Mendel University in Brno

MENDELU Working Papers
in Business and Economics
101/2025

X Bots and Earnings Announcements

Jan Hanousek, Jr., Jan Hanousek, 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., Hanousek, J., Sokolov, K. (2025). X Bots and Earnings Announcements. *MENDELU Working Papers in Business and Economics* 101/2025. Mendel University in Brno. Cited from: http://ideas.repec.org/s/men/wpaper.html

Abstract

Hanousek Jan Jr., Hanousek Jan, Sokolov Konstantin: **X Bots and Earnings Announcements.**

This paper relies on shocks to the CAPTCHA technology to study the effects of social media bot activity. We observe that bots drive a large amount of attention to corporate 𝕏 accounts around earnings announcements. In line with theoretical research, bot activity is a significant predictor of investor disagreement, which is persistent long-term. Moreover, social media bots increase analyst dispersion for the following earnings announcement. The failure of the CAPTCHA implementation technology leads to negative abnormal returns, and the failure of CAPTCHA bypassing technology leads to positive abnormal returns.

Key words

Investor disagreement, analyst 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

Jan Hanousek, Mendel University in Brno and CEPR, London. E-mail: jan.hanousek@mendelu.cz

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

Acknowledgements

The authors thank Boone Bowles, Anthony Cookson, Olivier Dessaint, Gerard Hoberg, Charles Martineau, and Dmitriy Muravyev for their insightful comments. We are also grateful for comments from participants of the Future of Financial Information Conference, Eastern Finance Association Meeting, UBC Summer Conference, Waseda Finance Workshop, and the seminar participants at the University of Kansas.

Introduction

A "bot" is a piece of software that imitates human communication. Multiple studies show that bots can be weaponized to manipulate and misinform the public.[1] We shed light on the outcomes of social media bot activity around earnings announcements.[2] Specifically, we rely on more than 4,500 unique shocks to the Completely Automated Public Turing test to tell Computers and Humans Apart (CAPTCHA) to identify the bot activity on 𝕏. The CAPTCHA is designed with a narrow purpose to curb bots. We illustrate that disruptions to CAPTCHA implementation lead to more intense bot activity. In what follows, we refer to these disruptions as "CAPTCHA bugs." In contrast, disruptions to CAPTCHA solving lead to lower bot activity. We refer to such disruptions as "SOLVER bugs." Such shocks are technological in nature, and they are unlikely to be reversely caused by the information in company-specific earnings announcements.

The theoretical literature suggests that corporate tweet[3] promotion by bots may induce disagreement among investors. Rabin and Schrag (1999) offer a model of confirmatory bias to illustrate that individual agents selectively adopt information to confirm their heterogeneous priors. The theory of Van den Steen (2011) further shows that agents with heterogeneous priors will become more confident about the precision of their priors as they receive a greater amount of information. The general theoretical framework of Bohren and Hauser (2021) shows that entrenched disagreement will likely arise even if the agents perceive the information correctly. This type of disagreement occurs due to misinterpreting peer response to such information.

The above theoretical literature implies that disagreement can increase as agents receive more data. Therefore, based on these models, we hypothesize that intense bot activity leads to investor disagreement.

The alternative hypothesis of how bots can affect capital markets arises from the technological consequences of bot activity. 𝕏 relies on advertising to generate income, and as a result, it uses algorithms to recommend posts and topics to users that they might be interested in to increase the amount of time the users spend on the site. The number of likes or retweets is thus a crucial variable that 𝕏 uses to decide which topics are popular and trending. Since 𝕏 cannot verify immediately which activity is being driven by bot attention, bot activity can lead to the given tweet being recommended to more users, even those not following the corporate 𝕏 account. The messages put out by the corporation will have a broader reach. Therefore, contrary to our main hypothesis, bot activity may alleviate investor disagreement.

Consistent with the theoretical implications of Rabin and Schrag (1999), Van den Steen (2011), and Bohren and Hauser (2021), we find that bot activity is a significant predictor of investor disagreement. As such, bots play a divisive role among investors. In particular, CAPTCHA bugs[4] lead to a 1.96% increase in abnormal volume and a 0.68% increase in the standard deviation of returns. Conversely, SOLVER bugs cause a decrease in abnormal volume of 1.09% and a 0.17% decrease in the standard deviation of returns.

We further show that not only investors but also analysts are affected by bot activity. When bot activity increases, it manipulates the popularity and reach of a tweet. Analysts rely on this manipulated popularity when making estimates. As a result, analyst dispersion, another indicator of investor disagreement, significantly increases in the next quarter following high bot activity. For example, a spike in SOLVER bugs will typically decrease analyst dispersion by 5.59%.

Moreover, we find that the effect of the bot activity is persistent for longer time horizons. First, an exogenous spike in bot activity typically leads to an increase in volatility that lasts up to the subsequent earnings announcement. Second, bot activity affects stock performance. An increase in CAPTCHA bugs during the current earnings announcement leads to lower returns in the following month. Specifically, a disruption to bot detection technology around the dates of the current earnings announcement will typically cause the stock to underperform by 2.7% in the following month. Conversely, an increase in SOLVER bugs will lead to stock overperformance by 1.7% in the two months following the earnings announcement. Overall, our results indicate that bot activity has persistent impacts that are economically significant.

The paper contributes to the growing literature on investor disagreement. Previous research shows that social media increases investor disagreement (Antweiler and Frank, 2004; Cookson and Niessner, 2020). Traditional finance literature finds that disagreement induces trading (e.g., Hirshleifer, 1977; Diamond and Verrecchia, 1981). Antweiler and Frank (2004) were the first to show that a message board activity and disagreement can predict market volatility and trading volume. Contemporary literature provides a deeper empirical look into how social media impacts disagreement. Cookson and Niessner (2020) use the investing platform StockTwits and analyze the differences between model-based and information-based disagreement. Cookson et al. (2023) show that users seek information that verifies their priors, thus leading to echo chambers and more profound disagreement. Hirshleifer et al. (2024) show that disagreement can persist even after quarterly earnings announcements in areas with more social transmission of ideas, as social connectedness leads to both news being incorporated into prices and divergence in opinion. While finance research on bot activity is sparse, Stella et al. (2018) and Broniatowski et al. (2018) show that bot activity is linked to conflict and polarization. As a result, we hypothesize that bot activity will have a polarizing effect on investors. In line with this hypothesis, we find that bot activity increases abnormal trading volume, volatility, and analyst dispersion.

Social media platforms and bot developers engage in technological warfare.[5] Bot developers employ Artificial Intelligence (AI) tools to solve the CAPTCHA, while social media platforms implement the CAPTCHA with AI tools to counteract bots. For example, Botright, one of the tools to solve CAPTCHAs, has the following description: "Solving your Captchas for free with AI."[6] On the opposing side, the description of the CAPTCHA implementation tool Color-captcha reads, "Dynamically generated, AI-resistant CAPTCHA images."[7] As such, both implementation and solving tools use AI. Our analysis relies on bugs in such CAPTCHA implementation and CAPTCHA solving tools to identify exogenous shocks to bot activity.[8]

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 in Section 4. Section 5 shows robustness to alternative specifications. Lastly, Section 6 concludes the paper.

