

ORIGINAL ARTICLE

Assembling the Networks and Audiences of Disinformation: How Successful Russian IRA Twitter Accounts Built Their Followings, 2015–2017

Yini Zhang^{1*}, Josephine Lukito², Min-Hsin Su³, Jiyoun Suk³,
Yiping Xia³, Sang Jung Kim³, Larissa Doroshenko⁴, & Chris Wells⁵

¹Department of Communication, University at Buffalo, Buffalo, NY, USA

²School of Journalism and Media, University of Texas at Austin, Austin, TX, USA

³School of Journalism & Mass Communication, University of Wisconsin-Madison, Madison, WI, USA

⁴College of Arts, Media and Design, Northeastern University, Boston, MA, USA

⁵Division of Emerging Media Studies, Boston University, Boston, MA, USA

This study investigates how successful Russian Internet Research Agency (IRA) Twitter accounts constructed the followings that were central to their disinformation campaigns around the 2016 U.S. presidential election. Treating an account's social media following as both an ego network and an audience critical for information diffusion and influence accrual, we situate IRA Twitter accounts' accumulation of followers in the ideologically polarized, attention driven, and asymmetric political communication system. Results show that partisan enclaves on Twitter contributed to IRA accounts' followings through retweeting; and that mainstream and hyperpartisan media assisted conservative IRA accounts' following gain by embedding their tweets in news. These results illustrate how network dynamics within social media and news media amplification beyond it together boosted social media followings. Our results also highlight the dynamics fanning the flames of disinformation: partisan polarization, media fragmentation and asymmetry, and an attention economy optimized for engagement rather than accuracy.

Keywords: Social Media Following, Disinformation, Political Polarization, Attention Economy, Asymmetry

doi: 10.1093/joc/jqaa042

As political communication has moved online, information manipulation through digital media has become more prevalent and sophisticated, targeting

*Corresponding author: Yini Zhang; e-mail: yzhang20@buffalo.edu

search engine query results, faking grassroots support, and artificially boosting visibility on social media (Metaxas & Mustafaraj, 2012; Ratkiewicz *et al.*, 2011). The Russian disinformation campaign on social media around the 2016 presidential election season has become a well-studied case of information manipulation. While most research on the operations of the Russian Internet Research Agency (IRA) has assessed IRA social media tactics and potential influence (e.g., Badawy *et al.*, 2018, 2019; Bail *et al.*, 2020), this article focuses on how the most successful impostor IRA Twitter accounts built the followings upon which their influence was based.

Several highly successful IRA Twitter accounts like @TEN_GOP transformed themselves from average Twitter users with a few hundred followers in late 2015 to microcelebrities commanding over one hundred thousand followers before their discovery in September 2017. They also appeared dozens of times in news media, mainstream and fringe alike (Lukito *et al.*, 2020), and interacted with thousands of American Twitter users, including prominent politicians, pundits, and strategists (O'Brien, 2017). By measures of news media exposure and social media engagement, these few IRA accounts and the larger disinformation campaign of which they were a part were highly successful. Integral to this success was the accounts' large Twitter followings, which (a) facilitated the direct dissemination of strategic messages by serving as their audiences; (b) functioned as a set of individuals likely to amplify their messages through retweeting and other means; and (c) provided visible engagement metrics that looked to other observers, including journalists, like evidence of authority and authenticity.¹

We conceptualize social media following as both an ego network and an audience, which ultimately operates as a critical asset of communicative power. We then take a quantitative case study approach to analyze how the four most successful English-speaking IRA Twitter accounts (quantified in terms of retweets; Cha *et al.*, 2010) constructed their followings. In doing so, this article makes two sets of contributions: the first concerns discussion about how social media networks evolve ecologically (McPherson *et al.*, 2001) and how attention and audiences are captured in an interlinked, hybrid media system (Chadwick, 2017; Webster, 2014). The second adds to our understanding of disinformation in the broader configuration of the political communication system. Research on the information disorder (Wardle & Derakhshan, 2017) tends to focus on diffusion within certain platforms. However, we must recognize that most (dis)information campaigns involve not a single platform or diffusion dynamic, but multiple parts of the political communication system as well as the interplay between institutional, cultural, and technological undercurrents in the system (Benkler, Faris, & Roberts, 2018; Bode & Vraga, 2018; Jamieson, 2020).

Using a combination of three datasets—tweets related to IRA accounts from a random 1% Twitter archive, partisan and mainstream news media's uptake of IRA accounts' tweets, and a unique set of politically active Twitter users' engagement with IRA accounts—and applying time series modeling and computational quantitative analysis, we find that the engagement with partisan Twitter networks was key

to the most successful IRA accounts' assemblage of followers. For IRA accounts masquerading as conservatives, the gain in Twitter following was also partly fueled by hyperconservative outlets' and mainstream media's amplification of their messages. Thus, the accumulation of a large social media following rests on both network dynamics within social media and news media amplification beyond it. By demonstrating how partisan Twitter enclaves and news media helped coalesce Twitter following, our findings reveal mechanisms of network and audience building in the deeply connected political communication system. Moreover, our results provide empirical evidence showing how features of the American political communication system—deep partisan polarization, news media's drive for attention, and the conservative media ecosystem in particular—fed the success of political disinformation actors.

In what follows, we describe the larger political communication context and connect it to disinformation on social media. Then we discuss the nature of social media following and its importance for any actor to spread messages and grow influence. We next introduce hypotheses about how the forces of the political communication system contributed to the IRA accounts' accumulation of Twitter followers.

