	3,644	23,092	178	419	1301
Bot-Level Activity						
Bot ratio (likes)	10,273	0.170	0.061	0.127	0.162	0.204
Abnormal bot activity indicator (3-day window)	10,240	0.267	0.442	0.000	0.000	1.000
Abnormal bot activity indicator (7-day window)	10,270	0.245	0.430	0.000	0.000	0.000
Firm Characteristics						
Size	10,270	11.020	1.435	9.881	10.916	12.067
Book to market	10,270	0.306	0.352	0.092	0.213	0.407
Illiquidity	9,925	0.000	0.000	0.000	0.000	0.000
R_{m-1}	10,271	1.007	0.091	0.959	1.009	1.055
$R_{[m-3,m-2]}$	10,271	1.014	0.121	0.951	1.017	1.077
$R_{[m-6,m-4]}$	10,269	1.038	0.155	0.950	1.032	1.125
$R_{[m-12,m-6]}$	10,263	1.084	0.243	0.939	1.069	1.202
Earnings and Forecasts by Analysts						
Earnings surprise	9,225	0.102	0.574	0.011	0.068	0.172
Analyst dispersion	9,514	0.073	0.628	0.021	0.038	0.077
Number of analysts covering a stock	9,554	3.104	0.319	2.890	3.091	3.332
Indicator for the fourth quarter	10,273	0.227	0.419	0.000	0.000	0.000
Investor Reaction to Earnings						
Volatility surrounding earnings (-7,7)	10,258	0.023	0.014	0.014	0.020	0.028
Volatility surrounding earnings (-7,0)	10,259	0.021	0.014	0.011	0.017	0.025
Volatility surrounding earnings (0,7)	10,259	0.026	0.018	0.015	0.022	0.033
AbLogVol	10,268	0.117	0.530	-0.233	0.042	0.392
Absolute analyst dispersion	9,170	0.050	0.980	0.021	0.038	0.077
Abnormal discretionary expenses	8,309	-0.091	0.678	-0.067	-0.019	0.001

Table 2. The effect of bot activities on the standard deviation of the returns

This table analyzes the impact of bot-level activities on return volatility (defined as the standard deviation of returns) over the specified period. For observations before the earnings announcement date, the volatility is computed for the period (-7,0). Similarly, for the period after the QE, the volatility is calculated using the period (0,7). Columns (1) and (2) employ the whole data set, i.e., one week before and after the QE. Column (3) consists of only observation from the before period, i.e., (-7,0), while in column (4), the corresponding time window is (0,7). All columns contain the 2SLS regression results computed using daily data on bot activity proxied by the total bot ratio, which is defined as the ratio of suspended accounts to total accounts. All variables are described in Appendix A2. The specification of IV regression is available in Equation (1), and the results of the first stage are provided in Appendix A3. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	No Controls (1)	Full (2)	Before QE (3)	After QE (4)
Bot ratio	0.217***	0.098***	0.121***	0.094***
	(0.018)	(0.020)	(0.029)	(0.027)
Size		-0.003***	-0.002***	-0.004***
		(0.000)	(0.000)	(0.000)
Book to market		0.001	0.002**	0.000
		(0.001)	(0.001)	(0.001)
Dummy fourth quarter		-0.001***	-0.002***	-0.000
		(0.000)	(0.000)	(0.001)
Illiquidity		41.847***	29.833***	53.491***
		(4.987)	(6.477)	(6.285)
Analyst dispersion		-0.001**	-0.001*	-0.001*
		(0.000)	(0.001)	(0.000)
Analyst number		0.004***	0.002*	0.007***
		(0.001)	(0.001)	(0.001)
R_{m-1}		-0.005*	-0.002	-0.005
		(0.003)	(0.003)	(0.004)
$R_{[m-3,m-2]}$		-0.013***	-0.016***	-0.011***
		(0.002)	(0.003)	(0.004)
$R_{[m-6,m-4]}$		-0.004***	-0.003*	-0.004**
		(0.001)	(0.002)	(0.002)
$R_{[m-12,m-6]}$		0.000	0.001	-0.001
		(0.001)	(0.001)	(0.001)
Constant	-0.031***	0.026***	0.014*	0.028***
	(0.004)	(0.006)	(0.008)	(0.009)
2SLS	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry FE	No	Yes	Yes	Yes
R-squared	0.07	0.255	0.295	0.294
N (observations)	10,226	9,138	4,466	4,672