1 The bot industry

Social media bots are accounts controlled by software designed to complete specific tasks. These tasks typically involve generating messages to sway public opinion (Yang et al., 2019) and amplifying certain narratives (Keller et al., 2020). Bessi and Ferrara (2016) provide evidence that such tactics are effective. Modern bots possess the intelligence necessary to automatically engage in natural conversations (Assenmacher et al., 2020). The widespread availability and reduced cost of AI tools, as well as the ease of implementation (Brundage et al., 2018; Weidinger et al., 2022), lead to large adoption rates for social media bots (Yang and Menczer, 2023). The threats posed by AI have been articulated in the past (e.g., Bommasani et al., 2021; Yamin et al., 2021), and there is notable anecdotal evidence (e.g., Hanley and Durumeric, 2023).

The more human-like AI bots appear, the better they can convince people of the messages they are propagating (Spitale et al., 2023). For example, AI can decide which account under its control will propagate a given message and how frequently and at what time the account posts new messages to avoid appearing suspicious. Naturally, this setting allows the AI to learn, since it can observe which accounts receive the most responses.

𝕏 asserts that earning people's trust is the company's core value.[9] This trust, however, may be eroded if a large proportion of 𝕏 activity comes from inauthentic users. Officially, 𝕏 has reported in its 10-Q form that only 5% of the activity is attributable to bots (e.g., the 10-Q report from March 31, 2022). However, this figure has been heavily scrutinized.[10] One of the most publicized criticisms came from Elon Musk before he acquired 𝕏, as he claimed bot activity to be far more than 5% previously reported by 𝕏 (Duffy and Fung, 2022). Furthermore, Timber and Dwoskin (2018) report that 𝕏 suspended more than one 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 by 𝕏, in which 𝕏 officials committed to providing a better platform for healthy civic discourse.[11] However, this was not the first of assurances from 𝕏, with the company pledging to fight spam in a 2012 press release.[12] For example, Twitter's co-founder, 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."[13]

The problem of bot accounts on social media is well-established in 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.

Companies are aware of the importance of social media presence and reach. Frequently, this awareness leads to the purchase of bots by their marketing and public relations agencies to meet their goals faster (Confessore et al., 2018). The negative consequences of such a decision are limited, as 𝕏 does not typically suspend users suspected of buying bots (Confessore et al., 2018). This is because while 𝕏 can detect and block possible bot activity, it does not typically know who is responsible for a bot purchase. Moreover, it is essential to note that 𝕏 may benefit from bot accounts since the individuals controlling them need to purchase the commercial 𝕏 API that allows automation.

On several occasions, 𝕏 has acknowledged that the bot activity constitutes a problem (e.g., Timber and Dwoskin, 2018). To address the issue, 𝕏 has developed several initiatives to increase transparency and offer tools to encourage research on bot activity. Notably, it created a website focusing on platform manipulation that documents the number of spam reports and bot challenges.[14] However, these initiatives were suspended after 𝕏 became private, and they have not been reinstated since the beginning of 2022.

Moreover, in February 2023, 𝕏 announced the termination and removal of free academic API access.[15] Until this time, it had given researchers access to 10 million tweets monthly for no charge, and the platform's data had been used in as many as 17,500 academic papers since 2020.[16] However, 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 per month. As the data for this study was collected in January 2023, it gives us a unique opportunity to shed light on 𝕏 bot activity with the most recent complete and affordable dataset for academic research.

2 Bot aktivity and CAPTCHA

2.1 CAPTCHA

Our analysis requires addressing the concern over the endogeneity of the bot activity on the firm's 𝕏 account. We take advantage of technological shocks to bot detection. One of the most common ways to limit bot activity is to employ the Completely Automated Public Turing test to tell Computers and Humans Apart (CAPTCHA). The application of CAPTCHA technology is very narrow, and the issues that disrupt this technology occur randomly. Therefore, the shocks to the CAPTCHA are unlikely to affect the company-specific financial variables through means other than bots. As such, the dynamics of disruptions to CAPTCHA serve as an instrument for the bot activity.

Specifically, we collect unique data on innovations in CAPTCHA technology. These data have two main constituents. The first constituent comprises technological innovations that improve CAPTCHA. Such innovations serve the purpose of curbing bot activity. In contrast, the second constituent comprises innovations in CAPTCHA solving technologies. Such innovations allow bot developers to circumvent CAPTCHA challenges and thus intensify bot activity. We use these shocks in CAPTCHA implementation and CAPTCHA solving technologies to identify the impact of bot activity.

We rely on the information publicly available on www.github.com to measure shocks to the CAPTCHA. GitHub is the largest developer platform for storing open-source codes. Each application has a repository, which includes documentation and all codes and resources needed to run it. Appendix A1 provides an example of a GitHub repository focusing on solving CAPTCHA. According to Hoffmann et al. (2024), proprietary software developers often borrow codes from open-source repositories, which allows them to reduce operating costs by as much as 3.5 times. Furthermore, GitHub has a material impact on entrepreneurship (Wright et al., 2023) and innovation (Conti, et al., 2024) dynamics. Consistent with the literature, we show that the dynamics of GitHub pick up the evolution of both CAPTCHA solving and implementing technologies.

We use textual analysis to categorize repositories posted on GitHub as focusing on CAPTCHA implementation and those focusing on CAPTCHA solving. We provide a detailed description of data processing in Appendix A2. We end up with 1,595 repositories in the CAPTCHA implementation and 2,642 in the CAPTCHA solving group. To estimate the dynamics of CAPTCHA implementation and solving technologies, we parse GitHub repositories for issues (i.e., bugs) reported and solved for each repository. We then aggregate bugs in these two categories into a series of unresolved bugs. The variable *SOLVER_Bugs* is the number of unresolved issues in the CAPTCHA solving repositories. Similarly, *CAPTCHA_Bugs* is the number of unresolved issues in the CAPTCHA implementation repositories.

Figure 1 plots the dynamics of *SOLVER_Bugs and CAPTCHA_Bugs* over the sample period. The correlation between the two series is positive $(\rho=0.4)$. As such, bugs in both CAPTCHA implementation and solving technologies tend to appear simultaneously. Despite this positive correlation, the next section shows that these variables have the opposite economic effect on bots.

Figure 1. Shocks to the CAPTCHA technology

This figure shows shocks to the CAPTCHA technology, which affects bot detection. SOLVER bugs is the number of unresolved issues in the GitHub repositories containing CAPTCHA solving solutions. CAPTCHA bugs is the number of unresolved issues in the GitHub repositories containing CAPTCHA implementation solutions. We show their movements and impacts consistently with the main specification, i.e., with log(X) transformation and adjusted for annual means.

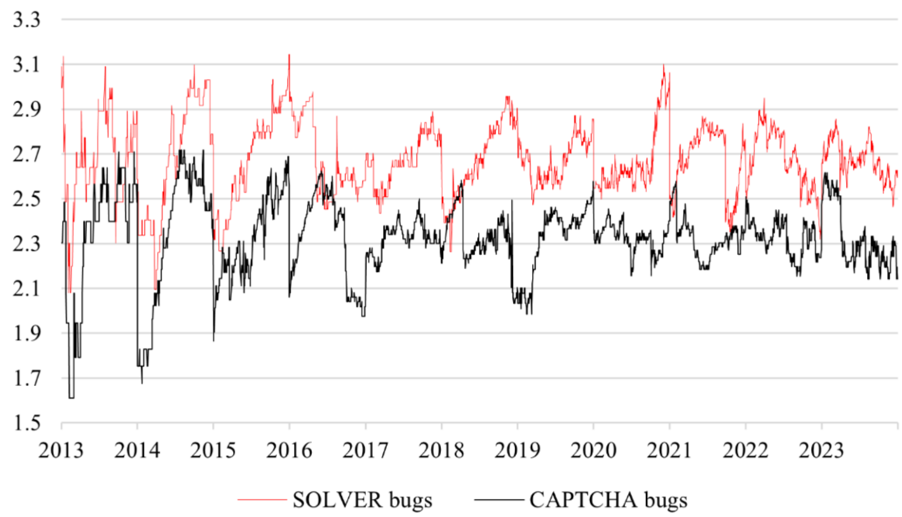

2.2 Instrument validity

In this section, we illustrate that the shocks to CAPTCHA implementation and solving technology significantly impact 𝕏 bots. We begin by scraping official corporate 𝕏 accounts of the S&P 500 companies. Specifically, we rely on the API offered by 𝕏 to scrape all original tweets.[17] In April 2013, the SEC allowed US-listed firms to disclose publicly 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 2 shows that the corporate 𝕏 accounts experienced an overall activity spike around the quarterly earnings announcements. Since the majority of corporate activity on 𝕏 is happening around earnings announcements, we focus on bot activity around the earnings announcement dates. Precisely, we follow Yang et al. (2020) and focus on a period of 15 days surrounding the quarterly earnings announcement date.