The new political communication system: Old players and new entrants in an attention economy, embedded in asymmetrical media ecosystems

In the past two decades, a host of digital media outlets and social media platforms have emerged, featuring more partisan slant and overt conflict than was common in the broadcast era (Baum & Groeling, 2008). Those new entrants compete with legacy outlets for attention by operating on the logics of inexpensive production, network distribution, and niche targeting (Klinger & Svensson, 2015). This competition exists in tandem with cooperation and coadaptation between the older and newer media, resulting in an interdependence between actors in the media system (Chadwick, 2017).

In the meantime, average citizens have become empowered to influence news agendas through connective actions organized over digital platforms (Bennett & Segerberg, 2013) or creative expressions during high-profile media events (Freelon & Karpf, 2015). As pathways to capture attention are no longer exclusively access to traditional media gates (Tufekci, 2013), social media has become a tool for growing one's audience and gaining attention (Tufekci & Wilson, 2012). Even “decentralized networks of people” (Boyd, 2017)—like Anonymous and 4chan communities—can exploit the media system by manipulating trending topics or memes. This new configuration opens opportunities for democratic movements, fake news outlets, and the fringes of ideological/cultural spheres to advance issue agendas, define issue frames, and amplify their voices (e.g., Freelon, McIlwain, & Clark, 2016; Marwick & Lewis, 2017; Vargo, Guo, & Amazeen, 2018).

A direct outcome of such a dazzling array of actors entering the political communication system is the cacophony of information and, consequently, the scarcity

of attention (Webster, 2014; Wu, 2017). Within the attention economy, power lies in the ability to command attention. The pressure to attract attention shapes communicators' behaviors, including chasing objects that bring attention (Karpf, 2016) and producing viral and click-bait content (McGregor, 2019; Munger, 2020).

Furthermore, the media system is asymmetrical, as illustrated in discrepant structures, norms, and cultures of the liberal and conservative media ecosystems (Benkler, Faris, & Roberts, 2018). While the former is undergirded by truth-seeking journalistic norms and tightly coupled with politically centrist legacy journalistic institutions, the latter is detached from the center and insulated by the doctrines of identity and ideological purity, making it more susceptible to the influence of ideologically-driven activism, fake news, and disinformation campaigns (Benkler, Faris, & Roberts, 2018; Bovet & Makse, 2019; Freelon, Marwick, & Kreiss, 2020; Hjorth & Adler-Nissen, 2019).

In short, a media system that is highly ideological, permeable, and attention-driven opens up opportunities for a myriad of actors to innovate ways of inserting themselves into the attention ecology and to disseminate messages and accumulate influence, in good and bad faith. And the asymmetry between the liberal and conservative media ecosystems suggests the likelihood of divergent processes of information flow and influence accrual.

Social media disinformation

Our cases, successful IRA Twitter accounts, are examples of the much broader category of disinformation circulated through social media that has attracted significant academic attention from various disciplines (e.g., Benkler, Faris, & Roberts, 2018; Bennett & Livingston, 2018; Starbird, Arif, & Wilson, 2019). We use the term disinformation to refer to fabricated, manipulated, impostor, or false-context information created by particular agents to cause harm and gain political, financial, social, and psychological advantages (Wardle & Derakhshan, 2017).

According to Tucker *et al.* (2018), a wide range of agents have been involved in producing political disinformation. Having a long history of running disinformation campaigns, Russia has recently extended its disinformation efforts from former Soviet countries like Ukraine (Vanderhill, 2013) to Western democracies like the U.S., using cyber-attacks, trolling, hacking, and promoting socially conservative ideologies (Jamieson, 2020; Ziegler, 2018). Hyperpartisan media sites are also a prominent originator of political disinformation and misinformation and their stories are much more widely shared on Facebook than stories from purely fake news websites (Faris *et al.*, 2017; Silverman *et al.*, 2017). Even mainstream media can recirculate and validate disinformation by inadvertently incorporating disinformation originating from foreign government (Faris *et al.*, 2017; Tucker *et al.*, 2018). Furthermore, communities of interest, i.e., enclaves, tend to consume, sustain, and amplify disinformation (Del Vicario *et al.*, 2016).

Notably, in a deeply connected media system, agents of political disinformation might play overlapping, competitive, and interactive roles in the diffusion process (Tucker *et al.*, 2018). As such, to examine social media disinformation operations, we need to broaden the scope of investigation beyond social media and recognize the role played by the larger information ecosystem (Benkler, Faris, & Roberts, 2018; Jamieson, 2020). This is echoed by Boyd (2017, para. 23), who lamented the deeply flawed aspects of the political communication system that make it susceptible to information manipulation: “Our media, our tools, and our politics are being leveraged to help breed polarization by countless actors who can leverage these systems for personal, economic, and ideological gain.” This leads to our central question about the accumulation of social media following: how do malicious actors take advantage of the dynamics of the political communication system to gain the followings that would enable the wide distribution of disinformation?

Social media following

An account’s social media following, or the number of other accounts that follow an account, can be understood from several perspectives. One way to understand social media following is as an ego network: the set of individuals with which an account has direct ties. Structurally, this ego network facilitates an efficient and effective diffusion of information (Bakshy *et al.*, 2012), with network size positively related to greater and faster information flows (Halberstam & Knight, 2016). When an account’s Twitter followers interact with it by liking, retweeting, or commenting, they potentially promote its tweets to their own ego networks. This not only contributes to messages diffusion, but also brings more followers through the “transitivity of attention” and social recommendation, in which nonfollowers might see and resonate with messages forwarded by their contacts (Golder & Yardi, 2010). As a result, a large follower base creates access to networks of social ties, resources, and influence, generating a “multiplier effect” that strengthens the influence of the central actor (Susarla *et al.*, 2012).