Table 3. The effect of bot activities on the abnormal trading volume

This table analyzes the impact of bot-level activities on abnormal trading volume (AbLogVol) on the corresponding day. Column (3) consists of only observation from the (-7,0) period, while in column (4), the window is (0,7). All columns contain the 2SLS regression results computed using daily data on bot activity proxied by the *Total bot ratio*, which is defined as the ratio of suspended accounts to total accounts. All variables are described in Appendix A2. The specification of IV regression is available in Equation (1), and the results of the first stage are provided in Appendix A3. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	No Controls (1)	Full (2)	Before QE (3)	After QE (4)
Bot ratio	4.171***	2.400***	1.956**	3.106***
	(0.515)	(0.599)	(0.884)	(0.820)
Size		-0.001	-0.017*	0.015*
		(0.007)	(0.009)	(0.009)
Book to market		-0.057***	-0.095***	-0.020
		(0.019)	(0.026)	(0.028)
Dummy fourth quarter		0.108***	0.111***	0.095***
		(0.013)	(0.019)	(0.018)
Illiquidity		72.412	99.968	58.702
		(135.448)	(190.156)	(190.931)
Analyst dispersion		-0.039***	-0.031*	-0.053***
		(0.013)	(0.017)	(0.014)
Analyst number		0.035	0.080**	-0.014
		(0.026)	(0.037)	(0.036)
R_{m-1}		0.038	0.063	-0.027
		(0.083)	(0.119)	(0.114)
$R_{[m-3,m-2]}$		-0.553***	-0.638***	-0.483***
		(0.068)	(0.095)	(0.094)
$R_{[m-6,m-4]}$		-0.120***	-0.184***	-0.069
		(0.039)	(0.057)	(0.051)
$R_{[m-12,m-6]}$		0.083***	0.182***	-0.011
		(0.027)	(0.038)	(0.037)
Constant	-0.925***	0.049	0.478	-0.286
	(0.122)	(0.212)	(0.303)	(0.286)
2SLS	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry FE	No	Yes	Yes	Yes
R-squared	0.021	0.068	0.103	0.079
N (obs)	10,235	9,147	4,469	4,678

Table 4. Bot activity and managerial short-termism

This table analyzes the association between the bot-level activities during QE and CEO myopic behavior (i.e., short-termism). The dependent variable is the abnormal discretionary expenses (Abnormal Expenses, defined as the regression residuals in Equation (3)). These serve as a proxy for short-termism since cutting discretionary expenses leads to artificially higher EPS. All columns contain the 2SLS regression results computed using daily data on bot activity proxied by the *Total bot ratio*, which is defined as the ratio of suspended accounts to total accounts. All variables are described in Appendix A2. The specification of IV regression is available in Equation (1), and the results of the first stage are provided in Appendix A3. Columns (3) and (4) factor the results for subsamples of firms' alignment with the analyst expectations. Column (3) contains the results for the firms that are in line with analyst expectations, and column (4) contains the results for firms that are beating their expectations. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	No Controls (1)	Full (2)	Earnings are in Line (3)	Positive Surprise (4)
Bot ratio	-1.206**	-1.476**	-0.006	-2.678***
	(0.533)	(0.636)	(1.003)	(0.603)
Size		-0.000	0.013	-0.003
		(0.008)	(0.010)	(0.012)
Book to market		0.005	-0.028	-0.017
		(0.018)	(0.030)	(0.016)
Dummy fourth quarter		-0.140***	-0.166***	-0.049*
		(0.024)	(0.030)	(0.029)
Illiquidity		448.492**	1018.800***	72.319
		(174.898)	(296.146)	(157.434)
Analyst dispersion		-0.008	0.010	-0.030**
		(0.007)	(0.008)	(0.013)
Analyst number		0.031	0.152**	-0.143***
		(0.040)	(0.061)	(0.041)
R_{m-1}		0.381***	0.683***	0.005
		(0.093)	(0.133)	(0.130)
$R_{[m-3,m-2]}$		-0.146**	-0.362***	0.209**
		(0.064)	(0.094)	(0.082)
$R_{[m-6,m-4]}$		0.103***	0.056	0.146***
		(0.033)	(0.051)	(0.046)
$R_{[m-12,m-6]}$		0.055**	0.100***	-0.001
		(0.025)	(0.037)	(0.038)
Constant	0.385***	-0.384	-1.409**	0.832***
	(0.140)	(0.481)	(0.716)	(0.246)
2SLS	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry FE	No	Yes	Yes	Yes
R-squared	0.036	0.112	0.149	0.109
N (obs)	8,286	8,237	4,900	3,337