Figure 2. The number of company tweets 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, while (0; 7) represents the week after the earnings. We control for the company fixed effects.

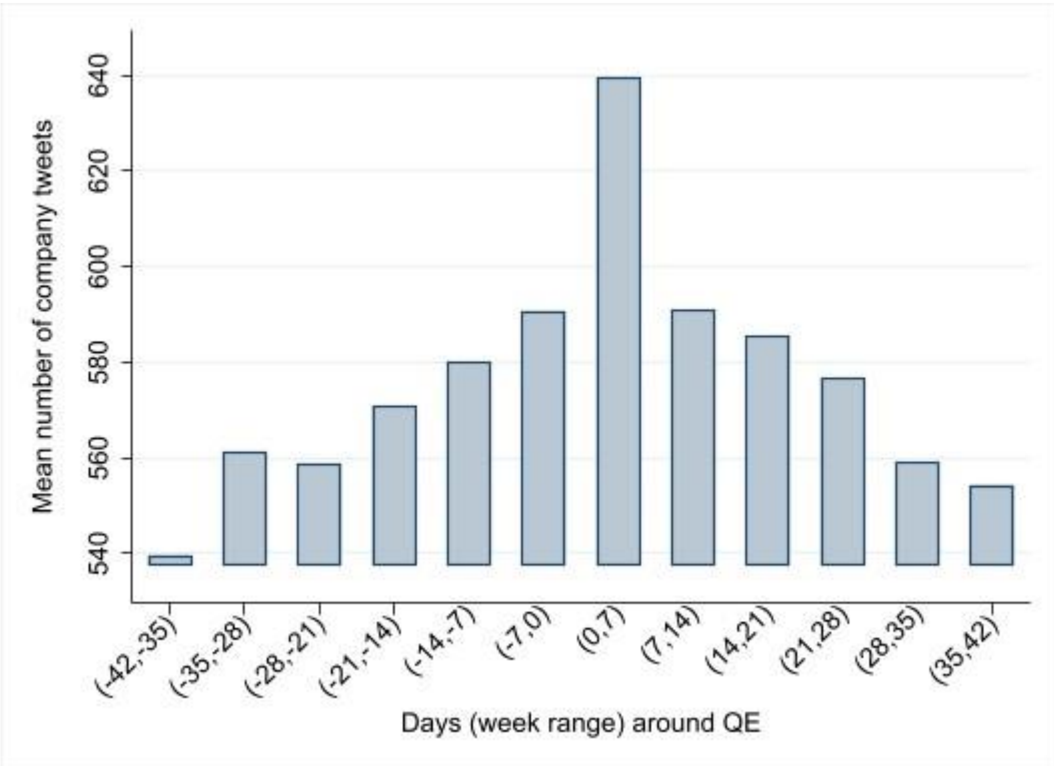

We proceed with measuring 𝕏 bot activity around earnings announcements. The bot identification procedure involves the following steps. First, we collect data on all 𝕏 users who liked or reposted the original S&P 500 companies' tweets. Then, we check whether 𝕏 has later suspended these users. We acknowledge that 𝕏 suspends accounts for reasons 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, the accounts engaging in these illicit activities are unlikely to react to corporate earnings or be suspended for failing the CAPTCHA tests. We provide more detailed information and a description of the data collection process in Appendix A3.

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 these, 7,641,950 (20.3%) likes were made by bot accounts. Not all bot accounts may be suspended by 𝕏, which may result in underestimating bot activity. We address this concern in Section 5.1.

We estimate the intensity of bot activity as *BotRatio*.

$$BotRatio_{i,t} = \frac{BotLikes_{i,t}}{TotalLikes_{i,t}} \tag{1}$$

where $BotLikes_{i,t}$ is the number of likes by suspended accounts on company *i* tweets that were tweeted on day *t*. Similarly, $TotalLikes_{i,t}$ is the total number of likes on the official tweets of company *i* that were tweeted on day *t*.[18] 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.

Table 1 presents the summary statistics. The average *BotRatio* is 17%. This figure is more prominent than the 5% baseline level reported by 𝕏's 10-Q report and is greater than the estimate of 11% provided by Duffy and Fung (2022). This result indicates that the 𝕏 accounts of S&P 500 companies experience a relatively higher level of attention from bots than other users' accounts.

Table 1. Sample descriptive statistics

This table provides the descriptive statistics for the measures of bot activity, dependent variables, and control variables used in the analysis. Appendix A4 contains detailed definitions of the variables.

	Mean	**SD**	**P25**	**Median**	**P75**
The Social Reach of Tweets					
Number of company tweets per day	2.254	4.102	1	1	2
Number of tweet likes per day	3,644	23,092	178	419	1301
Bot-Level Activity					
CAPTCHA bugs	4.069	0.144	3.951	4.078	4.174
SOLVER bugs	4.108	0.529	3.611	4.007	4.615
Bot ratio	0.170	0.032	0.149	0.165	0.187
Firm Characteristics					
Size	11.02	1.435	9.881	10.916	12.067
Book to market	0.306	0.352	0.092	0.213	0.407
Illiquidity	0.475	0.603	0.112	0.269	0.625
Rm−1Rm−1	0.007	0.091	-0.041	0.009	0.055
R[m−3,m−2]Rm−3,m−2	0.014	0.121	-0.049	0.017	0.077
R[m−6,m−4]Rm−6,m−4	0.038	0.155	-0.05	0.032	0.125
R[m−12, m−6]Rm−12, m−6	0.084	0.243	-0.061	0.069	0.202
Earnings and Forecasts by Analysts					
Analyst dispersion	0.073	0.628	0.021	0.038	0.077
Number of analysts covering a stock	3.104	0.319	2.89	3.091	3.332
Reaction to Earnings					
Volatility	0.024	0.016	0.013	0.019	0.030
AbLogVol	0.117	0.53	-0.233	0.042	0.392
Absolute analyst dispersion	0.05	0.98	0.021	0.038	0.077
N (observations)	10,273				

We proceed by testing whether disruptions to CAPTCHA are affecting the *BotRatio.* Specifically, we run the following regression.

$$BotRatio_{i,t} = \alpha_i + \beta_1 SOLVE\_Bugs_t + \beta_2 CAPTCHA\_Bugs_t + \epsilon_{i,t} \quad (2)$$

Where the variable $SOLVE\_Bugs_t$ is the natural logarithm of the total number of unresolved bugs on day *t* for the GitHub repositories that aim to solve CAPTCHA, and $CAPTCHA\_Bugs_t$ is the natural logarithm of the total number of unresolved bugs on day *t* for the GitHub repositories that aim to implement CAPTCHA. We report the results in Table 2.