An account’s social media following is also its audience, giving attention to the primary account’s messages (Wells *et al.*, 2020). Social media following constitutes a direct and stable communication channel for an account to reach their audiences without the filtering previously done by mass media. This means that social media accounts can repeatedly tap into the rich mine of attention of their followers in the future, without mediation. Such structural links between communicators and audiences have long been coveted by traditional subscription-based content providers and new entrants to the media system. For example, YouTube channel hosts profit off of their subscriber base (Burgess & Green, 2018), and President Trump has boasted about his direct access to the American people through Twitter (McCaskill, 2017).

The ability to capture attention and assemble a network of followers is directly linked to communicative power in an attention economy (Webster, 2014; Wu,

2017). An account's social media following signals a person's social status and influence (Marwick, 2013). Networked publics often rely on social media's built-in metrics to gauge a user's popularity, authority, and connectivity (e.g., Fu, 2016): a large following increases perceived competence and trustworthiness of an account owner (Jansen *et al.*, 2009), providing social proof of its prominence and value. Additionally, the market logic inherent in social media platforms promotes and perpetuates a hierarchy of influence based on visibility and attention. Being able to command and maintain a large audience's attention is rewarded with more traffic and attention by popularity-based algorithms and ranking systems (Gillespie, 2011). These mechanisms form a positive feedback system, as gaining more followers makes others more likely to follow such an account.

By treating an account's social media following as both an ego network and an audience, we can fully appreciate why social media following is critical for information dissemination and has become an important indicator of power in the networked public sphere. With this in mind, we turn to examining how one's social media following comes together on Twitter.

Network effects in social media

As social media following is an ego network, its formation can be influenced by network dynamics. A network comprised of nodes (i.e., objects) and ties (i.e., connections between nodes). Two dynamics related to tie formation are likely to be present in the development of social media followings, the bottom-up social contagion model driven by average users versus the top-down broadcast model driven by influentials (Goel *et al.*, 2016).

Average users in homophilous networks

Networks of ordinary people—characterized by homophily—can drive information diffusion (Del Vicario *et al.*, 2016). Network homogeneity facilitates diffusion within networks: one is more likely to disseminate information shared by like-minded peers due to similarity or social influence (Aral *et al.*, 2009). This might trigger diffusion beyond the network: Watts and Dodds (2007) argue that a critical mass of ordinary users might be more important than a few influentials in initiating a global cascade. Therefore, if an IRA account wanted to gain followers, it may be effective to first win over a group of susceptible people who, as “early adopters,” “make up the critical mass via which local cascades become global” (Watts & Dodds, 2007, p. 446). Existing research offers some support in this direction: when tweets propagate in Twitter networks via retweeting, the underlying social network structure changes, evidenced in formation (and deletion) of new ties (Hutto, Yardi, & Gilbert, 2013). As an account's tweet flows through a homogenous network via retweeting, endorsement from within the network facilitates growth in its following. Such a process might further increase the account's follower count by improving its engagement metrics, which can influence the followers of those who retweeted its

tweets. That is, benefiting from retweeting by a multitude of average users, an account can be amplified to reach and influence the networks of those who retweeted its message, exponentially expanding its potential reach.

All these suggest that if IRA accounts sought to accrue followers, they would benefit immensely from first enticing like-minded accounts to retweet their content. As IRA accounts turn retweeting users into audience and amplifiers of disinformation, they gain more followers due to this bottom-up network effect. Therefore, we hypothesize that the more retweets IRA accounts gain, the more followers they accumulate (H1); and that the greater is the size of follower networks of those retweeting IRA accounts, the more follower's IRA accounts gain (H2). While H1 and H2 examine the effect of retweets, we still need to ascertain if homogenous networks propelled the retweeting of IRA accounts. Therefore, we ask the question: who retweeted the IRA accounts? (RQ1).

Influential actors

Above and beyond retweets by ordinary people, influentials might also affect follower gain. The two-step flow theory suggests that information diffuses to the masses through introduction, endorsement, and adoption by “opinion leaders” or “influentials” (Katz *et al.*, 2017). Influential actors occupy an advantageous network position where information can be broadcast to especially large and diverse audiences (Burt, 1999). Empirical results show that online diffusion tends to follow this broadcast model (Goel *et al.*, 2016). The same logic might apply to how IRA accounts can gain followings through influential actors on Twitter: not only are more Twitter users likely to be exposed to an IRA account's message through its appearance in an established account's retweet, but they are also more likely to follow this IRA account because of the follow-worthiness signaled by the established account. For example, prominent Trump supporters like Donald Trump Jr. and Jack Posobiec, as well as far-right media figures like Ann Coulter and Dinesh D'Souza retweeted @TEN_GOP, one of the most prominent IRA accounts, repeatedly (O'Brien, 2017), which likely lent authenticity to this impostor account. This is a top-down network effect—influential actors on Twitter with large followings bridge the tie between IRA accounts and their future followers through certification of authenticity. We thus hypothesize that influential actors' retweeting of IRA accounts will increase IRA accounts' followings (H3).

