Algorithmic Agents in the Hybrid Media System: Social Bots, Selective Amplification, and Partisan News about COVID-19

Zening Duan [1], Jianing Li [2], Josephine Lukito[3], Kai-Cheng Yang[4], Fan Chen [5], Dhavan V. Shah[1], and Sijia Yang[1]

[1]School of Journalism and Mass Communication, University of Wisconsin-Madison, Madison, WI, USA
[2]Department of Communication, University of South Florida, Tampa, FL, USA
[3]School of Journalism and Media, University of Texas at Austin, Austin, TX, USA
[4]Luddy School of Informatics, Computing, and Engineering, Indiana University Bloomington, Bloomington, IN, USA
[5]Google LLC, New York, NY, USA

Social bots, or algorithmic agents that amplify certain viewpoints and interact with selected actors on social media, may influence online discussion, news attention, or even public opinion through coordinated action. Previous research has documented the presence of bot activities and developed detection algorithms. Yet, how social bots influence attention dynamics of the hybrid media system remains understudied. Leveraging a large collection of both tweets (N = 1,657,551) and news stories (N = 50,356) about the early COVID-19 pandemic, we employed bot detection techniques, structural topic modeling, and time series analysis to characterize the temporal associations between the topics Twitter bots tend to amplify and subsequent news coverage across the partisan spectrum. We found that bots represented 8.98% of total accounts, selectively promoted certain topics and predicted coverage aligned with partisan narratives. Our macro-level longitudinal description highlights the role of bots as algorithmic communicators and invites future research to explain micro-level causal mechanisms.

Open Science Framework awards

Open Materials
The components of the research methodology needed to reproduce the reported procedure and analysis are publicly available for this article.

Open Data
Digitally shareable data necessary to reproduce the reported results are publicly available for this article.

Keywords: Social Bots, The Covid-19 Pandemic, Hybrid Media System, Issue Polarization, Social Media, Attention Dynamics, Structural Topic Modeling, Time-Series Analysis

Corresponding author: Sijia Yang; email: syang84@wisc.edu

https://doi.org/:10.1093/hcr/hqac012

Social media have become a major channel in which to disseminate news, discuss politics, and engage in collective action (Keller et al., 2020). As artificial intelligence (AI) technologies increasingly permeate into the modern information ecology—in forms of virtual agents, chat bots, social bots, and many others (Guzman & Lewis, 2020)—the hybrid media system today not only consists of human actors such as journalists, politicians, and social media users (Chadwick, 2013), but also nonhuman, algorithmic communicators. Among them, social bots, algorithmic agents that amplify certain viewpoints and interact with selected actors on social media, often try to pose as human actors (Boshmaf et al., 2013). In some contexts, these algorithmic actors account for about two-thirds of tweeted links to popular websites (Wojcik et al., 2018). These algorithmic amplifiers and interaction agents can be strategically coordinated and act in an orchestrated manner (Meta, 2018).

Existing research has made impressive progress in describing the patterns of bot activities, locating their origin, and developing algorithms for bot detection (Sayyadiharikandeh et al., 2020). With that said, systematic research into the dynamic relationships among bots, news media, and the public remains scarce. The growing influence of bots warrants more scholarly efforts to better understand such algorithmic communicators' roles in directing online information flows and enriching the notion of hybridity characteristics in today's media system. Two questions loom large: (1) What topics within a complex issue do bots amplify and ignore? (2) When bots do amplify topics, what parts of the media ecology respond to this prompting? Particularly concerning is the growing body of evidence that reveals social bots can amplify malicious disinformation spreaders' capacity to influence online information and discourses (Ferrara, et al., 2016), often by promulgating sensational, misleading, and low-quality content across a broad spectrum of social issues ranging from presidential elections (Keller & Klinger, 2019) and vaccine debates (Broniatowski et al., 2018) to the #BlackLivesMatter movement (Stewart et al., 2018) and the ongoing COVID-19 pandemic (Ferrara, 2020).

It is also likely that the activity of social bots affects journalistic practices, given reporters' attentiveness to Twitter activity and interactions (Wells et al., 2020). While longstanding journalistic norms value verification and triangulation of sources, tweets are often used as "interchangeable building blocks" in a story, rather than sources whose words need to be verified (Molyneux & McGregor, 2021). Moreover, metrics of activity signal audience interest and tend to spur further coverage in an attention loop, especially among outlets concerned with economic viability (Vu, 2014). Despite editors insisting they guard against using audience metrics in news selection, research has found the most-viewed storylines received greater attention in those same outlets (Welbers et al., 2016; cf. Zamith, 2018). Online comments, popularity metrics, and forms of social endorsements have also been tested on audience responses as well, with outcomes pointing toward reinforcement of selective exposure tendencies and perceptions of information quality (Messing & Westwood, 2014; Chung, 2017, cf. Waddell, 2018). As such, posts from Twitter bots, especially those from authentic-seeming accounts, can display cultural competence, or coordinate to amplify particular topics can be picked up by news media professionals and media audiences as a sign of public endorsement. Even simple amplification can be seen as a cue of traction and legitimacy (Wells et al., 2020).

Research on AI agents operating within the hybrid media system needs to recognize the role of these automated communicators and how they try to operate within and exert influence throughout the communication ecology. That said, researchers disagree over the degree of intelligence the current generation of social bots in fact possess. Some researchers found that the majority of available bot services and software merely provide simple and repetitive automation, although they also admit that more sophisticated bots do exist (Assenmacher et al., 2020). Given the rapidly evolving nature of natural language models (e.g., GPT-3) and deep learning technologies, social bots may soon become more human-like and intelligent than the current generation. Hence, reflecting the broad spectrum of AI technologies, social bots range from simpler, low-intelligent bots who may follow some basic rules to conduct actions such as repeatedly mentioning certain accounts (Ratkiewicz et al., 2011) to highly sophisticated ones who adopt advanced deep neural networks to generate more credible content and profile images (Cresci, 2020). Studying even a subset of social bots that operate along this continuum, as we do in this study, can provide important insights about these coordinated operation within the broader hybrid media system.

Like human or organizational actors in the hybrid media system, social bots stand as an active player in the attention economy (Webster, 2014). Bots can execute repetitive tasks in bulk and with impressive efficiency (Grimme et al., 2017), such as engaging with or even hijacking hashtags, following and retweeting influencers, and algorithmically generating original content. Their capacities to influence issue attention can be used for ill or for good (Edwards et al., 2014). They may pose a threat to the online information ecosystem and civic discourses (Ferrara et al., 2016) by promulgating conspiracy theories and misinformation (Starbird, 2019), distorting civic information (Alothali et al., 2018), spreading low-credibility content (Shao et al., 2018), and exacerbating polarization (Bail et al., 2018). Meanwhile, bots can also facilitate or distort public health campaigns by disseminating personalized health prevention messages (Jamison et al., 2019). Debates about their impacts on information flows and attention dynamics remain, with some arguing that bots are less important than verified accounts in disseminating information about contentious political issues on Twitter (González-Bailón & Domenico, 2021). Accordingly, we adopt a hybrid media system perspective and aim to uncover social bots' dynamic interactions with news organizations as well as with human actors against the backdrop of multifaceted and multilayered messaging about the early phases of the COVID-19 pandemic.

Leveraging a large-scale collection of COVID-19 tweets ($N = 1{,}657{,}551$) and related news stories ($N = 50{,}356$) during the early phase of the pandemic (March to May 2020), we combined bot detection, structural topic modeling, and time-series analysis to examine how bots selectively amplified or downplayed certain topics over others in comparisons with humans and news media across the political spectrum. Our results reveal that social bots strategically amplified topic salience in human-generated content more so than in news coverages. Further, we uncovered evidence that such bot amplification then predicted news content in a manner that feeds the ideology of the partisan media outlets. Our findings suggest social bots are a scalable, active, and important set of algorithmic communicators in the hybrid media system.

Social Bots in the Hybrid Media System

Social bots' entry into the information ecology and their complex interactions with other human communicators suggest that social bots both reflect and contribute to the hybridity of today's media system. As argued by Chadwick (2013), the notion of hybridity rejects a binary perspective that delineates the boundary between "old" and "new" media, but instead emphasizes the enmeshment and interdependence between *older* and *newer* media logic—both in the sense of the dynamic interactions between different types of communicators and the co-existence and co-evolvement of the content they each produce. Viewed from this hybrid media system perspective, the current news-making process involves a more diverse set of actors and interactions than before (Chadwick, 2013; Wells et al., 2020). In this increasingly competitive media space, the public is not only influenced by human and organizational actors such as governments (King et al., 2017), low-credibility media (Shao et al., 2018), and political extremist groups (Ferrara, 2020), but also non-human, algorithmic actors such as bots.

There is a broad spectrum of bots on social media platforms varying in the level of their capacity to mimic human behaviors and carry out human-like conversations. In the current study, we focus on bots on the lower end of this spectrum for two reasons. First, less sophisticated social bots are more detectable and may have discernable relationships with humans operating within the media ecosystem. A better understanding of how basic automated agents operate can set the groundwork for estimating how more sophisticated AI-powered bots may shape human communication. Second, evidence shows that even these less sophisticated algorithmic agents can be powerful, especially when individual social bot accounts are strategically coordinated to form a botnet (Meta, 2018), acting in an orchestrated fashion to amplify certain voices and narratives (Pacheco et al., 2020). Experimental evidence shows that these less sophisticated bots can still exert considerable influence, like those increasing political polarization by sharing cross-cutting media content (Bail et al., 2018). Given the paucity of research on this topic, our study demonstrates the feasibility and importance of combining large-scale computational analyses and time-series modeling to start determining the topics bots selectively amplify and the parts of the media ecology appearing to respond to coordinated activities by algorithmic communicators.