Table 2. CAPTCHA shocks and bot ratio

The table shows that shocks to CAPTCHA implementation and solving technologies have impact on 𝕏 bots. Specifically, the table reports estimated coefficients in Equation (1). The dependent variable is the proportion of likes given by suspended bot accounts in the total number of likes for a corporate tweets on the given day (Bot ratio). CAPTCHA bugs are the disruptions to the CAPTCHA technology. Higher value of CAPTCHA bugs implies less restricted bot activity. SOLVER bugs are the disruptions to the CAPTCHA solving technology. Higher value of SOLVER bugs implies more restricted bot activity. Appendix A4 contains detailed definitions of the variables. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
CAPTCHA bugs	0.034***	0.022***	0.029***	0.018***
	(0.004)	(0.005)	(0.003)	(0.004)
SOLVER bugs	-0.053***	0.002	-0.048***	-0.008**
	(0.001)	(0.004)	(0.001)	(0.003)
Constant	0.249***	0.138***	0.184***	0.122**
	(0.015)	(0.027)	(0.048)	(0.050)
Industry fixed effects	No	No	Yes	Yes
Year fixed effects	No	Yes	No	Yes
F-statistic	1204.38	408.42	197.38	195.51
R-squared	0.190	0.241	0.442	0.473
Adjusted R-squared	0.190	0.241	0.439	0.471
N (Observation)	10,273	10,273	10,270	10,270

The coefficient estimates reported in Table 2 confirm that the shocks to CAPTCHA implementation and solving technology impact bot activity. We observe that $SOLVER_{Bugs_t}$ impacts the bot ratio negatively. Disruptions in the algorithms employed by bots to solve CAPTCHA decrease bot activity. Typically, a 1% increase in $SOLVER_{Bugs_t}$ typically causes a drop in the *BotRatio* of about 1.3%. Conversely, $CAPTCHA_{Bugs_t}$positively impacts bot ratio, which shows that bot developers can exploit disruptions in CAPTCHA to avoid detection. A 1% increase in $CAPTCHA_{Bugs_t}$ leads to an increase in *BotRatio* by 0.8%.

In the following analysis, we use $SOLVER_{Bugs_t}$ and $CAPTCHA_{Bugs_t}$ directly as exogenous variables to capture the effect of bots on disagreement. Section 5.1 shows that our findings are robust to a two-stage least squares (2SLS) setup, where we use Equation (2) as the first-stage regression. ¨

Our main regression specification is as follows.

$$DepVar_{i,t} = \alpha_i + \beta_1 CAPTCHA\_Bugs_t + \beta_2 SOLVER\_Bugs_t + \epsilon_{i,t} \quad (3)$$

where $DepVar_{i,t}$ is one of the dependent variables for which we estimate the effects of investor disagreement associated with bot activities. Variables $CAPTCHA\_Bugs_t$ and $SOLVER\_Bugs_t$ capture exogenous shocks to bot activity. $Controls_{i,t}$ include size, book-to-market value ratio, the dummy equal to one for the announcement of annual results, illiquidity, analyst dispersion, number of analysts, and the company's past returns. Finally, we include year and industry fixed effects. The next section describes the rationale for selecting the control variables. We provide the estimation details in Appendix A4.

2.3 Sample contruction

Next, we define the dependent variables to analyze the investor disagreement and price effects of bot activity. First, 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 t120 and t20. Second, we follow Antweiler and Frank (2004) and use volatility computed as the standard deviation of stock returns as a measure of investor disagreement. 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 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 in 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 definitions of these control variables are available in Appendix A4.

3 Results

3.1 Bot aktivity and investor disagreement

We begin by testing how bot activity impacts investor disagreement. On the one hand, information proliferation by bots may reduce information asymmetry and mitigate investor disagreement. On the other hand, bots may exacerbate disagreement by ingraining heterogeneous confirmation biases (Rabin and Schrag, 1999 and Van den Steen, 2011) and further entrenching disagreement through the misinterpretation of peer response to information (Bohren and Hauser, 2021).

Following the literature on investor disagreement, we test the impact of bot activity on return and volume volatility (e.g., Hirshleifer, 1977). Cookson and Niessner (2020) argue for using abnormal volume rather than return volatility to study the effects of disagreement, as return volatility relies on observed trading patterns. They develop a measure of abnormal volume for each stock and day (AbLogVol), defined as the natural logarithm of share volume on the day minus the natural logarithm of average volume over days $t - 140$ to $t - 20$. Identical to their methodology, we analyze the impact of disagreement brought on by bot activity on abnormal volume, which we report in Table 3. We can see that bot activity significantly increases trading volume during the quarterly earnings period. Specifically, a 1% increase in CAPTCHA bugs, which reduces barriers to entry for bots, leads to an abnormal volume increase of 1.96%. Similarly, a 1% increase in SOLVER bugs leads to a decrease in the abnormal volume of 1.09%.[19]

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

This table analyzes the impact of bot activity on abnormal trading volume (AbLogVol). All variables are described in Appendix A4. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
CAPTCHA bugs	0.113**	0.188***		0.170***
	(0.049)	(0.050)		(0.051)
SOLVER bugs	-0.173***		-0.104**	-0.074*
	(0.038)		(0.041)	(0.042)
Size		-0.004	-0.002	-0.003
		(0.007)	(0.007)	(0.007)
Book to market		-0.060***	-0.061***	-0.061***
		(0.019)	(0.019)	(0.019)
Dummy fourth quarter		0.123***	0.113***	0.118***
		(0.013)	(0.013)	(0.013)
Illiquidity		37.181	74.049	43.370
		(136.537)	(135.450)	(136.211)
Analyst dispersion		-0.039***	-0.040***	-0.039***
		(0.013)	(0.014)	(0.013)
Analyst number		0.040	0.039	0.040
		(0.026)	(0.026)	(0.026)
Rm−1Rm−1		0.018	0.009	0.012
		(0.083)	(0.083)	(0.083)
R[m−3,m−2] Rm−3,m−2		-0.556***	-0.530***	-0.549***
		(0.068)	(0.067)	(0.068)
R[m−6,m−4]Rm−6,m−4		-0.119***	-0.111***	-0.116***
		(0.039)	(0.039)	(0.039)
R[m−12, m−6]Rm−12, m−6		0.085***	0.082***	0.085***
		(0.027)	(0.027)	(0.027)
Constant	0.643**	-0.156	0.947***	0.169
	(0.259)	(0.247)	(0.217)	(0.313)
Year FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
R-squared	0.037	0.067	0.066	0.067
N (observations)	10,266	9,171	9,171	9,171

However, to be consistent with the prior literature on investor disagreement, we also test the impact on return volatility, which we define as the standard deviation of daily returns. Table 4 shows that bot activity also positively and significantly increases volatility. Consistent with the hypothesis that bots induce disagreement, a 1% increase in CAPTCHA bugs increases the volatility of returns by 0.68%.