News media amplification

As social media following is an audience, its growth might also be facilitated by news media, which attract audience attention with content. In the high-choice and fragmented news ecology, audiences have limited attention to allocate across various sources (Webster, 2014) and can even avoid news content (Prior, 2007). This attention economy challenges news media, mainstream, and partisan alike,

to produce more engaging content in less expensive and more efficient ways to retain and gain audience.

In response, one notable new journalistic practice is a collaboration with social media (Bossio, 2017). News media attend to audience reactions through social media to understand what attracts attention and to produce clickable stories. News media also incorporate user-generated content with little editing and embed viral social media messages as “vox populi” (Beckers & Harder, 2016; McGregor, 2019). As social media provide a platform for news media to monitor audience attention and source new stories, news media outlets amplify social media. This practice increases the possibility of news media disseminating stories with disinformation.

Just as partisan news media amplify ideologically consistent sites or blogs through linking (Rojecki and Meraz, 2016), they also amplify social media content by producing stories that include ideologically appealing social media messages. A recent report by Silverman *et al.* (2017, para. 1) showed that in 2016 alone, at least 187 new websites on U.S. politics were launched, which “unleashed a golden age of aggressive, divisive political content that reaches a massive amount of people on Facebook.” By trafficking in sensational and outrageous political stories appealing to niche audiences, these sites rode on partisan rage to harvest attention (Baum & Groeling, 2008; Sobierai & Berry, 2011). Given that partisan media amplify social media, we hypothesize that like-minded partisan media’s uptake of IRA tweets leads to an increase in IRA accounts’ followers (H4). However, it is unlikely that opposing partisan media’s uptake of IRA tweets will bring followers to IRA accounts, since partisan media’s audience are probably not interested in the IRA accounts that do not resonate with them. Therefore, we do not hypothesize this relationship.

Under the cut-throat competition for attention, mainstream media might not be immune from amplification of social media content and dissemination of disinformation. In fact, mainstream media recirculated disinformation from foreign governments (Lukito *et al.*, 2020; Tucker *et al.*, 2018). By embedding IRA accounts’ tweets in their news stories, mainstream media might have inadvertently introduced IRA accounts to interested audiences, who became followers of IRA accounts. We thus hypothesize that mainstream media’s uptake of IRA tweets leads to an increase in IRA accounts’ followers (H5).

Method

Data

IRA-related tweets

From a Twitter archive that randomly collects 1% of Global Twitter stream, we collected all tweets between 2015 and 2017 posted by or retweeting the 2,752 Russian IRA accounts identified by Twitter and listed by the House Permanent Select Committee on Intelligence Minority.² Given that the number of retweets that one’s tweets receive is a powerful indicator of influence (Cha *et al.*, 2010), we selected the

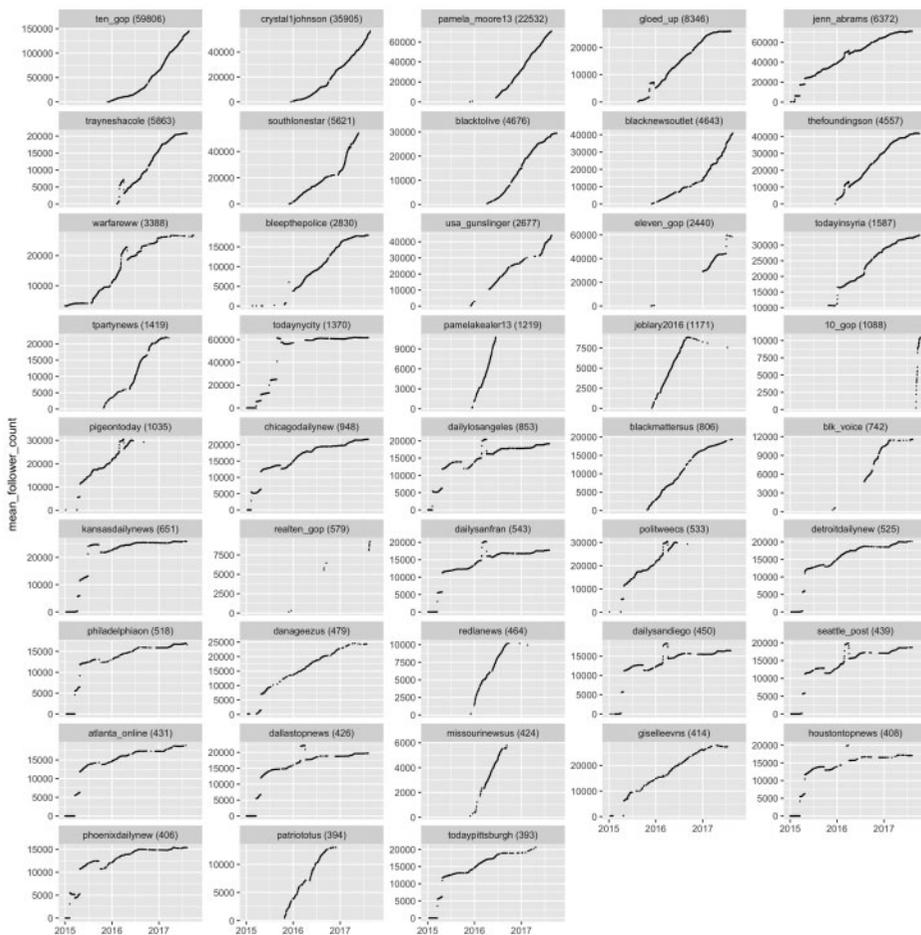


Figure 1 Daily follower counts of the English-tweeting IRA accounts among the top 60 most frequently retweeted IRA accounts (number of total retweets in parenthesis). Note the scales on the Y-axis are different. The top four accounts are selected for further analyses.

top four most retweeted English-speaking IRA accounts during this period, who turned out to be right and left trolls (Linville & Warren, 2020): @TEN_GOP and @Pamela_Moore13 posing as conservatives, and @Crystal1Johnson and @gloed_up imitating liberals. Not only did they amass the highest number of retweets, but their follower growth also seemed more steady and organic than that of the less frequently retweeted IRA accounts (Figure 1; see Appendix SI for details about the selection of accounts).