Feeding into the economy of attention in the hybrid media system, bots are often created with purposes to, for instance, attract attention and inflate popularity. They primarily influence the media system through the process of topic amplification, defined as bots' contribution of attention paid to particular topics, elevating other actors' (citizens, journalists, news outlets) perceptions of the topic's worthiness or significance. Topic amplification happens through selectively re-posting, retransmitting, and sometimes remixing certain voices of human actors such as political elites, journalists, influencers, and the public—and given that attention is a scarce public resource on social media, amplification of one topic also implies diversion from other important issues. Moreover, bots are efficient in identifying users or content already on track to virality, and they can be programmed to monitor user engagement metrics (e.g., counts of likes, shares) in real time and on a large scale. Given that the online attention economy is built on the bedrock of such measurables, indicative of rapidly evolving concentration and transition of public attention, metric-savvy bots in some sense have the late-mover advantage when competing for public attention with other media outlets.

Bot Actors and Partisan Media Bias

COVID-19 spread across the world in early 2020 and soon came to dominate the public conversation in the United States. In the early phase of the pandemic, the public's understanding of COVID-19 was chaotic, with divisive, and unfortunately, politicized assessments of the spread and severity of disease, as well as mitigation recommendations and policies. During the same period, we witnessed debates among news media outlets on a broad range of topics involving changes in daily life (Zhang, 2021), virus origins (Chen, Chen, Zhang, et al., 2020), prevention strategies (e.g., wear a mask, Sanders et al., 2021), and others. When one fell, another rose.

To illustrate the importance of social bots in the hybrid media system, we carried out a time-series analysis on how these automated agents can predict news organizations' attention allocation to a broad spectrum of issues related to the COVID-19 pandemic. Further, we compared these relationships across media outlets that span the partisan spectrum. Fueling the political divide around the pandemic, partisan media in the United States have demonstrated stark differences in their COVID-19 coverage. At the early stage of the pandemic, liberal media not only covered COVID-19 more attentively than conservative media, but also put more emphasis on the urgency of responding to the pandemic, stressing the science behind prevention strategies. A computational text analysis of partisan media coverage showed that liberal media most frequently mentioned Anthony Fauci, Donald Trump, lockdown, and social distancing, while conservative media usually talked about COVID-19 through a more foreign lens, focused on the Diamond Princess Cruise, Iran, and Korea (Bermejo et al, 2020).

Echoing Trump's downplaying of the pandemic, conservative media have minimized the severity of COVID-19 and propagated claims that the virus was lab-designed or that a cure already existed (Motta et al., 2020). While mainstream and liberal media have been acutely stressing the risks of the COVID-19 public health crisis, their coverage of COVID-19 was still often casted through a politicized and polarized lens, as seen in the frequent appearance of partisan politicians in news stories and distinct language patterns in coverage mentioning Republican and Democratic politicians (Hart et al., 2020).

Partisan media on the political left and right, aiming to advance their own issue agendas, may be distinctively influenced by Twitter bots. A growing line of literature has documented the routinization of tweets in the newsroom. Tweets can influence the editorial decisions on what to cover; journalists who frequently use Twitter see headlines shared in anonymous tweets as equally newsworthy as an Associated Press headline (McGregor & Molyneux, 2020). Journalists also take cues from the degree to which a message is retweeted or liked, seeing it as a signal of public sentiment or audience interest (Wells et al., 2020; Zhang et al., 2018). While longstanding journalistic norms value verification, tweets, retweets, and metrics of interaction volume are often used as cues for story writing (Molyneux & McGregor, 2021), influencing coverage decisions or even becoming part of a story (Xia et al., 2019). As partisan media are mainly concerned with offering an interpretive package for their audience, they are likely to selectively choose content from Twitter that fits their ideological lens and issue agenda (Levendusky, 2013).

How Bots Elevate Human Attention

Social bots may intentionally or inadvertently affect attention dynamics of human actors, often targeting a specific event or online group. An emerging body of research has found that social bots engaged in online discourse concerning national elections (Howard et al., 2018), social movement organizing (Stewart et al., 2018), and public health disputes (Broniatowski et al., 2018). Bots not only follow certain events, but also selectively target specific groups of strategic significance—for example, future voters to sway election results (Santini et al.,

2021) or online influencers to promote misinformation (Shao et al., 2018). Even bot activities that are not designed to target any specific issue or group, but rather to achieve goals such as maximizing visibility, may unexpectedly drive interaction on contentious issues.

No matter if social bots are intentionally or inadvertently promoting certain issues, they are often programmed to deploy three common strategies to attract attention from human actors on social media platforms. First, some bots automatically generate content through running natural language generation modules (Hill et al., 2015). Because people often rely on keywords to search for content of interest on social media, bot-generated posts may squeeze out authentic posts and incidentally expose people to artificial or artificially amplified content. A more problematic scenario occurs when bots attempt to hijack trending hashtags and keywords. A well-known case occurred in Syria during the Arab Spring: bots were employed to flood hashtags (e.g., #Syria) related to the Syrian civil war with irrelevant or opposing narratives to redirect the attention of users from criticizing the government (Alothali et al., 2018).

The second strategy commonly adopted by bots to attract attention is by directly interacting with, replying to, retweeting, mentioning, or quoting targeted human actors of importance, such as influencers occupying central positions in the network. Even bots located on periphery positions of the online network gain visibility by interacting with central influential users (Shao et al., 2018). Lastly, bots may grow, elevate, or downplay their own network of followers and wield increasing power as algorithmic influencers in shaping issue salience. Social bots are, in these ways, deeply integrated into the networks of human communication (Ferrara et al., 2016). This human–bot network integration enables bots to amplify posts and push information into the timelines of human users, winning attention.

Studies find that although social bots are not always "followed" or retweeted by human users, they constitute a non-trivial portion in online discussion of contentious topics such as presidential elections (Bessi & Ferrara, 2016) and childhood vaccinations (Broniatowski et al., 2018). An emerging line of scholarship has only begun to compare the keywords (Keller et al., 2020) and sentiments (Bessi & Ferrara, 2016) posted by human-like and bot-like accounts. From this work, we have learned bots are much more likely to disseminate low-credibility information (Shao et al., 2018), negative valence content (Broniatowski et al., 2018), and polarizing political messages (Stella et al., 2018). Yet, we have little knowledge on how attention to topics transfers among different actors in the communication ecosystem, moving among human users, news media, and social bots, nor do we know their overtime relationships.

Given the strategies summarized above, social bots may affect attention dynamics either intentionally or inadvertently. Social bots seem likely to amplify the voices of human users around the COVID-19 pandemic, especially in disseminating information of broader interest, sharing sensationist content, and highlighting contentious topics with high degree of partisan polarization. This suggests that topic prevalence in human-generated tweets is also likely to predict, selectively, subsequent topic prevalence in bot-generated tweets. As social bots evolve to adopt more automated text generation technologies (Dathathri et al., 2020; Zhang, Sun, Galley, et al., 2019), this dependence on human discourse is likely to diminish, moving bots from amplification engines toward autonomous communicators.

Of course, social bots, especially if coordinated, may be able to sway news attention, especially by ideologically aligned outlets that seek out the "building blocks" for partisan narratives and have metrics-driven imperatives (Molyneux & McGregor, 2021; Welbers et al., 2016). Likewise, various forms of online endorsements have been found to shape judgments of broader audiences, reinforcing selective exposure tendencies and information quality perceptions (Messing & Westwood, 2014; Chung, 2017). Nonetheless, given the limited work on such bi-directional relationships—i.e., the particular human and news content social bots tend to amplify *and* how that amplification shapes downstream human discourse and news coverage—we pose a series of interrelated research questions that motivate this research on social bots, general Twitter users, and news coverage:

RQ1: What is the volume of information generated by bot-like accounts in the observed Twitter discourse on the COVID-19 outbreak?

RQ2: What are the distinctive topical foci of bot-like accounts, compared to (a) human accounts (b) and partisan and mainstream news media coverage?

RQ3: Through investigating the bi-directional overtime relationship between topics emphasized by human accounts and Twitter bots: (a) which human discourse topics are amplified by Twitter bots and (b) how does that amplification shape downstream human discourse?

RQ4: Through investigating the bi-directional overtime relationship between topics emphasized in news coverage and Twitter bots: (a) which news content topics are amplified by Twitter bots and (b) how does that amplification shape downstream news coverage?

Methods

Data Collection

Our main dataset consists of (1) about 0.01% sample of all COVID-19-related tweets ($N = 1{,}657{,}551$) matched up with our keywords list containing 181 n-grams spanning three months from March 1st to May 31st, 2020, and (2) a parallel corpus of COVID-19 news stories ($N = 50{,}356$) from the same time span. Third-party commercial data vendor Synthesio was used to gather the corpus of COVID-19 tweets (Chen et al., 2021). This three-month span started one day after the first U.S. death was reported and ended approximately one week after the death of George Floyd, a point after which the public's attention began shifting to other issues (Hart et al., 2020). We built a comprehensive list of English keywords based on expert knowledge, systematic web search, and similar keyword lists shared by Twitter (Twitter, 2020) and researchers (Chen, Lerman, & Ferrara, 2020). We followed suggestions by Kim et al. (2016) to screen and iteratively update the keyword list based on relevance and frequency. In total, 181 keywords were used for retrieval and listed in Table A1, Online Appendix A. To support open science and facilitate replication, we released our R/Python codes, analytical datasets including news story URLs and tweet IDs, and other supplementary materials to an online folder housed by the Open Science Framework, see https://doi.org/10.17605/OSF.IO/2C65S.