Table 4. The effect of bot activity on the standard deviation of the returns

This table analyzes the impact of bot activity on volatility defined as the standard deviation of returns. All variables are described in Appendix A4. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
CAPTCHA bugs	0.004***	0.004***		0.004***
	(0.001)	(0.001)		(0.001)
SOLVER bugs	-0.001		-0.001	-0.000
	(0.001)		(0.001)	(0.001)
Size		-0.003***	-0.003***	-0.003***
		(0.000)	(0.000)	(0.000)
Book to market		0.001	0.001	0.001
		(0.001)	(0.001)	(0.001)
Dummy fourth quarter		-0.000	-0.001	-0.000
		(0.000)	(0.000)	(0.000)
Illiquidity		41.000***	41.780***	41.030***
		(5.000)	(4.999)	(5.002)
Analyst dispersion		-0.001**	-0.001**	-0.001**
		(0.000)	(0.000)	(0.000)
Analyst number		0.004***	0.004***	0.004***
		(0.001)	(0.001)	(0.001)
Rm−1Rm−1		-0.005*	-0.005*	-0.005*
		(0.003)	(0.003)	(0.003)
R[m−3,m−2] Rm−3,m−2		-0.013***	-0.013***	-0.013***
		(0.002)	(0.002)	(0.002)
R[m−6,m−4]Rm−6,m−4		-0.003**	-0.003**	-0.003**
		(0.001)	(0.001)	(0.001)
R[m−12, m−6]Rm−12, m−6		0.000	0.000	0.000
		(0.001)	(0.001)	(0.001)
Constant	0.003	0.032***	0.053***	0.034***
	(0.007)	(0.007)	(0.006)	(0.008)
Year FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
R-squared	0.164	0.252	0.251	0.252
N (observations)	10,257	9,162	9,162	9,162

Overall, these results suggest that bot activity significantly impacts both the return volatility and the abnormal volume. As such, bots lead to more considerable investor disagreement, which is in line with the theoretical implications of Rabin and Schrag (1999), Van den Steen (2011), and Bohren and Hauser (2021). In the next section, we show that bots lead to disagreement among analysts as well.

3.2 Bot aktivity and analyst dispersion

The results so far could be attributed to retail investors, who might be swayed by social media and the artificial increase in popularity caused by bot activity (e.g., Bradley et al., 2023). In fact, professional investors may not even consider the number of likes or retweets, as this signal is largely information-free. One would expect this signal to matter only for relatively unsophisticated market participants. Consistent with this, Foucault et al. (2011) show that speculative retail traders induce volatility. In this section, we explore whether 𝕏 bots affect sophisticated market participants.

Since analysts are often viewed as more professional market participants than retail investors, disagreement among them allows us to test whether bot activity causes disagreement among more sophisticated investors. In fact, Diether et al. (2002) suggest that analyst dispersion indicates disagreement among such investors. Table 5 explores how bot activity impacts future analyst dispersion following earnings announcements.

The results further confirm that the presence of bots significantly affects future analyst dispersion. Namely, an increase in unresolved bugs for detectors means an easier bot entry and is associated with a high rise in future analyst dispersion. Consistent with the results in the previous sections, the coefficients for the unresolved bugs for detectors show higher significance. For example, a 1% increase in CAPTCHA bugs causes a 5.59% increase in analyst dispersion. This indicates that analysts likely consider social media popularity when preparing the reports. Since bot activity is not observable to analysts, it can lead to overestimating the popularity (boosted by bots) and biased estimates in analyst reports. Therefore, sophisticated market participants, such as analysts, are also affected by bot activity.

Table 5. The effect of bot activity on future analyst dispersion

This table shows the impact of bot activity on absolute analyst dispersion for the following quarterly earnings. All variables are described in Appendix A4. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
CAPTCHA bugs	0.049	0.102***		0.006
	(0.031)	(0.039)		(0.042)
SOLVER bugs	-0.329***		-0.406***	-0.404***
	(0.068)		(0.088)	(0.091)
Size		0.008	0.009	0.008
		(0.015)	(0.015)	(0.015)
Book to market		-0.108**	-0.110**	-0.110**
		(0.049)	(0.049)	(0.049)
Dummy fourth quarter		0.039	0.009	0.009
		(0.025)	(0.027)	(0.027)
Illiquidity		3316.543***	3360.974***	3359.810***
		(606.347)	(603.403)	(602.950)
Analyst dispersion		0.011	0.011	0.011
		(0.021)	(0.021)	(0.021)
Analyst number		-0.008	-0.009	-0.009
		(0.038)	(0.038)	(0.038)
Rm−1Rm−1		-0.894***	-0.933***	-0.933***
		(0.264)	(0.268)	(0.268)
R[m−3,m−2] Rm−3,m−2		0.457***	0.499***	0.499***
		(0.142)	(0.147)	(0.147)
R[m−6,m−4]Rm−6,m−4		0.152**	0.176**	0.176**
		(0.077)	(0.080)	(0.081)
R[m−12, m−6]Rm−12, m−6		-0.012	-0.007	-0.007
		(0.045)	(0.044)	(0.044)
Constant	0.563**	-0.502*	1.295***	1.268**
	(0.225)	(0.264)	(0.459)	(0.528)
Year FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
R-squared	0.052	0.080	0.083	0.083
N (observations)	9,047	8,692	8,692	8,692

3.3 Long-term impact of bot aktivity

This section explores whether bot activity can impact volatility and performance over the long term. We consider various time windows from one week after earnings until two months after earnings.[20] Table 6 shows the relation between bot activity, return volatility, and stock performance.

Table 6. The long-term effect on return and volatility

This table shows the impact of CAPTCHA technological shocks on the long-term return and volatility. We report results for four different periods, from one week after earnings to two months after earnings. For each period, we report the results for both volatility and performance. Volatility is defined as return volatility, calculated as the standard deviation of returns in the specified period after the earnings. For performance, we use the cumulative abnormal return (CAR), calculated as the product of abnormal returns for the given period after the earnings, reported in %. All variables are described in Appendix A4. Firm controls are identical to the previous specifications used in Tables 3, 4, and 5. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	1 week		3 weeks		1 month		2 months	
	SD	CAR	SD	CAR	SD	CAR	SD	CAR
CAPTCHA bugs	0.001	-2.065***	0.002**	-3.115***	0.003***	-2.694***	0.005***	0.032
	(0.002)	(0.791)	(0.001)	(0.831)	(0.001)	(0.898)	(0.001)	(1.125)
SOLVER bugs	0.001	1.815***	-0.002**	1.522**	-0.003***	1.698**	-0.007***	4.678***
	(0.001)	(0.651)	(0.001)	(0.747)	(0.001)	(0.784)	(0.001)	(0.942)
Constant	0.043***	1.157	0.027***	8.818	0.020***	9.499*	0.027***	-3.462
	(0.011)	(4.838)	(0.007)	(5.443)	(0.006)	(5.731)	(0.005)	(7.153)
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.250	0.036	0.329	0.042	0.335	0.051	0.405	0.053
N (observations)	9,036	9,036	9,033	9,033	9,033	9,033	9,023	9,023

We use daily data to analyze the long-term impact consistently with previous models. The effect seems to increase and persist even two months after earning announcements. The volatility begins to rise in the weeks following an earnings announcement with high bot activity and further exacerbates during the next two months. The opposite is the case when the bot activity is low. This result further shows that bot activity can have lasting effects and affect market stability. We observe similar patterns using the cumulative abnormal return (CAR) as a variable of interest.[21] An increase in CAPTCHA bugs, which increases bot supply and thus disagreement, is associated with a future drop in performance. Conversely, increased SOLVER bugs are related to increases in performance, likely due to lower volatility. Namely, a 1% increase in CAPTCHA bugs causes return underperformance by 2% in the first week, 2.4% in the third week, and 1.95% in the first month. Similarly, a 1% increase in SOLVER bugs leads to an increase in abnormal returns by 1.8% in the first week, 1.19% in the third week, and 1.24%

in the first month. As such, disagreement amplified by bot activity negatively harms future performance.