News media

To identify media outlets that referenced or quoted the IRA accounts, we conducted a search of 218 media outlets using MediaCloud, a media archive that scrapes RSS

data from thousands of news-oriented websites. Our list of outlets was compiled from four collections: the media ecosystem list (Faris *et al.*, 2017), the BuzzFeed hyperpartisan media list (Silverman *et al.*, 2017), the Vox list of IRA-tweet-citing media (Romm & Molla, 2017), and the MediaCloud “U.S. Top Digital Native News” and “U.S. Top Online News” collections. We only retained the media outlets that had a stable stream of publications in MediaCloud, and classified them into mainstream, hyperliberal, and hyperconservative based on existing classifications in those collections (see Appendix SII for details).

A sample of Twitter users who retweeted the IRA accounts

To answer RQ1 about who retweeted IRA accounts, we turned to a unique collection of tweets from 375,725 accounts who followed Donald Trump on Election Day 2016 (3% of his following at that time; that these accounts were following Trump is not relevant to the present analysis) and were classified based on their full following lists (Zhang *et al.*, 2018).³ Each account’s tweets, up to 3,200 per account, were downloaded in March 2017. Each account was categorized using spectral clustering based on similarity in who those accounts followed. We focus on the six politically active categories of accounts: “Trump supporters,” “Far right conservatives,” “Alt-Right,” “Mainstream conservatives,” “Liberals,” and “Mainstream politics.” “Trump supporters” were those who primarily followed Trump’s family members and surrogates. “Far-right conservatives” followed both Trump circle and far-right political figures, like Alex Jones and Mike Cernovich. “Alt-Right” followed white nationalists but not Trump surrogates. “Liberals” followed liberal media, journalists, and politicians. “Mainstream conservatives” followed established political and media accounts on the right, while “Mainstream politics” followed such accounts from the center and the center-left and right. Though not ideal, this sample provides an accurate clustering of election-interested accounts across the political spectrum and contains their tweets around the peak of the IRA disinformation campaign.

Measures

Follower count

We extracted the IRA accounts’ follower counts embedded in the metadata of their tweets and other accounts’ retweets of IRA tweets from the 1% Twitter data. When we had multiple measures of follower count within a day, we averaged the follower counts by day. As can be seen in Figure 1, @TEN_GOP, @Crystal1Johnson, and @gloed_up gained followers steadily throughout 2016 and 2017, whereas @Pamela_Moore13 saw less uptick in follower numbers before mid-2016.

Number of retweets

We computed the daily number of retweets by counting how many times each IRA account was retweeted by other Twitter accounts, as observed in our 1% sample.

Size of retweet network

We defined each account's "daily retweet network" as the set of accounts that retweeted a given IRA account per day, as observed in the 1% sample. We then computed a "size of retweet network" variable by adding together the follower counts of all accounts in retweet network. While there can be overlap between the followers of IRA-retweeting accounts, this measure represents the amount of potential exposure or impressions each account gained each day through retweeting.⁴

Retweeting by influential actors

The verified status Twitter granted to users reflects these accounts' relative prominence on Twitter. We identified all verified Twitter users in the 1% sample who retweeted the IRA accounts and treated them as influential actors. In total, 30 verified users retweeted @TEN_GOP, 2 retweeted @Pamela_Moore13, 63 retweeted @CrystalJohnson, and 3 retweeted @gloed_up. These actors ranged from politicians, activists, media organizations, journalists, pundits, to those in the entertainment and music industries. We computed the number of times each IRA account was retweeted by influential actors, per day.

News media coverage

We tabulated the number of articles referencing an IRA account each day by media type: hyperliberal, mainstream, and hyperconservative. Altogether, @TEN_GOP's tweets appeared in 24, 36, and 484 stories run by hyperliberal, mainstream, and hyperconservative outlets, respectively; for @Pamela_Moore13 those numbers were 10, 8, and 151; @CrystalJohnson's tweets were embedded in 12, 5, and 12 stories from hyperliberal, mainstream, and hyperconservative outlets; and for @gloed_up those numbers were 7, 4, and 4. It is noteworthy that the conservative IRA accounts received much more media attention than the liberal accounts did. Given our sampling from across the news media ecology, this imbalance reflects asymmetry in media uptake of conservative and liberal IRA tweets.

Time series modeling

The data collection process above yielded daily counts of the following variables for each IRA account: retweet count, size of retweet network, number of retweeting by influential actors, number of hyperconservative media stories, number of hyperliberal media stories, number of mainstream media stories, and follower count. For missing data in follower count, retweet count, and size of retweet network, we applied data imputation techniques, including linear interpolation (for 1-, 2-, or 3-day gaps) and forecasting (for gaps longer than 3 days).