When searching for relevant news stories, we modified the keywords list, which can be found in Table A2, Online Appendix A, by removing digispeaks (e.g., "flatteningthecurve"), internet slangs (e.g., "trumppandemic") and misspellings (e.g., "corono virus"). These terms were unlikely to be adopted by journalists in formal news reporting. We collected matched news stories from Media Cloud, an open-source platform collecting the URLs of millions of online news stories published in multiple languages. Six news outlets were selected to represent the media landscape in the United States, with two conservative outlets (*Fox News* and *Breitbart*), two center-left outlets *(The New York Times [NYT]* and *The Washington Post [WP]*), and two liberal outlets (*MSNBC* and *HuffPost*). Two leading authors manually verified the URL domains and parsed qualified URLs to retrieve full content and metadata using a web crawler tool called Zyte. In total, we successfully collected full data of all the qualified URLs returned by Media Cloud. For more details in URL processing, please see Table B1, Appendix B.

Bot Detection

To examine the roles of social bots, we needed to first identify bot accounts. Since it is infeasible to manually annotate millions of Twitter accounts, we used Botometer v4 (botometer.org) (Sayyadiharikandeh et al., 2020), a supervised machine learning-based system, for bot detection. Botometer v4 examines over 1,000 features from the profile and activity of a Twitter account to identify patterns indicating automated behaviors and produces a bot score ranging from 0 to 1. Accounts with higher bot scores are more bot-like.

We needed to determine a threshold for bot detection. Previously, different thresholds such as 0.5 (Sayyadiharikandeh et al., 2020; Yan et al., 2021) and 0.7 (Grinberg et al., 2019) were adopted. A higher threshold can reduce false positive errors (i.e., human accounts mistakenly labeled as bots) but identifies fewer true bots, while a lower threshold yields more identified bots but introduces more false positives simultaneously. Since our study focuses on estimating the activities of social bots, we prioritize the need to reduce false positives, even though setting a higher threshold would return fewer bot accounts and sacrifice statistical power. To balance these considerations, we followed Grinberg and colleagues (2019) and chose 0.7 as our threshold.[1] Note that this choice was likely to have introduced more false negative errors (i.e., bot accounts mistakenly labeled as humans) in comparison to using 0.5 as the threshold. Therefore, we also carried out robustness checks using 0.5 as the threshold value. Summaries of the largely consistent findings are reported in the Results section and detailed in Online Appendix F. We also observed evidence of coordination patterns in our bot datasets (see Online Appendix C).

Structural Topic Modeling

The next step was to extract the embedded topics from tweets and news stories. We applied Structural Topic Model (STM) as it incorporates document-level covariates in modeling. Document-level covariates are allowed to affect topical prevalence, topical content, or both and are shown to improve inference and qualitative interpretability (Roberts et al., 2019). This is critical for our purposes because we need to incorporate author types (e.g., bots, news outlets, human Twitter users) and time stamps in estimating topic prevalence and changes over time.

We merged tweets (all types of tweets including retweet and reply) and news data into a single corpus to train the STM model. Two covariates, the date when a document (a tweet or a news story) was published and author type (i.e., human Twitter users, Twitter bots, or one of the six media outlets) were incorporated as independent variables. To preprocess the text, we removed stop words, punctuations, URL links, and applied stemming. To make our results more robust, we identified a list of Twitter accounts ($n = 646$; e.g., @dailykos, @washingtonpost, @foxnews; see Table B2, Appendix B for the full list) belonging to influential news outlets and removed them from our analyses. This was to minimize the possibility that observed relationships were driven by news entities' own social media posts.

We adopted a data-driven iterative approach to select the proper number of topics (Roberts et al., 2019). A set of model diagnostics (e.g., held-out likelihood, semantic coherence, and residual) jointly suggested that an appropriate number of topics in our corpus would fall into the range between 40 and 60. We compared models and chose $K = 50$ as the number of topics based on these performance criteria. Subsequent validation and categorization of these topics, including removal of "noisy" or irrelevant content, is detailed in the Results section. The details concerning the set-up and estimation of structural topic model are in Figure D1, Appendix D. The STM results provided estimated average differences between author pairs (e.g., bots vs. human Twitter users, bots vs. each of the news outlets) across the time span of this study along with 95% confidence intervals. These estimates were interpreted to answer RQ2a and RQ2b.

Time-Series Analysis

To construct the time series datasets for a given topic, a given actor category, and a given day, we first retrieved the estimated topic proportions from the STM model output for each of the documents posted by actors of that category (e.g., all bot accounts). This document-level topic proportion amounts to the estimated posterior probability of observing this topic given the document. This numeric estimate ranges between 0 and 1 and is often interpreted as document-level topic prevalence. Next, we aggregated all the raw topic proportions for all the documents produced by actors in a given category (e.g., bots) on a given day and denoted this aggregated daily topic prevalence estimate θ. Lastly, we calculated a normalized daily topic prevalence estimate φ to address variations in the total volume of documents between actor types, using equation (1) below, where θmin and θmax referred to the minimum and maximum aggregated daily topic prevalence estimate for a given actor category over the 92-day period:

$$\varphi = \frac{\theta - \theta min}{\theta max - \theta min} \quad (1)$$

For our time series models, the unit of analysis was one day. Since information in the media system about the COVID-19 pandemic evolved rapidly over time, a granular unit of analysis is better suited to capture the nuanced dynamics of social media activities. Furthermore, day-level analyses are common to studies on news media coverage given the news production cycle (Wells et al., 2019). Although we could perform even more granular analyses (e.g., hour-by-hour) of our social media data, day-level is more appropriate when we need to merge social media activities with news coverage, which often operates on a 24-hour cycle.

Using the bot detection results at 0.7 threshold, we constructed five time series for each topic: Twitter bots, human Twitter users, conservative media (i.e., *Fox News* and *Breitbart*), liberal media (i.e., *MSNBC* and *HuffPost*), and centrist media (i.e., *NYT* and *WP*). All the time series spanned from March 1 to May 31. In total, there were 92 data points in each series.

Following this data wrangling, we then built vector autoregression (VAR) models to estimate the relationships between these five different time series for each topic. VAR models can incorporate lagged values of past observations and help establish the temporal order among multiple time series (Box-Steffensmeier et al., 2014; Wells et al., 2019). While VAR models cannot determine causality, they are one of the best available tools for time series prediction. These analyses were replicated with bot detection results at 0.5 threshold as a robustness check.

Results

Our analyses are three-fold: first, we present descriptive statistics of bot activities as well as suggestive evidence of bot coordination to answer RQ1; second, we summarize prevalent topics embedded in tweets and media coverage, and more importantly, we highlight the distinctive topical focus of Twitter bots as compared with other actors in the hybrid media system to answer RQ2a and RQ2b; lastly, we provide empirical evidence to quantify the degree to which the salience of a broad set of topics amplified by Twitter bots potentially predicted, and were predicted by, online human discourses and news outlets, respectively, from time series analyses (e.g., VAR models). This addresses RQ3a, RQ3b, RQ4a, and RQ4b. We assess the robustness of our main findings with regards to the threshold value for bot detection; the key findings were largely replicated. Details about the robustness check can be found in Online Appendix I.

Bot Activity Patterns

In general, bot accounts occupied a relatively small proportion of the COVID-19 tweet dataset ($N = 711{,}205$ unique accounts): Botometer labeled 63,843 accounts (8.98%) as Twitter bots at the bot score threshold of 0.7. During the three-month period, on average, a human Twitter user in our dataset posted two COVID-related tweets (*Mean* =1.99, *Median* = 1.00, *SD* = 3.47), while a bot account posted nearly six (*Mean* = 5.85, *Median* = 2.00, *SD* = 10.41). Twitter allows three types of tweets: Original tweet (including original posts and replies), retweet (retransmitting an existing tweet from another account without any comment), and quote (retweets with comment). Our results revealed that human Twitter users tended to post slightly more original content (20.5% vs. 19.8%, human vs. social bot, same below), fewer retweets (69.4% vs. 75.4%), and a higher volume of comments (10.1% vs. 4.8%). This pattern reflects bots' tendency towards retweeting rather than generating original content as the latter requires more technical sophistication, though the amount of original content sharing does reflect some level of sophistication in the types of bots detected using Botometer.

Topical Focus of Twitter Bots

STM was used to identify the prevalent topics in texts. Topics were represented as a multinomial distribution over words and their subjective meanings can be interpreted by manually examining groups of words with the highest topic loadings as well as semantic differentiation from other topics (e.g., the FREX scores). STM also made it easy to read example documents with the highest estimated proportions of the focal topic. Four authors, including the first author, independently reviewed and labeled all 50 topics, including whether they were irrelevant to COVID-19, not in English and therefore excluded, or noisy to the point of being uninterpretable. All disagreements were resolved collectively, with 17 topics discarded and 33 topics retained for further analyses. These retained topics were grouped into three clusters to facilitate presentation: Daily Life, Political/Societal, and Public Health (see Table 2). We labeled topics based on their corresponding top tokens and high-ranked sample documents (both tweets and news articles). This was a human-in-loop process requiring qualitative understanding and subjective interpretation. To make this process transparent, the full list of topics along with corresponding tokens is in Figure D2, Appendix D and sample documents is in Appendix E.