4 Robustness tests

4.1 Two-stage least squares and underestimations of bots

This section presents several robustness checks of our results, methodology, and identification. Our analysis relies on exogenous shocks to CAPTCHA technology to estimate bots' impacts on disagreement. While Table 2 shows that shocks to CAPTCHA technology impact bot activity, some concerns and questions regarding firm-specific effects may remain.

We begin by illustrating that CAPTCHA indeed affects investor disagreement through the bot activity measured at the individual firm level. Specifically, we rely on the *BotRatio* to measure the shocks to the firm-specific bot activity. We employ a two-stage least squares (2SLS) approach, using *BotRatio* as our primary independent variable. To account for possible endogeneity arising from shocks to bot availability through CAPTCHA innovations, we use specification (2) and instrument the *BotRatio* using shocks to CAPTCHA technology. More specifically, the outcome of the first stage regression is summarized in Column (1) of Table 2. The regression specification in Column (1) has the F-statistics exceeding 1,200 and R^2 of 0.19. This confirms the validity of CAPTCHA bugs and SOLVER bugs as instruments.[22]

To estimate the firm-level impact of bots on investors' disagreement, we use a modified Equation (3) in which the *BotRatio* measures bot activity on a firm level. Our main regression specification is as follows.

$$DepVar_{i,t} = \alpha_i + \beta_1 \widehat{BotRatio}_{i,t} + Controls_{i,t} + YearFE + IndustryFE + \epsilon_{i,t} \qquad (4)$$

where $DepVar_{i,t}$ denotes one of the dependent variables for which we estimated the effects of investor disagreement associated with bot activities in Equation (3). Using specification (4), we replicate our results from Tables 3, 4, and 5 using the (instrumented) *BotRatio* to measure bot activity around the particular firm's quarterly earnings (QE). We present the results in Table 7.

Table 7. 2SLS with bot ratio

This table replicates main results using Bot ratio as the main independent variable. Here, we repeat the results for Tables 3–5. Column (1) contains the effects on a standard deviation of the returns, Column (2) lists abnormal trading volume, and Column (3) shows future analyst dispersion. All columns contain the 2SLS regression results computed using daily data on bot activity. The first stage of the IV regression is presented in Table 2, Column (1). All variables are described in Appendix A4. All variables and the sample are the same as in Tables 3-5. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	AbLogVol (1)	SD (2)	Analyst dispersion (3)
Bot ratio	2.394***	0.038**	5.676***
	(0.637)	(0.018)	(1.187)
Size	-0.003	-0.003***	0.007
	(0.007)	(0.000)	(0.015)
Book to market	-0.061***	0.001	-0.110**
	(0.019)	(0.001)	(0.049)
Dummy fourth quarter	0.112***	-0.001*	0.020
	(0.013)	(0.000)	(0.025)
Illiquidity	61.065	41.595***	3319.067***
	(135.381)	(4.990)	(603.516)
Analyst dispersion	-0.040***	-0.001**	0.012
	(0.013)	(0.000)	(0.021)
Analyst number	0.040	0.004***	-0.008
	(0.026)	(0.001)	(0.038)
Rm−1Rm−1	0.007	-0.005*	-0.927***
	(0.083)	(0.003)	(0.267)
R[m−3,m−2] Rm−3,m−2	-0.536***	-0.013***	0.473***
	(0.067)	(0.002)	(0.143)
R[m−6,m−4]Rm−6,m−4	-0.112***	-0.003**	0.166**
	(0.039)	(0.001)	(0.079)
R[m−12, m−6]Rm−12, m−6	0.084***	0.000	-0.004
	(0.027)	(0.001)	(0.043)
Constant	1.255***	0.055***	-1.282***
	(0.170)	(0.005)	(0.276)
2SLS	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes
R-squared	0.067	0.252	0.082
N (observations)	9,171	9,162	8,692

We observe that the results using *BotRatio* align with the previous findings, which in this case are extended for firm-specific bot interference. Using a similar approximation for the elasticities used earlier, we observe that increased *BotRatio* leads to significant increases in abnormal volume, volatility of returns, and analyst dispersion. The effect is economically sizable, as for example, a 1% increase in *BotRatio* leads to an increase of 1.4% in abnormal volume and 0.27% volatility of returns.[23] Higher bot activity on the firm's official corporate accounts is associated with higher disagreement among investors and analysts.

A potential concern with our identification is how well 𝕏 detects bots and whether some are left undetected. While 𝕏 does not provide an exact methodology for detecting bots, it specifies that it uses a holistic approach that includes using data that is not publicly available.[24] Current bot detection literature primarily focuses on using publicly available user data and metadata instead and employs one of several possible machine-learning algorithms.[25] According to Cresci et al. (2015), one of the crucial factors in detecting bot accounts is the account's age. More importantly, age indirectly improves bot detection, as the account's higher age allows the bot to follow and react to more content, which increases the friend-follower ratio. This is further supported by other research, e.g., Yang and Menczer (2023), who use the 200 most recent actions by accounts to assess them and conclude that younger accounts might be more difficult to detect.

An important implication of Cresci et al. (2015) is that due to 𝕏 being able to use otherwise inaccessible data and apply its proposed models globally, it should be able to identify suspect accounts and issue a bot challenge if it observes a sufficient amount of account activity. We use the benchmark of 200 tweets, which Yang and Menczer (2023) recommend as a sufficient number of tweets to detect bots. In our sample, the accounts post tweets once per business day on average, yielding a period of around one year necessary for a bot to be detected. While bot accounts are arguably more active than official corporate 𝕏 accounts, and thus a shorter window might suffice, we employ the one-year benchmark. After this period, 𝕏 should identify all bots on its platform. Therefore, for robustness checks, we restrict our sample to observations ending in 2021, which is one year before the original data ends. This sample only contains bot ratios of posts where most bots will be identified. We then repeat our analysis using this sample to allay the concern that some of the bots may be undetected. The results remain robust and are not statistically different from the result of our total sample. We report the results in Table 8.

Table 8. Robustness to underestimation of bots

This table replicates our results from Table 6 using the sample ending in 2021. This allows us to only consider tweets that should have all bot activity fully detected. Column (1) contains the effects on a standard deviation of the returns, Column (2) indicates abnormal trading volume, and Column (3) displays future analyst dispersion. All columns contain the 2SLS regression results computed using daily data on bot activity. The first stage of the IV regression is presented in Table 2, Column (1). All variables are described in Appendix A4. Robust standard errors are in parentheses, and ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	AbLogVol (1)	**SD (2)**	**Analyst dispersion (3)**
Bot ratio	1.805***	0.035**	5.812***
	(0.643)	(0.018)	(1.205)
Size	-0.005	-0.003***	0.009
	(0.007)	(0.000)	(0.016)
Book to market	-0.057***	0.002***	-0.108**
	(0.020)	(0.001)	(0.050)
Dummy fourth quarter	0.103***	-0.000	0.015
	(0.013)	(0.000)	(0.027)
Illiquidity	38.313	39.409***	3380.741***
	(148.213)	(5.477)	(610.560)
Analyst dispersion	-0.034*	-0.000	-0.023
	(0.018)	(0.000)	(0.031)
Analyst number	0.023	0.004***	-0.005
	(0.028)	(0.001)	(0.039)
Rm−1Rm−1	0.013	-0.004	-0.927***
	(0.090)	(0.003)	(0.279)
R[m−3,m−2] Rm−3,m−2	-0.489***	-0.008***	0.453***
	(0.073)	(0.003)	(0.148)
R[m−6,m−4]Rm−6,m−4	-0.119***	-0.002	0.191**
	(0.041)	(0.001)	(0.082)
R[m−12, m−6]Rm−12, m−6	0.105***	0.001	0.003
	(0.028)	(0.001)	(0.045)
Constant	0.157	0.043***	-0.952***
	(0.181)	(0.006)	(0.256)
2SLS	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes
R-squared	0.056	0.251	0.083
N (observations)	8,424	8,416	8,378