A vector autoregression (VAR) model was used to determine the time-ordered relationship between the aforementioned variables. VAR captures the interdependencies between endogenous variables by fitting a multivariate time series regression of each endogenous variable on lags of itself and lags of all other endogenous

variables. To account for all potential relationships, we treated every variable as endogenous. To use a VAR model, we first differenced the time series to identify non-stationary time series, as integrated components may produce incorrect estimates (Box-Steffensmeier *et al.*, 2014). The VAR models were then used to conduct Granger causality tests, which better estimate whether lags of one variable can be used to forecast another variable longitudinally (Groshek, 2011). We used Wald two-way Granger causality tests, a bivariate test of causality. Impulse response functions (IRFs) from the VAR models provide further information regarding the longer-term effects of one variable (X) on another (Y), by testing the impact or “shock” of X on Y (Swanson & Granger, 1997). IRFs are particularly useful for determining the statistical significance, magnitude, and temporal pattern of one variable to a one standard deviation increase in another variable, while controlling for other variables in the model. As customarily done in time series analysis, we present Wald two-way Granger causality tests and the relevant IRFs.

Results

What drives follower gain?

H1 and H2 hypothesize that large numbers of retweets and retweet networks would lead to growth in IRA accounts’ followers. Wald two-way Granger causality tests (Table 1) show that across all four IRA accounts, both the number of retweets and the size of daily retweet networks Granger caused follower gain. IRFs for all four accounts (Figures 2–5), which further control for other variables in the model, corroborate the effect of retweets. For @TEN_GOP, a one standard deviation increase in retweets ($SD = 73$) would increase its followers by about 70 at time $t + 1$ and about 150 at time $t + 3$.⁵ Similarly, a one standard deviation increase in retweets of @Pamela_Moore13 ($SD = 38$) and @Crystal1Johnson ($SD = 65$) would increase their followers by around 20 at time $t + 1$ and around 70 at time $t + 3$. For @gloed_up, a one standard deviation increase in retweets ($SD = 16$) would increase its followers by about 6 at time $t + 1$ and around 17 at time $t + 2$. These results provide strong support for H1. IRFs partially confirm the effect of the retweet network size (H2). For @TEN_GOP and @Crystal1Johnson, retweet network size increased the IRA account’s follower account, but the duration and magnitude of the effect varied by account. However, the 95% confidence interval band of the shock did not cross 0 for @Pamela_Moore13 and @gloed_up.

H3 hypothesizes that influential actors’ retweeting of IRA accounts would increase IRA accounts’ Twitter followings. Granger causality tests demonstrate that influential actors’ retweeting of IRA messages led to a statistically significant growth in followers for @TEN_GOP and @Crystal1Johnson. However, IRFs shows the 95% confidence interval band of the shock cleanly above 0 only for @TEN_GOP. A one standard deviation increase of influential actors’ retweeting of @TEN_GOP ($SD = 0.2$) would increase its followers by about 70 at time $t + 2$. Therefore, H3 receives

Table 1 Two-Way Granger Causality Wald tests

Relationship	Conservative Accounts				Liberal Accounts			
	@TEN_GOP		@Pamela_Moore13		@Crystal1Johnson		@gloed_up	
	<i>F</i>	<i>p</i> -value	<i>F</i>	<i>p</i> -value	<i>F</i>	<i>p</i> -value	<i>F</i>	<i>p</i> -value
Number of retweets → followers	58.726	0.000***	104.280	0.000***	66.932	0.000***	6.979	0.008**
Size of retweet network → followers	45.783	0.000***	58.441	0.000***	52.654	0.000***	5.604	0.018*
Retweeting by influential actors → followers	38.950	0.000***	0.453	0.502	13.838	0.000***	0.455	0.500
Hyperconservative media uptake → followers	20.832	0.000***	6.439	0.012*	0.242	0.623	0.610	0.435
Hyperliberal media uptake → followers	9.436	0.002**	0.000	0.988	0.863	0.353	0.141	0.708
Mainstream media uptake → followers	10.100	0.002**	4.981	0.026*	0.572	0.450	0.045	0.832
Followers → number of retweets	0.686	0.408	0.261	0.610	0.000	0.983	0.124	0.725
Followers → size of retweet network	1.107	0.293	0.231	0.631	0.401	0.527	0.151	0.698
Followers → retweeting by influential actors	0.135	0.713	0.018	0.894	0.260	0.611	0.000	0.993
Followers → hyperconservative media uptake	0.476	0.490	0.268	0.605	0.033	0.856	0.001	0.971
Followers → hyperliberal media uptake	0.084	0.773	0.009	0.924	0.087	0.768	0.012	0.914
Followers → mainstream media uptake	0.298	0.586	0.066	0.798	0.001	0.982	0.000	0.988

p* < .05, *p* < .01, ****p* < .001

partial support. Despite the inconclusiveness of the results for @Pamela_Moore13, @Crystal1Johnson, and @gloed_up, it is notable that retweeting by influential American communicators enhanced the following of @TEN_GOP, above and beyond the contributions of retweets and retweet network in the model (see Appendix SIII for the verified accounts that retweeted @TEN_GOP, the majority of whom seemed to be conservatives).