Model outputs revealed distinctive topical focus between bots and other non-bot actors. As shown in Figure 1, a total of 33 named topics were sequentially presented and ranked by most significant differences in topic prevalence between Twitter bots and human Twitter users. Figure 1 also illustrates the comparisons between bot discourse with one liberal, one centrist, and one conservative news outlet, respectively. Full results can be found in Appendix D. We observed that Twitter bots emphasized political/societal topics over other topical themes. For instance, a broad range of political/societal topics such as *Blaming China* (T3; $b = -0.022$, $p < .001$), *Trump Response Failure* (T4; $b = -0.009$, $p < .001$), *Voting & Election* (T16; $b = -0.008$, $p < .001$), *Hydroxychloroquine Misinformation* (T17; $b = -0.005$, $p < .001$), *Economy and Market* (T33; $b = -0.002$, $p < .001$), *Protest & Arrest* (T36; $b = -0.004$, $p < .001$), *Senators Stock Scandal* (T41; $b = -0.003$, $p < .001$), *Critiques of Media and Misinformation* (T42; $b = -0.007$, $p < .001$), and *Relief Bills* (T43; $b = -0.009$, $p < .001$) were discussed more frequently by bots than human Twitter users.

Following the same protocol, we examined topical differences between bots and the news outlets separately. Appendix D displays these topical association patterns. In summary, across all six media outlets, three topics were consistently discussed more by bots (*bs* range: $[-0.035, -0.005]$, all *ps* $< .001$), those being: *Blaming China* (T3), *Vaccine* (T30), and *Senators Stock Scandal* (T41). Second, results show a relatively homogenous pattern when comparing the topical associations between Twitter bots and the centrist media outlets. Three out of eight daily life topics, nine out of fourteen political/societal topics, and eight out of eleven public health topics were associated with bots (*bs* range: $[-0.030, -0.002]$, all *ps* $< .001$) more than either of the two centrist media outlets (i.e., *NYT* and *WP*).

By comparison, topical associations are less consistent when comparing Twitter bots with partisan media outlets, especially liberal media outlets (i.e., *MSNBC* and *HuffPost*). For instance, topics such as *Panic Shopping* (T2; $b = -0.008$, $p < .001$ on *MSNBC*, $b = 0.000$, $p = 0.731$ on *HuffPost*), *Protest and Arrest* (T36; $b = -0.012$, $p < .001$ on *MSNBC*, $b = -0.001$, $p = 0.484$ on *HuffPost*) and *Masking and Personal Hygiene* (T46; $b = -0.011$, $p < .001$ on *MSNBC*, $b = 0.002$, $p = .050$ on *HuffPost*) were discussed more by bots when compared to *MSNBC*, but not significantly different when compared to *HuffPost*. Two daily life topics (i.e., *Critique of Wishful Thinking*, T22; *Hope and Faith*, T26), three out of fourteen political/societal topics (i.e., *Blaming China*, T3; *Hydroxychloroquine Misinformation*, T17; and *Senators Stock Scandal*, T41), and two out of eleven public health topics (i.e., *Spring Break Super Spreader*, T18; and *Vaccine*, T30) were more consistently associated with bot discourse (*bs* range: $[-0.035, -0.003]$, all *ps* $< .001$) than both liberal media outlets. For conservative media outlets *Breitbart* and *Fox News*, we found three out of eight daily life topics (e.g., *Food and Cooking*, T14), half of fourteen political/societal topics (e.g., *Panic Shopping*, T41), and seven out of eleven public health topics (e.g., *Institutional Distrust and*

COVID Denial, T5) were consistently discussed more by bots (*bs* range: [−0.026, 0], all *ps* $<$ .001) than both conservative media outlets.

Time-Series Analyses

Examining bi-directional temporal relationships of topic volume prevalence among social bots, human actors, and journalists would enhance our understanding of how bots influence attention dynamics in the hybrid media system. For each of the 33 named topics, we constructed the time series data on topic prevalence for each actor following the protocol detailed in the Methods section. As we only considered the interpretable topics—33 out of 50—our analysis presents the results of 33 VAR (1) models examining relationships at a one-day lag.

Vector autoregression (VAR) models were built to address RQ3a, RQ3b, RQ4a, and RQ4b. To use VAR models, it is necessary to pre-process the data to remove nonstationarity and seasonality. The former occurs when the statistical properties of a time series are consistent across the time frame. To address nonstationarity, we constructed univariate ARIMA models of each time series (time series with an order of integration that is one or greater are nonstationary) and first-differenced time series with an integrated component (Wei, 1981). The latter is seasonality, characterized by a cyclical pattern in a time series (Cleveland & Tiao, 1976). We found no evidence for significant seasonality in our time series.

Following pre-processing, we tested different lag structures ranging from one to eleven days for each VAR model and used the Akaike information criterion (AIC) to determine the most appropriate number of lags (with AIC, a lower score suggests a more optimal model, see Findley, 1985). We then constructed VAR (1) models for each topic. To facilitate the interpretation of model results, we followed the literature (Benati & Surico, 2009; Soroka, 2002) and applied two post-hoc analyses: Granger causality tests and impulse response functions. Though inadequate to establish causal relationships, Granger causality tests can help examine the predictive validity of how one variable forecasts the subsequent occurrence of another while adjusting for lagged past observations (Toda & Phillips, 1994). Impulse response functions (IRFs) were used to assess the persistence of statistically significant temporal relationships. In IRFs, the independent variable is treated as an impulse to examine how a "shock" of the impulse impacts the dependent variable (Lütkepohl, 2010). Given the focus on quantifying the magnitude of temporal persistence, IRF results are often discussed in temporal units such as days. To conduct the two-way Granger causality tests, we applied the Toda–Yamamoto method (Toda & Yamamoto, 1995). These results are confirmed using Prais–Winsten regressions, reflecting the robustness of the results (see Appendix G).

Take Topic 2 *Panic Shopping* as an example. In Table 1, the Granger Causality tests suggested a unidirectional relationship from human Twitter users to Twitter bots ($X^2 = 10.29$, $p = .001$), but not between bots and any other media outlets. To examine the human-bot relationship under this topic further, IRFs were used to analyze how a shock in human Twitter users' activities impacted bot activities. Figure 2 displayed eight IRFs (from all nonbot actors to bots and from bots to all nonbot actors). The IRFs show that Twitter bots responded to human activities at $t = 1$ and this effect was substantial and lasted for ten days. In other words, a shock in human activities on the topic *Panic Shopping* at $t = 1$ increased bots' activity immediately and this amplification effect took over 10 days to diminish.

Table 1 Granger Causality Tests on the Topic 2 Panic Shopping

	X^2	*p-value*		X^2	*p-value*
Human → Bot	10.29	.001	Bot → Human	0.18	.673
Conservative media → Bot	0.01	.929	Bot → Conservative media	1.03	.311
Centrist media → Bot	0.85	.358	Bot → Centrist media	3.37	.067
Liberal media → Bot	2.10	.147	Bot → Liberal media	0.01	.911

We summarized results from the time series analyses in Table 2. On the left of this table, we placed four non-bot actors as potential predictors of future bot activities; and on the right, we placed the same set of actors as outcome variables potentially responding to prior bot activities. These non-bot actors included human Twitter users, conservative, centrist, and liberal media outlets, respectively. An orange tile in Table 2 indicates a significant positive relationship, blue is a negative relationship, and white is a nonsignificant relationship. To better quantify the duration of the shock on the temporal relationship, the number of days before this relationship vanished was tallied and displayed within the tile. This number in each colored tile can be interpreted as the persistence of a temporal relationship of the related amplification.

Overall, bots selectively amplified human-generated tweets more than media coverage. Human discourses on Twitter positively predicted bot activities with regard to five out of the eight topics in the Daily Life cluster: *Panic Shopping* (T2; $X^2 = 10.294$, $p = .001$), *School Closures & Changes* (T19; $X^2 = 8.608$, $p = .003$), *Critique of Wishful Thinking* (T22; $X^2 = 9.020$, $p = .003$), *Hope & Faith* (T26; $X^2 = 5.719$, $p = .017$), and *National Sports Cancellation* (T34; $X^2 = 34.756$, $p < .001$). The persistence of these relationships (from +3 to +10 days) was generally substantial. Under the topic *Hope & Faith*, for example, VAR simulations on average bots positively responded to a shock in human activity up to three days in the past. For the remaining four daily life topics, a shock in human activity predicted increased bot activities up to ten days later.

We also found similar bot amplification on several human-generated political/societal and public health topics, though with varying levels of persistence of temporal relationships. For instance, bot activity rose with an increase in human discussions on three topics falling into the Political/Societal category: *Travel Bans and Global Lockdown* (T6, $X^2 = 22.451$, $p < .001$) and *Call for Mitigation Strategies* (T48, $X^2 = 5.349$, $p = .021$) lasting ten days, and *Economy and Market* (T33, $X^2 = 4.152$, $p = .042$) with a smaller impact lasting four days. Most topics in the Public Health cluster did not show significant bot amplification with the exceptions including *Testing Availability* (T31, $X^2 = 7.676$, $p = .006$) and *COVID-19 Research* (T47, $X^2 = 4.186$, $p = .041$). On both topics, the shock's impacts were estimated to last three days. Centrist media negatively predicted bot activities under four topics: *National Sports Cancellation* (T34, $X^2 = 4.410$, $p = .036$), *Travel Bans and Global Lockdown* (T6, $X^2 = 12.223$, $p < .001$), *Captain Crozier Scandal* (T21, $X^2 = 4.940$, $p = .026$), and *Economy and Market* (T33, $X^2 = 4.320$, $p = .038$). In other words, when centrist media covered the four above-mentioned topics, bots emphasized them less.