Table 5. The long-term effect on return volatility (SD)

This table shows the impact of bot activity around the quarterly earnings announcement on the long-term return volatility of the returns. Before and after periods are defined using seven calendar days around the QE. Specifically, the period before the QE covers the window (-7,0), while the period after the QE covers the interval (0,7). The dependent variable is returns volatility, calculated by taking the standard deviation of returns in the specified period after the earnings. All variables are described in Appendix A2. All columns contain the 2SLS regression results computed using daily data on bot activity proxied by the *Total bot ratio*, which is defined as the ratio of suspended accounts to total accounts. The specification of IV regression is available in Equation (1), and the results of the first stage are provided in Appendix A3. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	SD 1 week		SD 2 weeks		SD 3 weeks		SD 1 month		SD 2 months		SD 3 months	
	before	after	before	after	before	after	before	after	before	after	before	after
Bot ratio	0.098***	0.100***	0.112***	0.083***	0.109***	0.064***	0.098***	0.040**	0.066***	0.019	0.047***	0.007
	(0.034)	(0.033)	(0.030)	(0.026)	(0.026)	(0.022)	(0.023)	(0.020)	(0.019)	(0.015)	(0.017)	(0.013)
Constant	0.023**	0.030***	0.009	0.028***	-0.001	0.024***	-0.006	0.027***	0.004	0.035***	0.009	0.034***
	(0.011)	(0.011)	(0.008)	(0.008)	(0.008)	(0.007)	(0.008)	(0.006)	(0.007)	(0.006)	(0.006)	(0.005)
2SLS	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.252	0.260	0.296	0.305	0.332	0.336	0.347	0.341	0.406	0.401	0.426	0.422
N (obs)	4,466	4,672	4,466	4,672	4,464	4,671	4,464	4,671	4,459	4,666	4,458	4,665

Table 6. The effect of the bot activity on future analyst dispersion

This table analyzes the impact of bot-level activities on absolute analyst dispersion for the following quarterly earnings. The absolute analyst dispersion is defined as the absolute value of the standard deviation of analysts' forecasts scaled by the average of their projections for the quarterly earnings per share before the quarter-end date. All columns contain the 2SLS regression results computed using daily data on bot activity proxied by the *Total bot ratio*, which is defined as the ratio of suspended accounts to total accounts. All variables are described in Appendix A2. The specification of IV regression is available in Equation (1), and the results of the first stage are provided in Appendix A3. Columns (3) and (4) factor the results for subsamples of firms' alignment with the analyst expectations. Column (3) contains the results for the firms in line with analyst expectations, and column (4) contains the results for firms beating their expectations. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	No Controls (1)	Full (2)	Earnings are in line (3)	Positive Surprise (4)
Bot ratio	2.303***	1.987**	3.418***	-0.186
	(0.795)	(0.909)	(1.296)	(1.007)
Size		0.007	0.059***	-0.084***
		(0.015)	(0.018)	(0.019)
Book to market		-0.109**	-0.058	-0.130***
		(0.047)	(0.078)	(0.037)
Dummy fourth quarter		0.019	0.033	-0.005
		(0.024)	(0.030)	(0.030)
Illiquidity		3,156***	4,398***	839**
		(562.8)	(779.0)	(355.7)
R_{m-1}		-0.905***	-1.093***	-0.591
		(0.240)	(0.215)	(0.478)
$R_{[m-3,m-2]}$		0.428***	0.286*	0.718***
		(0.134)	(0.164)	(0.222)
$R_{[m-6,m-4]}$		0.151**	0.041	0.171
		(0.077)	(0.053)	(0.138)
$R_{[m-12,m-6]}$		-0.007	0.073**	-0.099
		(0.042)	(0.035)	(0.076)
Constant	-0.490***	-0.515*	-1.064**	0.529
	(0.187)	(0.310)	(0.452)	(0.408)
2SLS	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry FE	No	Yes	Yes	Yes
R-squared	0.027	0.083	0.134	0.076
N (obs)	9,149	8,790	5,210	3,618