H4 hypothesizes a positive effect of like-minded partisan media’s uptake of IRA tweets on IRA accounts’ follower gain. Granger causality tests show that, for the two conservative IRA accounts, hyperconservative media’s quoting their tweets did Granger cause their follower growth. However, when controlling for other factors in the model using IRFs (Figures 2 and 3), an increase in conservative media’s attention to @TEN_GOP or @Pamela_Moore13 would not significantly impact its follower count. This suggests that conservative media’s influence might be mediated, which is examined below. For the two liberal IRA accounts, the appearance of their

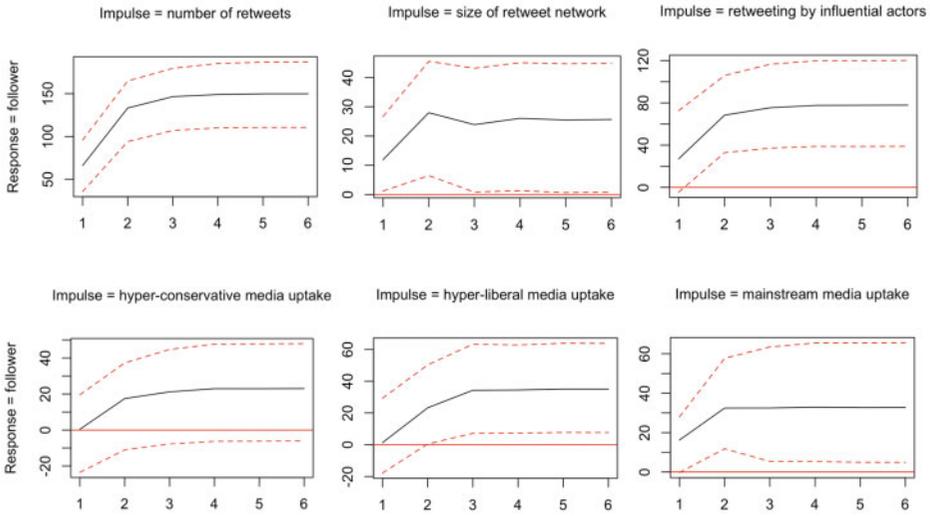


Figure 2 IRFs for @TEN_GOP (conservative), 95% bootstrap confidence interval (dotted ribbon), 100 runs. Effects are cumulative.

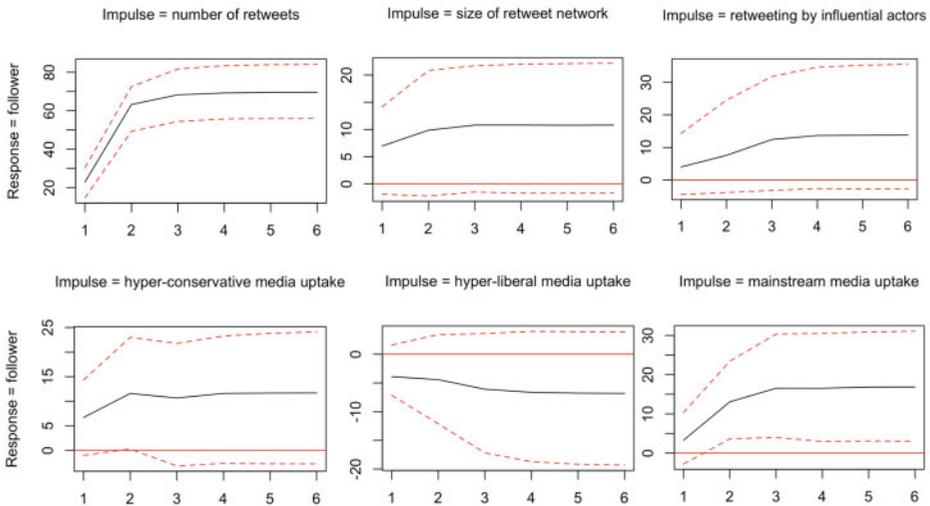


Figure 3 IRFs for @Pamela_Moore13 (conservative), 95% bootstrap confidence interval (dotted ribbon), 100 runs. Effects are cumulative.

tweets in hyperliberal media’s stories did not have any effect on follower growth, both when testing for Granger causality or stimulating IRFs in liberal media quoting. Based on such evidence, H4 is not supported.

H5 concerns mainstream media uptake and follower growth and receives partial support. Mainstream media’s uptake of @TEN_GOP’s and @Pamela_Moore13’s

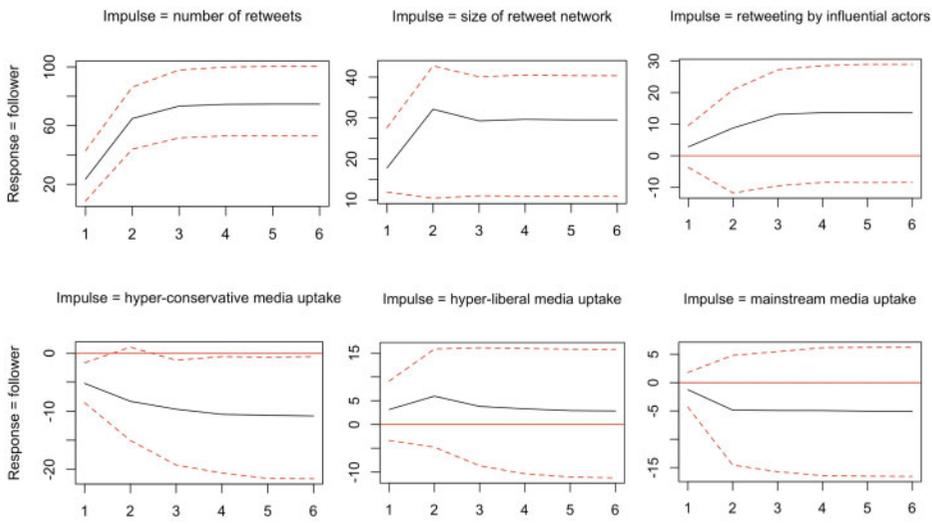


Figure 4 IRFs for @Crystal1Johnson (liberal), 95% bootstrap confidence interval (dotted ribbon), 100 runs. Effects are cumulative.