In contrast, human users did not appear to respond to amplified topics by Twitter bots. Most results were nonsignificant, suggesting human Twitter users would not take cues from

Table 2 Summary of All Time-Series Analyses (with Bot Threshold at 0.7)

Human	Conservative	Central	Liberal	*to Bots*	*Bots to*	Human	Conservative	Central	Liberal
Daily Life (8)									
+10				**Panic Shopping (T2)**					
				Dog Immunity (T7)					+10
				Symptom Experience (T11)					+10
				Food & Cooking (T14)					+10
+10				**School Closures & Changes (T19)**					
+10				**Critique of Wishful Thinking (T22)**		−3			
+3				**Hope & Faith (T26)**					
+10		−10		**National Sports Cancellation (T34)**		−10			
Political/Societal (14)									
				Blaming China (T3)					
				Trump Response Failure (T4)					+10
+10		−4		**Travel Bans & Global Lockdown (T6)**		−10			
				Voting & Election (T16)					
				Hydroxychloroquine Misinformation (T17)					
		−10		**Captain Crozier Scandal (T21)**					+4
				Racism & Xenophobia (T25)			+10		
+4		−4		**Economy & Market (T33)**					
				Protest & Arrest (T36)					
				Senators Stock Scandal (T41)					
				Critiques of Media & Misinformation (T42)					
				Relief Bills (T43)					
+10				**Call for Mitigation Strategies (T48)**					
				Workers & Unemployment (T49)					

(*Continued*)

Table 2 (continued)

Human	Conservative	Central	Liberal	*to Bots*		*Bots to*	Human	Conservative	Central	Liberal
Public Health (11)										
					Institutional Distrust & COVID Denial (T5)					
					Frontline Gratitude & Support (T12)					
					Case Reporting (T15)					
					Spring Break Super Spreader (T18)					
					Vaccine (T30)					
+3					**Testing Availability (T31)**					+2
					Epicenter Cities (T38)					
					PPE & Hospital Supplies (T39)					
					Safety Orders & Guidelines (T45)					
					Masking & Personal Hygiene (T46)			+10		
+3					**COVID-19 Research (T47)**					

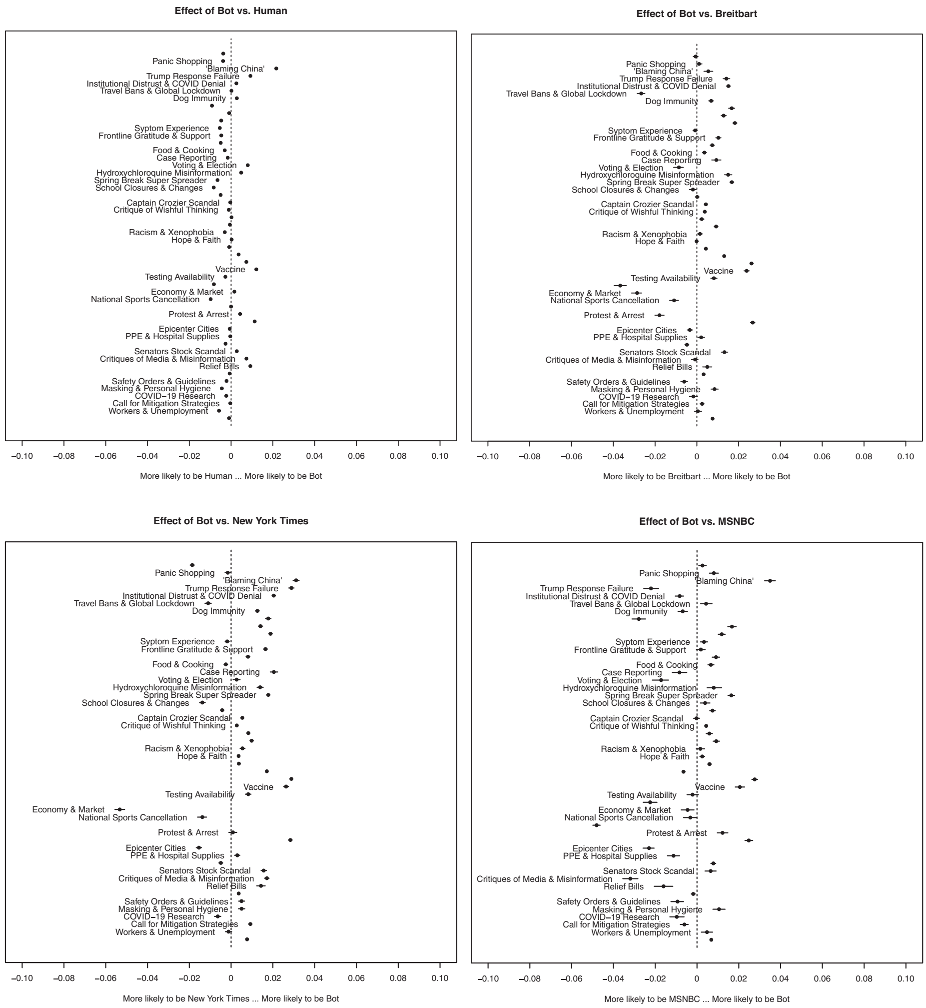

Figure 1 *Topical Association Difference Between Twitter Bots (Right) and Some Non-Bot Actors (Left).*

Note. See full-size version in Online Appendix D.

bot-driven discourse. In fact, bot activities negatively predicted human conversation under three topics: *Travel Bans & Global Lockdown* (T6, $X^2 = 7.700$, $p = .006$), *Critique of Wishful Thinking* (T22, $X^2 = 6.602$, $p = .010$), and *National Sports Cancellation* (T34, $X^2 = 5.360$, $p = .021$), suggesting that once amplified by bots, these topics were not taken up again by humans.

More interestingly, news organizations responded to bot-driven discourse in a heterogeneous way. On one hand, liberal media picked up topics that stressed disease severity, safeguards, and preventative measures, along with topics damaging to Donald Trump, such as *Dog Immunity* (T7, $X^2 = 5.932$, $p = .015$), the severity of *Symptom Experiences* (T11, $X^2 =$

5.149, p = .023), *Food and Cooking* tips during shutdown (T14, $X^2 = 5.047$, p = .025), *Testing Availability* in the early pandemic (T31, $X^2 = 8.105$, p = .004) as well as *Trump's Response Failure* (T4, $X^2 = 6.183$, p = .013) and the firing of *Captain Crozier* (T21, $X^2 = 9.706$, p = .002). For most of these topics, liberal media positively responded to a shock in bot activity that lasted ten days before diminishing.

Conversely, conservative media positively reacted to bots on only two topics: *Racism and Xenophobia* (T25, $X^2 = 4.644$, p = .031) and *Masking and Personal Hygiene* (T46, $X^2 = 4.314$, p = .038). Both topics fit within the conservative media narratives that identify race and ethnicity as problematic categories and emphasize personal responsibility rather than government mandates for COVID-19 management. The predictive power of social bots lasted for up to ten days. In contrast, centrist media rarely responded to bot activity, further suggesting that bot emphasis served as a cue prompting partisan media to cover COVID-19 in ways that aligned with partisan news narratives.

Discussion

Our findings reveal how social bots—i.e., algorithmic agents that interact with selected actors on social media—emphasized certain COVID-19 topics over others, amplified certain voices in the hybrid media system, and predicted partisan news coverage around ideological narratives in the early days of the pandemic. The fact that these automated communicators provided signals to follow-up news coverage of partisan news outlets suggests social bots hold sway beyond the Twitter platform. Our findings thus extended Chadwick's (2013) notion of hybrid media by empirically demonstrating how social bots may serve as conduits transferring topical salience from online human discourses to certain news outlets' issue agenda.

We posed four research questions: (1) What is the volume of information generated by Twitter bots? (2) What topics do Twitter bots emphasize compared to human users and partisan and mainstream news coverage? (3) What are the overtime relationships between topics emphasized by human accounts and Twitter bots? (4) What are the overtime relationships between topics emphasized in news coverage and Twitter bots? Using bot detection, structural topic modeling, and time series analysis, our findings suggested that Twitter bots were mainly amplification engines, retweeting most often (75.4% of observed content), though nearly a quarter of social bot posts were original content (19.8%) or commenting of others' posts (4.8%). Comparing accounts identified as social bots to non-bot actors, we not only observed different interaction patterns, but also distinctive topical emphasis relative to humans and news media.

Theoretical and Practical Implications

Our study revealed how topics discussed by Twitter bots responded to content production by human and new outlets and, in turn, predicted the attention allocation from these actors. Systematic characterization of the way these algorithmic agents discursively participated in the macro-level temporal dynamics connecting social media discourses with news coverage within the hybrid media ecosystem is overdue. Our observations indicate that Twitter bots selectively amplify a broad set of topics, reflecting the prevalence of bots in the information ecology, and that overtime relationships between bots and news coverage tend to fall along