Table 7: Excessive bot activity effects around QE

This table reports the Average Treatment Effect on Treated (ATET) results measuring the impact of excessive bot activities. The treatment group (T=1) is the companies with bot activity above the 75 percentile, while the control group is those with bot activity below the 45 percentile. We employ the nearest neighbor matching procedure *tefect* implemented in Stata. The treatment and control groups were exactly matched on the same QE period, and the other covariate was the company size. Matching is conducted using nearest neighbor matching on the common support, using Mahalabish distance. The standard errors of the ATET (in parentheses) are computed with the robust option (at least two suitable matches for each treated). ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively. At the bottom of each column, we present a balancing summary comparing differences in means and variance ratios for the treated and control samples. The remaining balancing test details, including the estimated kernel density graphs for firm size, are available upon request.

	Bot activity before QE			Bot activity after QE		
	(1)	(2)	(3)	(4)	(5)	(6)
	AbLogVol	Sd (-7,0)	Absolute analyst dispersion	AbLogVol	Sd (0,7)	Absolute analyst dispersion
ATET	0.170***	0.004***	0.222***	0.224***	0.002**	0.072
	(0.041)	(0.001)	(0.080)	(0.038)	(0.001)	(0.063)
p-value	<0.001	<0.001	0.005	<0.001	0.046	0.258
Mean of column variable (full sample)	0.404	0.020	0.164	0.277	0.026	0.156
Mean of column variable (treated)	0.541	0.023	0.301	0.444	0.028	0.222
ATET/mean (treated)	31.4%	17.4%	73.8%	50.5%	7.1%	32.4%
Number of treated	534	534	480	538	538	526
Number of observations	1,493	1,494	1,335	1,518	1,515	1,352
Balance summary (size)						
Standardized difference	-.013	-.011	-.015	0.008	0.009	-0.004
Variance ratio	1.093	1.088	1.081	1.076	1.076	1.036

Appendix

A1. Detection of malicious automation by 𝕏

To identify bots, we determine the number of suspended accounts that liked or reposted a given tweet by a company. We obtain a list of all accounts that liked or reposted a tweet by the company. As suspended users have missing information, we can calculate the ratios of bots versus normal users. This approach is supported by the official 𝕏 guidelines[29] that state: “Most[30] of the accounts we suspend are suspended because they are spammy, or just plain fake…”

The main benefit of this approach is that relying on 𝕏 to classify bot accounts will yield the lowest type 2 error, meaning that suspended users are very likely to be bots. That is because 𝕏 does not rely only on language analysis, which it argues is a too restrictive approach.[31] Instead, it uses more comprehensive analysis, and when it detects that an account might be engaged in manipulative behavior, it will send out a bot challenge[32] that the user needs to pass. This bot challenge consists of completing tasks that should be simple for human users to do, such as resetting a password or passing a CAPTCHA test, but which would be difficult or costly for bot accounts to solve. Accounts that fail to complete a challenge within a specified period of time are automatically suspended. It should be noted that while this approach will yield the lowest type 2 error, it could also only represent the lower bound of bot activity.

[29] Available at https://help.twitter.com/en/managing-your-account/suspended-twitter-accounts, last accessed January 18, 2023.

[30] For completeness, the account might be suspended for other reasons than being a bot (e.g., violence, terrorism, and child sexual exploitation, to name a few). We argue, however, that accounts of individuals reacting to official corporate accounts will not be associated frequently with other reasons for suspension.

[31] The detailed discussion is available at https://blog.twitter.com/en_us/topics/company/2020/bot-or-not.

[32] More detailed information available at https://transparency.twitter.com/en/reports/platform-manipulation.html#2021-jul-dec.