4.2 Difference between high and low bot aktivity: ATET

Another potential concern is that due to our 2SLS approach and IV choice, our results could be driven by the global level of bot activity affected by CAPTCHA innovations rather than the distinctive bot activity for each firm. We use propensity score matching and the average treatment effect on treated (ATET) to address this concern. More specifically, we use the Rubin Causal Model (Holland, 1986). The model is based on two possible outcomes: with and without treatment.

$$y_{0i} = \mu_0 + \epsilon_{0i} \; and \; y_{1i} = \mu_1 + \epsilon_{1i} \tag{5}$$

Formally, it can be written as $y_{Ti} = \mu_T + \epsilon_{Ti,}$ where subscript T=1 denotes the treatment, and T=0 represents the control group. We use propensity score matching to estimate the ATET, where y_{0i} is estimated using the nearest-neighbor approach with an extensive set of controls. In our models, we precisely 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 bots' activity above the 75 percentile of the global level, which is defined as the average bot ratio in the 7-day window of all firms apart from the given firm. In comparison, the control group is represented by those with bot activity below the 45 percentile of the global level.[26] Table 9 reports the results of this analysis.

To facilitate better matching, we aggregate the bot activity data by calculating the maximum ratio for both the pre- and post-earning periods.[27] Our results confirm our previous regression model results. The treatment effect is statistically and economically significant, showing that bot activity substantially affects volatility, abnormal volume, and analyst dispersion. These results confirm our earlier results and causality interpretation.[28] The high bot activity (proxied by the difference between high and low *BotRatio*, i.e., the ATET) highly impacts abnormal volume and future analysts' dispersion. In contrast, the effect of bots on a standard deviation of returns remains much smaller.

Table 9: Average treatment effect on treated

This table reports the average treatment effect on treated (ATET) results measuring the impact of excessive bot activity. The treatment group (T=1) is companies with bot activity above the 75 percentile, while the control group is companies with bot activity below the 45 percentile. We employ the nearest neighbor matching procedure *teffect* 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) AbLogVol	**(2) SD (-7,0)**	**(3) Absolute analyst dispersion**	**(4) AbLogVol**	**(5) SD (0,7)**	**(6) Absolute analyst dispersion**
ATET	0.170*** (0.041)	0.004*** (0.001)	0.222*** (0.080)	0.224*** (0.038)	0.002** (0.001)	0.072 (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

5 Conclusion

Bot activity is an affordable and quick way for companies to attain a broad social reach. It provides them with a virtual megaphone to reach investors and artificially increases the number of likes and reposts a company receives on its 𝕏 posts. However, what consequences are incurred by 𝕏 bot activity? Our paper answers this question by studying the impact of bot activity during quarterly earnings announcements.

We document that bot activity significantly increases investor disagreement, which results in higher volatility and trading volume. Furthermore, the increase appears to be long-term, as the effect persists for up to three months following earnings. Moreover, we detected that bot activity could affect analysts since it significantly affects analyst dispersion in the following quarter. Our results are robust to various specifications, including 2SLS and a matching ATET approach indicating causal interpretation.

Since 𝕏 became private in 2022, many tools for studying or analyzing bot activity have been discontinued. For instance, the academic API for 𝕏 data has been discontinued, and academics have been priced out of 𝕏 and bot research.

While our research is limited to the quarterly earnings and constituents of S&P 500 companies due to data availability constraints, both corporations and individuals might employ bots over more extended periods due to their inexpensiveness. Moreover, the effect can be pronounced for smaller companies.

References

Angrist, J. D., Imbens, G. W., & Krueger, A. B. (1999). Jackknife instrumental variables estimation. *Journal of Applied Econometrics*, *14*(1), 57-67.

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. *Journal of Finance*, *59*(3), 1259-1294.

Assenmacher, D., Clever, L., Frischlich, L., Quandt, T., Trautmann, H., & Grimme, C. (2020). Demystifying social bots: On the intelligence of automated social media actors. *Social Media + Society*, *6*(3), 2056305120939264.

Bartov, E., Faurel, L., & Mohanram, P. S. (2018). Can Twitter help predict firm-level earnings and stock returns?. *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).

Bohren, J.A. and Hauser, D.N. (2021). Learning with heterogeneous misspecified models: Characterization and robustness. *Econometrica*, 89(6), pp.3025-3077.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.

Bradley, D., Hanousek Jr, J., Jame, R., & Xiao, Z. (2023). Place Your Bets? The Value of Investment Research on Reddit's Wallstreetbets. *Review of Financial Studies*, hhad098.

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.

Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., ... & Amodei, D. (2018). The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. *arXiv preprint arXiv:1802.07228*.

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

Conti, A., Peukert, C., & Roche, M. P. (2024). Beefing IT up for your Investor? Engagement with Open Source Communities, Innovation and Startup Funding: Evidence from GitHub. *Organization Science*, (ja).

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

Cookson, J. A., Engelberg, J. E., & Mullins, W. (2023). Echo chambers. The Review of Financial Studies, 36(2), 450-500.

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. *Journal of Finance*, *57*(5), 2113-2141.

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 www.edition.cnn.com/2022/10/10/tech/elon-musk-twitter-bot-analysis-cyabra/index.html

Fama, E. F. (1998). Market efficiency, long-term returns, and behavioral finance. *Journal of Financial Economics*, *49*(3), 283-306.

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*.

Foucault, T., Sraer, D., and Thesmar. D. (2011). Individual investors and volatility. *Journal of Finance* 66(4): 1369-1406.

Hanley, H. W., & Durumeric, Z. (2023). Machine-Made Media: Monitoring the Mobilization of Machine-Generated Articles on Misinformation and Mainstream News Websites. *arXiv preprint arXiv:2305.09820*.

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

Hirshleifer, D., Peng, L., & Wang, Q. (2024). News diffusion in social networks and stock market reactions. The Review of Financial Studies.

Hoffmann, M., Nagle, F., & Zhou, Y. (2024). The value of open source software. Harvard Business School Strategy Unit Working Paper, (24-038).

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

Keller, F. B., Schoch, D., Stier, S., & Yang, J. (2020). Political astroturfing on twitter: How to coordinate a disinformation campaign. *Political Communication*, *37*(2), 256-280.

Mitchell, M. L., & Stafford, E. (2000). Managerial decisions and long-term stock price performance. *Journal of Business*, *73*(3), 287-329.

Rabin, M. and Schrag, J.L. (1999). First impressions matter: A model of confirmatory bias. *Quarterly Journal of Economics*, 114(1),.37-82.

Spitale, G., Biller-Andorno, N., & Germani, F. (2023). AI model GPT-3 (dis) informs us better than humans. *arXiv preprint arXiv:2301.11924*.

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 www.washingtonpost.com/technology/2018/07/06/twitter-is-sweeping-out-fake-accounts-like-never-before-putting-user-growth-risk/

Van den Steen, E. (2011). Overconfidence by Bayesian-rational agents. *Management science*, *57*(5), 884-896.

Weidinger, L., Uesato, J., Rauh, M., Griffin, C., Huang, P. S., Mellor, J., ... & Gabriel, I. (2022, June). Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 214-229).

Wright, N. L., Nagle, F., & Greenstein, S. (2023). Open source software and global entrepreneurship. *Research Policy*, 52(9), 104846.