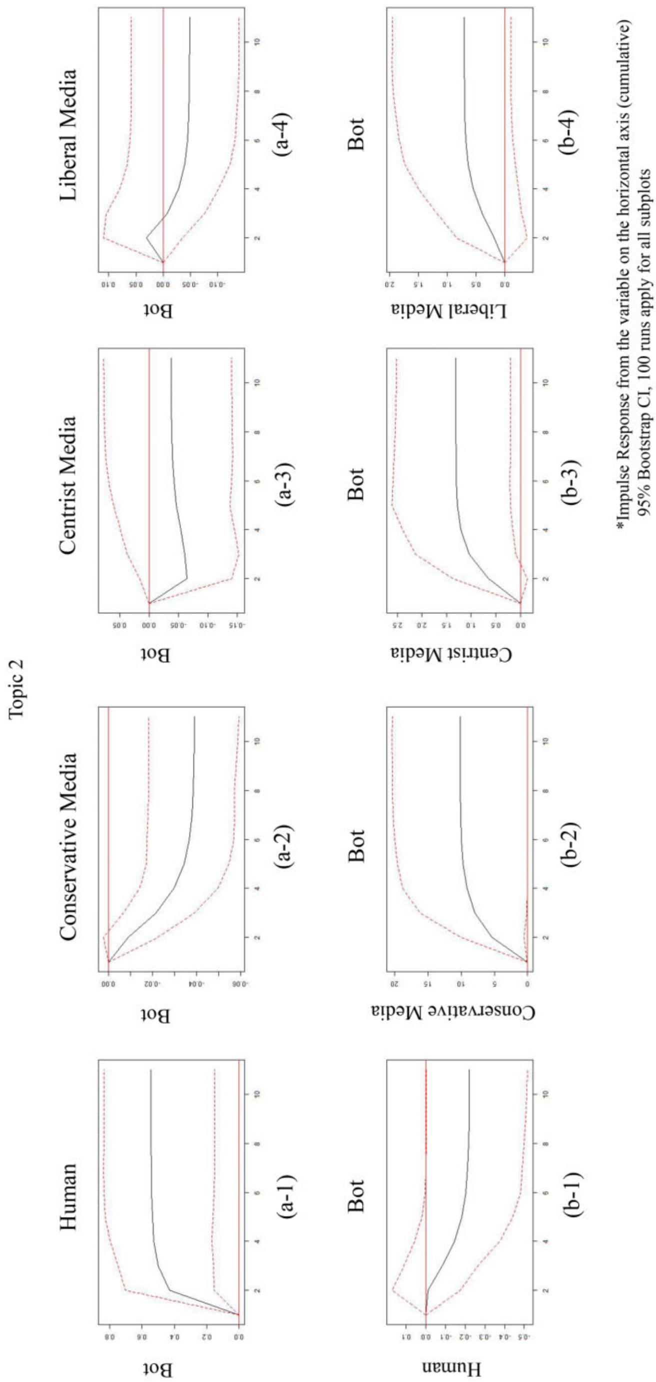

Figure 2 *Impulse Response Function Plots (Topic 2).*

ideological lines. Given that social media research often removes bot activities as "noise" (e.g., Deb et al., 2019), these findings highlight the importance of considering these algorithmic agents as players in the hybrid media system, who are acting to amplify certain viewpoints and topics, which in turn can shape partisan news coverage.

First, Twitter bots amplified the salience of issues in human-generated content more than news coverage. Although bot accounts only made up a relatively small proportion in the overall Twitterverse, bot activities nevertheless amplified voices from human users across a diverse set of topics, mainly involving Daily Life issues such as *Panic Shopping*, *National Sports Cancellations*, among others. Social bots also amplified human discourse on Political/Societal topics such as *Travel Bans & Global Lockdown*, and Public Health conversations such as *Testing Availability* and *COVID-19 Research*. This is clear, if somewhat mundane, evidence of bots' entry into the online information flow and their ramping up of selected human communications. Much of this reflect automated accounts designed to circulate updates or respond to events, simply "raising the volume" of human voices through algorithmic amplification. This imitation of human behaviors not only makes bot detection work more challenging but also underscores two contrasting issues: (1) that bot behaviors, especially when operating at scale, may elude awareness of ordinary people and media professionals alike absent detection software, and (2) that bot accounts, when examined at scale, may most readily reveal simple amplification activities, minimizing the sense of their potential impact on other actors in the hybrid media system. We also recognize that the detection of bot discourse as meaningfully distinguishable from human voices in the broad informational ecology will become more difficult as bots continue evolving to become more sophisticated and more "human-like" in their actions and content generation (Dathathri et al., 2020; Zhang, Sun, Galley, et al., 2019). These changes are likely to further empower these algorithmic communicators to move beyond imitation and amplification towards the production of distinctive and autonomous voices.

Second, we found that when Twitter bots amplify topics like *Travel Bans and Global Lockdown* and *Critique of Wishful Thinking*, human users subsequently reduce their posting on these topics, suggesting that the public may have re-allocated their attention to other issues. Besides amplification effects, we also observed avoidance and competition as evidenced by *negative* temporal relationships. For instance, when centrist media reported on the *Captain Crozier Scandal*, an event embarrassing to the Trump's administration, their increased emphasis was preceded by Twitter bots' avoidance of these topics. This suggests bots may avoid emphasizing topics once they are receiving widespread attention.

Third, and perhaps more consequential, our time-series analysis reveals that bot amplification predicts news content in a manner that supports the ideology of partisan media outlets. We observed a heterogeneous pattern, but a pattern nonetheless: Liberal media echoed disease severity and safeguards, which inherently treated the pandemic as real and serious, and critiques that allowed them to write about topics blaming Donald Trump when bot activities rose on these issues. At the same time, conservative media positively reacted to bots when they amplified content on, for example, *Racism and Xenophobia*, topics that allow conservatives to build media narratives and reframe these issues to serve their ideological goals. These patterns were largely replicated via robustness checks (see Appendix I, p. 71–72) using a more liberal 0.5 threshold for bot detection, adding additional 11,517 bot accounts to our analysis, or 18% more than at the 0.7 threshold.

These findings connect to recent work documenting the politicization of news coverage on mitigation policies and compliance with preventive measures around COVID-19 (Gadarian et al., 2021; Grossman et al., 2020). In this environment, partisan media, especially liberal outlets, emphasized ideologically aligned topics, seemingly using the cue of bot amplification to justify writing about issues that fit the ideological leanings of their audiences, their approach to disease coverage, and their desire to either critique or support President Trump and his handling of COVID-19. Admittedly, our findings cannot directly pin down the causal relationships between Twitter bots' topic amplification and observed partisan politicization with regards to these critical public health issues. That said, we did not find evidence that bots had picked up topics from partisan news outlets; rather, the temporal relationships appear to move in one-direction for partisan media—from bots to subsequent partisan coverage. It is also worth noting that the bi-directional relationships between bots and centrist media were much weaker. Taken together, these results demonstrated closer alignment between partisan media and Twitter bots during the early phase of the COVID-19 pandemic. Future research is warranted to investigate micro-level mechanisms that may give rise to these documented macro-level asymmetric temporal patterns, given that algorithmic influence appears to hold some sway over news content, especially coverage generated by liberal news media outlets.

Given the growing concern that social media metrics are often taken by newsrooms as traces of human interactions and opinion expressions without much scrutiny (Petre, 2021), one possibility is that partisan media outlets are reacting to digital traces left by bots. This is not to deny partisan media's own selectivity in choosing what topics to cover, as bots' amplification effects are likely to be conditional upon partisan media's own norms and issue agenda. What we argue is that the presence of bots and their selective topical foci produced digital traces that may have helped accelerate and deepen the magnitude of polarization in issue coverage by partisan media, and that these patterns may be asymmetric. By situating the discussion of Twitter bots within a heterogeneous media ecology where centrist, liberal, and conservative news media co-exist and differ in their news production norms and likelihood to respond to metrics of online engagement, our work highlighted the possibility that coordinated bot amplification can increase the malleability of attention dynamics connecting social media with news coverage.

As social bots evolve towards the appearance of human-like interaction, future research needs to seriously consider when and how coordinated bot activities may induce systematic biases in the sensing and coverage of public opinions in a hybrid media ecology. It may be that left-leaning journalists are more likely to treat social media platforms like Twitter as "purveyors of legitimated content" and "conduits to surface the words of 'ordinary people'" without the need for verification (Molyneux & McGregor, 2021), whereas right-leaning journalists and outlets may be less responsive to social media discourse as "vox populi" and more reactive to coverage on the left (Zhang et al., 2022). Future research must consider this possibility.

Limitations and Future Direction

Future research is encouraged to address several limitations in this study. First, while we took care to develop and polish the keyword list, public discourse about the pandemic evolved rapidly and new keywords probably need to be added if researchers plan to expand the time frame. Second, there is no perfect tool for bot detection and in this study, we had to rely on the results suggested by one highly useful and state-of-the-art tool, Botometer. The continuous development of more sophisticated detection techniques could help validate our results and we welcome future methodological advances. In addition, we adopted an established bot score threshold following previous research (Sayyadiharikandeh et al., 2020). Our robustness checks, in which we varied this threshold, suggest that the key findings reported in the main analyses are robust with reasonable model variations. Botometer helped label a relatively small yet non-negligible proportion (8.98%) of total Twitter accounts in our dataset as social bots, which, consistent with other studies (Assenmacher et al., 2020), currently functioned more as amplification engines than highly intelligent, human-like communicators. That said, similar to existing conversational chat bots (e.g., Yan, 2018), the next generation of social bots may widely incorporate deep learning technologies to acquire human-like intelligence for online discursive exchanges. Future research should continue to monitor the roles of social bots in the hybrid media ecology. Third, we selected six news outlets to represent the conservative, the liberal, and the centrist media in the United States, but future research should include a broader set of news outlets to further investigate the role of social bots in the more complex and diverse media environment. Lastly, although our study is descriptive in nature and cannot directly elucidate micro-level causal mechanisms, our findings warrant such efforts in future research and point to several promising directions. For example, it is worthwhile examining the direct interactions between bots, human users, and news organizations' Twitter accounts to better clarify their discursive influences on each other, and whether these patterns are symmetric or asymmetric across issues. To do so, using Twitter's streaming APIs is more appropriate than our archival retrieval approach as Twitter routinely removes suspected bot accounts. Another direction is to employ recently proposed personalized page-rank sampling method (Zhang et al., 2022) to identify potential "flocks" of bots and empirically quantify the degree of coordination among bot accounts. Our preliminary analyses provided initial suggestive evidence for bot coordination (see Appendix C), but more systematic research is needed.