A2. Variable definition

Variable	*Description*
Dependent variables	
Volatility	Volatility is defined as the standard deviation of daily return, calculated for various windows surrounding earnings. E.g., (-7,0) corresponds to a window from 7 days before earnings until earnings day, and volatility is calculated as the standard deviation of all daily returns during that period. Data source: CRSP.
AbLogVol	The abnormal log trading volume on date t for firm i. It is calculated as the difference between the log volume on date t and the average log volume from trading days t − 140 to t – 20, following Cookson and Niessner (2020). Data source: CRSP.
Absolute dispersion	The absolute value of the standard deviation of analysts' forecasts scaled by the consensus for the quarterly earnings per share before the quarter-end date. Source: I/B/E/S.
Abnormal expenses	Following Roychowdhury (2006), we define abnormal discretionary expenses as the residuals of the following model: $\frac{XSGA_{it}}{Assets_{i,t-4}} = \beta_1 * \frac{1}{Assets_{i,t-4}} + \beta_2 * \frac{Sales_{i,t-4}}{Assets_{i,t-4}} + \epsilon_{it}$, where $XSGA_{it}$ is the SG&A costs including marketing and R&D for firm *i* in quarter *t*. The model is estimated for every quarter and the Fama-French 48 industries separately.
Bot activity measures	
Total bot ratio	The total bot ratio is defined as the number of suspended accounts that liked any tweet of the corporation on the given day divided by the total number of accounts that liked a tweet of the corporation on the given day.
3-days (7-days) global level of bot activity	The global level of bot activity is calculated as the mean of the total bot ratios, excluding the given firm, in a three-day (seven-day) moving window.
Flag 3-days (7-days) abnormal bot activity	This is an indicator variable that is equal to one if the bot ratio for the given day is larger than the corresponding global level of bot activity.
Difference from abnormal bot activity	The difference between the total bot ratio and the global level of bot activity. The global level of bot activity is calculated as the mean of the total bot ratios, apart from the given firm, in a three-day or seven-day moving window.

Firm control variables	
Firm size	Firm size is the natural logarithm of the market value of equity. Data sources: CRSP and Compustat.
Book-to-market ratio	The book-to-market ratio is defined as book equity divided by market equity. Data sources: CRSP and Compustat.
Past profitability	The group of variables $\mathbf{R}_{m-1}$, $\mathbf{R}_{[m-3,m-2]}$, $\mathbf{R}_{[m-6,m-4]}$, and $\mathbf{R}_{[m-12,m-6]}$, which stand for returns over the last month, months 3 to 2, 6 to 4, and 12 to 6, respectively. Defined by Brennan et al. (2012). Data sources: CRSP and Compustat.
Illiquidity	Illiquidity is the sum of the absolute values of daily returns divided by the daily volume for the year, multiplied by 10^6. Defined by Amihud (2002). Data sources: CRSP and Compustat.
Analyst dispersion	The standard deviation of analysts' forecasts was scaled by the consensus for the quarterly earnings per share before the quarter-end date. Source: I/B/E/S.
Number of analysts covering a stock	The natural logarithm of one plus the number of analysts forecasting the quarterly earnings per share in the latest I/B/E/S consensus before the quarter-end date. Source: I/B/E/S.
Fourth quarter indicator	The indicator is equal to one if the given quarter is the fourth fiscal quarter. Source: I/B/E/S.
Earnings surprise	The earnings surprise is measured as the I/B/E/S reported quarterly earnings per share less the latest I/B/E/S consensus analyst quarterly earnings per share forecast just prior to the quarterly earnings announcement date, which is scaled by stock price as of the forecast date and multiplied by 100. Data sources: CRSP and I/B/E/S.

A3. Summary of the First-Stage Regressions - Dependent Variable Bot Ratio

The explanatory variables include the 3- and 7-day global level of bot activities and their squares. Additional control variables are the year and quarterly dummy variables to control for global levels of bot activity associated with elections and bot detection scrutiny. The predicted outcomes were very close. The model and results are robust to the quadratic specification of variables. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Bot ratio (1)
The 3-day global level of bot activity	-0.107* (0.059)
The 7-day global level of bot activity	0.817*** (0.072)
Constant	0.072*** (0.010)
Year and quarter FE	Yes
R-squared	0.271
Adjusted R-squared	0.270
N (Observation)	10,240