Yamin, M. M., Ullah, M., Ullah, H., & Katt, B. (2021). Weaponized AI for cyber attacks. *Journal of Information Security and Applications*, *57*, 102722.

Yang, K. C., Varol, O., Davis, C. A., Ferrara, E., Flammini, A., & Menczer, F. (2019). Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emerging Technologies*, *1*(1), 48-61.

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

Yang, K. C., & Menczer, F. (2023). Anatomy of an AI-powered malicious social botnet. *arXiv preprint arXiv:2307.16336*.

Appendix

Appendix A1. Example of GitHub repository

Below is an example of the open-source GitHub repository that contains a code to solve CAPTCHA and thus facilitate bot activity. It uses AI to solve CAPTCHAs that blocks bots from accessing websites.

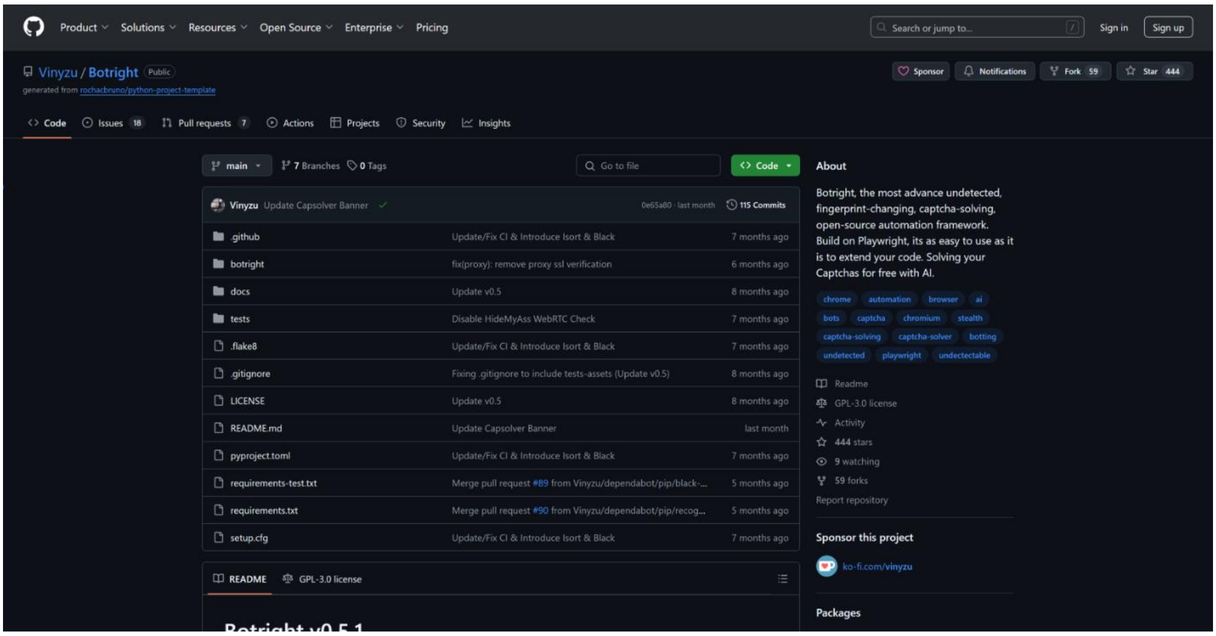

Appendix A2. GitHub data collection and shocks to CAPTCHA technology

To identify shocks to CAPTCHA technology, we collect information from GitHub. We focus on all repositories that have the word "CAPTCHA" as one of the keywords, which gives us 21,475 repositories. We further restrict the sample to those that have non-missing descriptions. We then use textual analysis to categorize repositories as those focusing on CAPTCHA implementation and those focusing on solving CAPTCHA. Specifically, we look for specific keywords. For CAPTCHA implementing repositories, we look for the words "generate," "detect," "implement," and "integrate" while controlling for various forms of the word. This results in 1,595 repositories for CAPTCHA implementation. For CAPTCHA solving repositories, we look for the words "block," "fill," "solve," "auto," "bypass," and "break," again controlling for various forms of the word. We manually classify those that contain both types of keywords. This leaves us with 2,642 CAPTCHA solving repositories. We only leave repositories that fit one of these two categories.

We focus on issues (i.e., bugs) reported for each repository. First, we remove bugs that were marked closed by the creator of the repository after less than 24 hours, as those are frequently either questions or do not constitute a serious issue. Next, we omit bugs that were raised or answered by a bot. GitHub allows users to call for automatic updates to the libraries used by developers by filing a bug. These are omitted, as they do not constitute an actual issue but rather a small, unrelated update. Lastly, we remove bugs that are never solved. These come from repositories that are no longer updated and do not accurately portray the technological advances and shocks to CAPTCHA technology.

Appendix A3. 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 the complete list of accounts from which the company liked or reposted a tweet. As suspended users have missing information, we can calculate the ratios of bots versus normal users. The official 𝕏 guidelines support this approach,[29] as they 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 1 error, meaning that suspended users are highly likely to be bots. That is because 𝕏 does not rely only on language analysis, as considers it too restrictive.[31] Moreover, established bot detection techniques are not able to detect AI-controlled bots or AI-generated text.[32] Instead, 𝕏 uses more comprehensive analysis, and when it detects that an account might be engaged in manipulative behavior, it will send out a bot challenge[33] 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. Since human users can easily complete any challenge, this approach should yield the lowest type 1 error.

A4. Variable definitions

Variable	Description
Dependent variables	
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.
Volatility	Volatility is defined as the standard deviation of daily returns. In our main specification, for days prior to quarterly earnings, we define volatility as the standard deviation of daily returns for the period of (-7,0), and for days post quarterly earnings, we define it as the standard deviation of daily returns for the period of (0,7). This approach was chosen for better variation of the dependent variable and to identify potential different sources of disagreement outlined in Cookson and Niessner (2020). In the section on long term effects, volatility is defined as the standard deviation of returns for the given period after quarterly earnings. Data source: CRSP.
Absolute analyst 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. Data source: I/B/E/S.
Cumulative abnormal return (CAR)	The cumulative abnormal return is defined as $CAR_{i,p} = \prod_{t=1\ to\ p}\left(1 + r_{i,t} - r_{mkt,t}\right) - 1,$ where $r_{i,t}$ is the return of the stock i for day t, $r_{mkt,t}$ is the return of the CRSP value-weighted index for day t, and p corresponds to the different periods from 1 week to 2 months. Data source: CRSP.
Bot activity measures	
Bot ratio	The bot ratio is defined as the number of likes by suspended accounts that liked any tweet from the corporation on the given day divided by the total number of likes of a tweet from the corporation on the given day. Data source: 𝕏.
CAPTCHA bugs	The number of unresolved bugs in GitHub repositories focusing on CAPTCHA implementing on the given day. We use log(X) transformation to account for the skewness. Higher number of CAPTCHA bugs implies less restricted bot activity. Data source: GitHub.

SOLVER bugs	The number of unresolved bugs in GitHub repositories focusing on CAPTCHA solving on the given day. We use log(X) transformation to account for the skewness. Higher number of SOLVER bugs implies more restricted bot activity. Data source: GitHub.
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 $R_{m-1}, R_{[m-3,m-2]}, R_{[m-6,m-4]}$, and $R_{[m-12,\ m-6]}$, which stand for returns in percent 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. Data 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. Data source: I/B/E/S.
Fourth quarter indicator	The indicator is equal to one if the given quarter is the fourth fiscal quarter. Data source: I/B/E/S.