Conclusion

Against the backdrop of intense social discussion surrounding the COVID-19 pandemic, we found non-negligible presence of bot activities in the Twittersphere. Combining automated bot detection, structural topic modeling, and time series analyses, our longitudinal study demonstrated the importance of considering social bots as an active player in today's hybrid media ecology. Rather than noises to remove, the social bots we examined selectively amplified certain human discourses and were predictive of subsequent news coverage in ways that suggest news outlets relied on bot-driven metrics when they matched their ideological slant. This dynamic may contribute to the long-term trend of partisan issue polarization in news coverage. This macro-level analysis invites future research to elucidate micro-level mechanisms giving rise to documented asymmetric temporal patterns, such as how newsrooms along the partisan spectrum may differentially react to social media metrics susceptible to bot influence. In an era of artificial intelligence, we provide novel insights into the roles played by social bots in contributing to the hybridity of today's media system.

Note

1. At the 0.7 threshold, Botometer failed to calculate the scores for about 1.74% of the accounts, which were either under protection status or removed at the time our data were collected. Those accounts were removed from analyses.

Acknowledgments

The authors thank Michael Wagner, Katherine J. Cramer, Yicheng Hu, Alex Zhi-Xiong Koo, Sicong Zhang, Xuanzhou Chen, and "Estelle" Ranran Zhu Mi for their support and insightful feedback to our project. We are grateful to Yiting Duan, Yiming Sui, and Xinyi Liu for their research assistance. We also thank the anonymous reviewers and editors for their review and helpful feedback.

Funding

This research was supported by the National Science Foundation's Convergence Accelerator program under award number 2137724, and grants from the University of Wisconsin - Madison Office of the Vice Chancellor for Research and Graduate Education, with funding from the Wisconsin Alumni Research Foundation, the William and Flora Hewlett Foundation, and the John S. and James L. Knight Foundation.

Conflict of Interest

The authors have no conflict of interest to report.

Supplementary Material

Supplementary material is available online at *Human Communication Research.*

References

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. https://doi.org/10.1177/2056305120939264

Alothali, E., Zaki, N., Mohamed, E. A., & Alashwal, H. (2018, November). *Detecting social bots on Twitter: A literature review.* 2018 International Conference on Innovations in Information Technology (IIT), Al Ain, United Arab Emirates. https://doi.org/10.1109/INNOVATIONS.2018.8605995

Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. F., Lee, J., Mann, M., Merhout, F. & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, *115*(37), 9216–9221. https://doi.org/10.1073/pnas.1804840115

Benati, L., & Surico, P. (2009). VAR analysis and the great moderation. *American Economic Review*, *99*(4), 1636–1652. https://doi.org/10.1257/aer.99.4.1636

Bermejo, F., Reggi, L., Tiribelli, S., & Zuckerman, E. (2020). *Coverage of COVID-19 and political partisanship-comparing across nations*. Retrieved from Media Cloud. https://civic.mit.edu/index.html%3Fp=2771.html

Bessi, A., & Ferrara, E. (2016). Social bots distort the 2016 US Presidential Election online discussion. *First Monday*, *21*(11–7). http://dx.doi.org/10.5210/fm.v21i11.7090

Boshmaf, Y., Muslukhov, I., Beznosov, K., & Ripeanu, M. (2013). Design and analysis of a social botnet. *Computer Networks*, *57*(2), 556–578. https://doi.org/10.1016/j.comnet.2012.06.006

Box-Steffensmeier, J. M., Freeman, J. R., Hitt, M. P., & Pevehouse, J. C. (2014). *Time series analysis for the social sciences*. Cambridge University Press. https://doi.org/10.1017/CBO9781139025287

Broniatowski, D. A., Jamison, A. M., Qi, S., AlKulaib, L., Chen, T., Benton, A., Quinn, S. C., & Dredze, M. (2018). Weaponized health communication: Twitter bots and Russian trolls amplify the vaccine debate. *American Journal of Public Health*, *108*(10), 1378–1384. https://doi.org/10.2105/AJPH.2018.304567

Chadwick, A. (2013). *The hybrid media system: Politics and power*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199759477.001.0001

Chen, E., Lerman, K., & Ferrara, E. (2020). Tracking social media discourse about the covid-19 pandemic: Development of a public coronavirus twitter data set. *JMIR Public Health and Surveillance*, *6*(2), e19273. https://doi.org/10.2196/19273

Chen, K., Chen, A., Zhang, J., Meng, J., & Shen, C. (2020). Conspiracy and debunking narratives about COVID-19 origination on Chinese social media: How it started and who is to blame. *arXiv* preprint. https://arxiv.org/abs/2011.08409

Chen, K., Duan, Z., & Yang, S. (2021). Twitter as research data: Tools, costs, skill sets, and lessons learned. *Politics and the Life Sciences*, 1–17. https://doi.org/10.1017/pls.2021.19

Chung, M. (2017). Not just numbers: The role of social media metrics in online news evaluations. *Computers in Human Behavior*, *75*, 949–957. https://doi.org/10.1016/j.chb.2017.06.022

Cleveland, W. P., & Tiao, G. C. (1976). Decomposition of seasonal time series: A model for the census X-11 program. *Journal of the American Statistical Association*, *71*(355), 581–587. https://doi.org/2285586

Cresci, S. (2020). A decade of social bot detection. *Communications of the ACM*, *63*(10), 72–83. https://doi.org/10.1145/3409116

Dathathri, S., Madotto, A., Lan, J., Hung, J., Frank, E., Molino, P., Yosinski, J., & Liu, R. (2020). Plug and play language models: A simple approach to controlled text generation. *arXiv* preprint. http://arxiv.org/abs/1912.02164

Deb, A., Luceri, L., Badaway, A., & Ferrara, E. (2019, May). *Perils and challenges of social media and election manipulation analysis: The 2018 us midterms.* In Companion Proceedings of the 2019 World Wide Web Conference (pp. 237–247). San Francisco, CA. https://doi.org/10.1145/3308560.3316486

Edwards, C., Edwards, A., Spence, P. R., & Shelton, A. K. (2014). Is that a bot running the social media feed? Testing the differences in perceptions of communication quality for a human agent and a bot agent on Twitter. *Computers in Human Behavior*, *33*, 372–376. https://doi.org/10.1016/j.chb.2013.08.013

Ferrara, E. (2020). #COVID-19 on Twitter: Bots, Conspiracies, and Social Media Activism. *arXiv* preprint. https://arxiv.org/abs/2004.09531v1

Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The rise of social bots. *Communications of the ACM*, *59*(7), 96–104. https://doi.org/10.1145/2818717

Findley, D. F. (1985). On the unbiasedness property of AIC for exact or approximating linear stochastic time series models. *Journal of Time Series Analysis*, *6*(4), 229–252. https://doi.org/10.1111/j.1467-9892.1985.tb00412.x

Gadarian, S. K., Goodman, S. W., & Pepinsky, T. B. (2021). Partisanship, health behavior, and policy attitudes in the early stages of the COVID-19 pandemic. *Plos One*, *16*(4), e0249596. https://doi.org/10.1371/journal.pone.0249596

González-Bailón, S., & De Domenico, M. (2021). Bots are less central than verified accounts during contentious political events. *Proceedings of the National Academy of Sciences*, *118*(11). https://doi.org/10.1073/pnas.2013443118

Grimme, C., Preuss, M., Adam, L., & Trautmann, H. (2017). Social bots: Human-like by means of human control. *Big Data*, *5*(4), 279–293. https://doi.org/10.1089/big.2017.0044

Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., & Lazer, D. (2019). Fake news on Twitter during the 2016 US presidential election. *Science*, *363*(6425), 374–378. https://www.science.org/doi/10.1126/science.aau2706

Grossman, G., Kim, S., Rexer, J. M., & Thirumurthy, H. (2020). Political partisanship influences behavioral responses to governors' recommendations for COVID-19 prevention in the United States. *Proceedings of the National Academy of Sciences*, *117*(39), 24144–24153. https://doi.org/10.1073/pnas.2007835117

Guzman, A. L., & Lewis, S. C. (2020). Artificial intelligence and communication: A human–-machine communication research agenda. *New Media & Society*, *22*(1), 70–86. https://doi.org/10.1177/1461444819858691

Hart P. S., Chinn S., & Soroka S. (2020). Politicization and polarization in COVID-19 news coverage. *Science Communication*, *42*(5) 679–697. https://doi.org/10.1177/1075547020950735

Hill, J., Ford, W. R., & Farreras, I. G. (2015). Real conversations with artificial intelligence: A comparison between human–human online conversations and human–chatbot conversations. *Computers in Human Behavior*, *49*, 245–250. https://doi.org/10.1016/j.chb.2015.02.026

Howard, P. N., Woolley, S., & Calo, R. (2018). Algorithms, bots, and political communication in the US 2016 election: The challenge of automated political communication for election law and administration. *Journal of Information Technology & Politics*, *15*(2), 81–93. https://doi.org/10.1080/19331681.2018.1448735

Jamison, A. M., Broniatowski, D. A., & Quinn, S. C. (2019). Malicious actors on Twitter: A guide for public health researchers. *American Journal of Public Health*, *109*(5), 688–692. http://doi.org/10.2105/AJPH.2019.304969

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. https://doi.org/10.1080/10584609.2019.1661888

Keller, T. R., & Klinger, U. (2019). Social bots in election campaigns: Theoretical, empirical, and methodological implications. *Political Communication*, *36*(1), 171–189. https://doi.org/10.1080/10584609.2018.1526238

Kim, Y., Huang, J., & Emery, S. (2016). Garbage in, garbage out: Data collection, quality assessment and reporting standards for social media data use in health research, infodemiology and digital disease detection. *Journal of Medical Internet Research*, *18*(2), e41. https://doi.org/10.2196/jmir.4738

King, G., Pan, J., & Roberts, M. E. (2017). How the Chinese government fabricates social media posts for strategic distraction, not engaged argument. *American Political Science Review*, *111*(3), 484–501. https://doi.org/10.1017/S0003055417000144

Levendusky, M. (2013). *How partisan media polarize America*. University of Chicago Press. https://press.uchicago.edu/ucp/books/book/chicago/H/bo16468853.html

Lütkepohl, H. (2010). *Impulse response function*. In S. N. Durlauf, & L. E. Blume (Eds), *Macroeconometrics and time series analysis* (pp. 145–150). The New Palgrave Economics Collection. Palgrave Macmillan. https://doi.org/10.1057/9780230280830_16

McGregor, S. C., & Molyneux, L. (2020). Twitter's influence on news judgment: An experiment among journalists. *Journalism*, *21*(5), 597–613. https://doi.org/10.1177/1464884918802975

Messing, S., & Westwood, S. J. (2014). Selective exposure in the age of social media: Endorsements trump partisan source affiliation when selecting news online. *Communication Research*, *41*(8), 1042–1063. https://doi.org/10.1177/0093650212466406

Meta. (2021). *Coordinated inauthentic behavior archives*. https://about.fb.com/news/tag/coordinated-inauthentic-behavior/

Molyneux, L., & McGregor, S. C. (2021). Legitimating a platform: Evidence of journalists' role in transferring authority to Twitter. *Information, Communication & Society*, 1–19. https://doi.org/10.1080/1369118X.2021.1874037

Motta, M., Stecula, D., & Farhart, C. (2020). How right-leaning media coverage of COVID-19 facilitated the spread of misinformation in the early stages of the pandemic in the US. *Canadian Journal of Political Science*, *53*(2), 335–342. https://doi.org/10.1017/S0008423920000396

Pacheco, D., Hui, P. M., Torres-Lugo, C., Truong, B. T., Flammini, A., & Menczer, F. (2020). Uncovering coordinated networks on social media: Methods and case studies. *arXiv* pre-print. https://arxiv.org/abs/2001.05658

Petre, C. (2021). *All the news that's fit to click: How metrics are transforming the work of journalists*. Princeton University Press. https://press.princeton.edu/books/hardcover/9780691177649/all-the-news-thats-fit-to-click

Ratkiewicz, J., Conover, M., Meiss, M., Gonçalves, B., Flammini, A., & Menczer, F. (2021). Detecting and tracking political abuse in social media. *Proceedings of the International AAAI Conference on Web and Social Media*, *5*(1), 297–304. https://ojs.aaai.org/index.php/ICWSM/article/view/14127

Roberts, M. E., Stewart, B. M., & Tingley, D. (2019). Stm: An R package for structural topic models. *Journal of Statistical Software*, *91*(1), 1–40. https://doi.org/10.18637/jss.v091.i02

Sanders, A. C., White, R. C., Severson, L. S., Ma, R., McQueen, R., Paulo, H. C. A., Zhang, Y., Erickson, J. S., & Bennett, K. P. (2021). Unmasking the conversation on masks: Natural language processing for topical sentiment analysis of COVID-19 Twitter discourse. *medRxiv*. https://pubmed.ncbi.nlm.nih.gov/34457171/

Santini, R. M., Salles, D., & Tucci, G. (2021). Comparative approaches to mis/disinformation| when machine behavior targets future voters: The use of social bots to test narratives for political campaigns in Brazil. *International Journal of Communication*, *15*, 24. https://ijoc.org/index.php/ijoc/article/view/14803

Sayyadiharikandeh, M., Varol, O., Yang, K. C., Flammini, A., & Menczer, F. (2020, October). *Detection of novel social bots by ensembles of specialized classifiers*. In Proceedings of the 29th ACM international conference on information & knowledge management, Virtual Event Ireland. https://doi.org/10.1145/3340531.3412698

Shao, C., Ciampaglia, G. L., Varol, O., Yang, K. C., Flammini, A., & Menczer, F. (2018). The spread of low-credibility content by social bots. *Nature Communications*, *9*(1), 1–9. https://doi.org/10.1038/s41467-018-06930-7

Soroka, S. N. (2002). Issue attributes and agenda-setting by media, the public, and policymakers in Canada. *International Journal of Public Opinion Research*, *14*(3), 264–285. https://doi.org/10.1093/ijpor/14.3.264

Starbird, K. (2019). Disinformation's spread: bots, trolls and all of us. *Nature*, *571*(7766), 449–450. https://doi.org/10.1038/d41586-019-02235-x

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. https://doi.org/10.1073/pnas.1803470115

Stewart, L. G., Arif, A., & Starbird, K. (2018, February). Examining trolls and polarization with a retweet network. In *Proceedings of ACM WSDM, workshop on misinformation and misbehavior mining on the web*, Los Angeles, CA, United States. http://faculty.washington.edu/kstarbi/examining-trolls-polarization.pdf

Toda, H. Y., & Phillips, P. C. (1994). Vector autoregression and causality: A theoretical overview and simulation study. *Econometric Reviews*, *13*(2), 259–285. https://doi.org/10.1080/07474939408800286

Toda, H. Y., & Yamamoto, T. (1995). Statistical inference in vector autoregressions with possibly integrated processes. *Journal of Econometrics*, *66*(1–2), 225–250. https://doi.org/10.1016/0304-4076(94)01616-8

Twitter. (2020). *COVID19 Stream Filtering Rules*. https://developer.twitter.com/en/docs/labs/covid19-stream/filtering-rules

Vu, H. T. (2014). The online audience as gatekeeper: The influence of reader metrics on news editorial selection. *Journalism*, *15*(8), 1094–1110. https://doi.org/10.1177/1464884913504259

Waddell, T. F. (2018). What does the crowd think? How online comments and popularity metrics affect news credibility and issue importance. *New Media & Society*, *20*(8), 3068–3083. https://doi.org/10.1177/1461444817742905

Webster, J. G. (2014). *The marketplace of attention: How audiences take shape in a digital age*. MIT Press. https://doi.org/10.7551/mitpress/9892.001.0001

Wei, W. W. (1981). Effect of systematic sampling on ARIMA models. *Communications in Statistics-Theory and Methods*, *10*(23), 2389–2398. https://doi.org/10.1080/03610928108828197

Welbers, K., Van Atteveldt, W., Kleinnijenhuis, J., Ruigrok, N., & Schaper, J. (2016). News selection criteria in the digital age: Professional norms versus online audience metrics. *Journalism*, *17*(8), 1037–1053. https://doi.org/10.1177/1464884915595474

Wells, C., Shah, D. V., Pevehouse, J. C., Foley, J., Lukito, J., Pelled, A., & Yang, J. (2019). The temporal turn in communication research: Time series analyses using computational approaches. *International Journal of Communication*, *13*, 4021–4043. https://ijoc.org/index.php/ijoc/article/view/10635

Wells, C., Shah, D., Lukito, J., Pelled, A., Pevehouse, J. C., & Yang, J. (2020). Trump, Twitter, and news media responsiveness: A media systems approach. *New Media & Society*, *22*(4), 659–682. https://doi.org/10.1177/1461444819893987

Wojcik, S., Messing, S., Smith, A., Rainie, L., & Hitlin, P. (2018, April 9). Bots in the Twittersphere. https://www.pewresearch.org/internet/2018/04/09/bots-in-the-twittersphere/

Xia, Y., Lukito, J., Zhang, Y., Wells, C., Kim, S. J., & Tong, C. (2019). Disinformation, performed: Self-presentation of a Russian IRA account on Twitter. *Information, Communication & Society*, *22*(11), 1646–1664. https://doi.org/10.1080/1369118X.2019.1621921

Yan, H. Y., Yang, K. C., Menczer, F., & Shanahan, J. (2021). Asymmetrical perceptions of partisan political bots. *New Media & Society*, *23*(10), 3016–3037. https://doi.org/10.1177/1461444820942744

Yan, R. (2018). "*Chitty-Chitty-Chat Bot": Deep Learning for Conversational AI. In Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence Survey track*. (pp. 5520–5526). https://doi.org/10.24963/ijcai.2018/778

Zamith, R. (2018). On metrics-driven homepages: Assessing the relationship between popularity and prominence. *Journalism Studies*, *19*(8), 1116–1137. https://doi.org/10.1080/1461670X.2016.1262215

Zhang, J. (2021). People's responses to the COVID-19 pandemic during its early stages and factors affecting those responses. *Humanities and Social Sciences Communications*, *8*(1), 1–13. https://doi.org/10.1057/s41599-021-00720-1

Zhang, Y., Chen, F., & Rohe, K. (2022). Social media public opinion as flocks in a murmuration: Conceptualizing and measuring opinion expression on social media. *Journal of Computer-Mediated Communication*, *27*(1). https://doi.org/10.1093/jcmc/zmab021

Zhang, Y., Sun, S., Galley, M., Chen, Y. C., Brockett, C., Gao, X., Gao, J., Liu, J., & Dolan, B. (2019). Dialogpt: Large-scale generative pre-training for conversational response generation. *arXiv* preprint. https://arxiv.org/abs/1911.00536

Zhang, Y., Wells, C., Wang, S., & Rohe, K. (2018). Attention and amplification in the hybrid media system: The composition and activity of Donald Trump's Twitter following during the 2016 presidential election. *New Media & Society*, *20*(9), 3161–312. https://doi.org/10.1177/1461444817744390