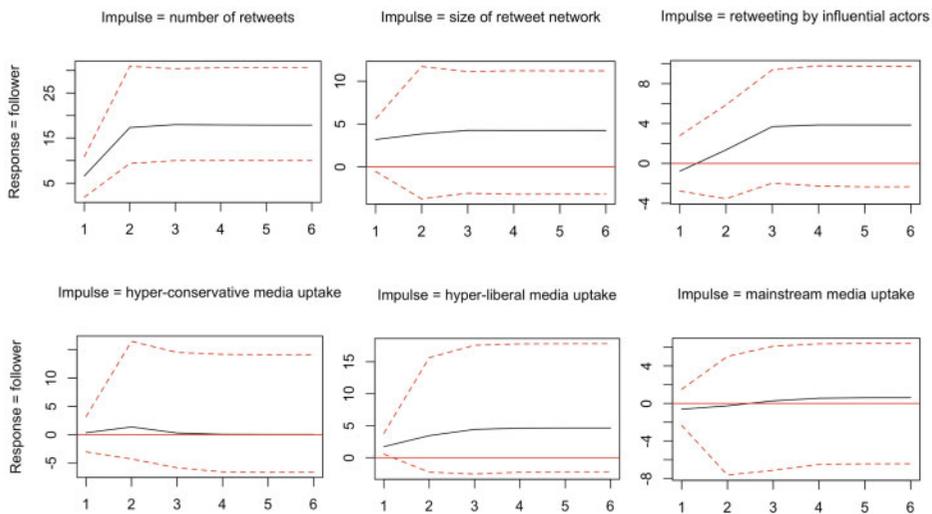


Figure 5 IRFs for @gloed_up (liberal), 95% bootstrap confidence interval (dotted ribbon), 100 runs. Effects are cumulative.

tweets Granger caused their followings to increase, with IRFs confirming this result. A one standard deviation increase in mainstream media’s coverage of @TEN_GOP (SD = 0.3) would increase its follower by around 30 at time $t + 2$; for @Pamela_Moore13 (SD = 0.2), it is about 15 at time $t + 3$. Given the significance of this effect, and the lack of (direct) relationship between hyperconservative media

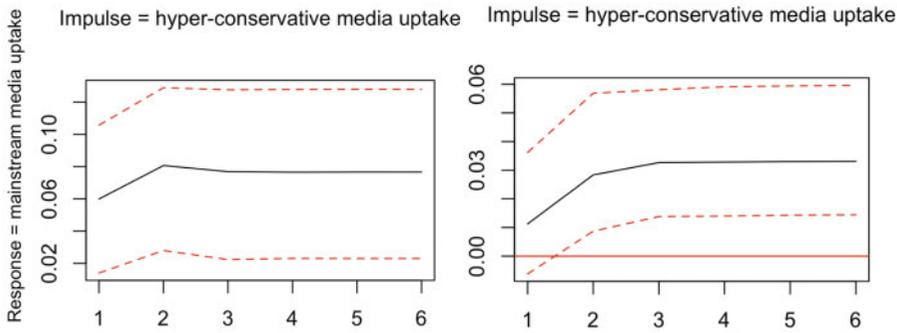


Figure 6 IRFs of hyperconservative media uptake on mainstream media uptake for @TEN_GOP (on the left) and @Pamela_Moore13 (on the right). Effects are cumulative.

and follower growth, we investigated the temporal relationship between mainstream and hyperconservative media uptake. We find that hyperconservative media's coverage of Granger caused mainstream media coverage for both @TEN_GOP and @Pamela_Moore13, suggesting that hyperconservative media's effect on their follower growth was mediated by mainstream media coverage. IRFs show that a standard deviation increase in hyperconservative media's coverage of @TEN_GOP ($SD = 1.3$) would increase mainstream media's coverage by about 0.08 at time $t + 2$; for @Pamela_Moore13 ($SD = 0.7$), the impact is weaker, about 0.03 at time $t + 2$ (Figure 6). However, mainstream media uptake, or any media uptake, did not have any effect on these two liberal accounts' follower growth. The different patterns of media amplification for the conservative and liberal IRA accounts probably result from the overall imbalance of media attention: there was simply not enough coverage of the liberal IRA accounts for any signal to be detected. This qualification, though, represents a key finding: in the case of the "liberal" IRA accounts, there was little possibility of follower growth deriving from media coverage, as they were not picked up by media.

Who retweeted IRA accounts?

RQ1 asks whether homogenous networks propelled the retweeting of IRA accounts, which they relied on to amass followings. Given the key role retweeting played in building IRA accounts' followers, a deeper understanding of this mechanism is critical. The present analysis is concerned with who retweeted the four IRA accounts based on a sample of political relevant Twitter users, whose political leanings are detected through clustering of their full friend lists.

The pattern of retweeting from the sample is clear: those retweeting the two conservative IRA accounts were mainly from the most extreme conservative groups—far-right conservatives and the alt-right, whereas liberals predominantly retweeted the two liberal accounts (Table 2). Far-right conservatives and the alt-right together made up over 70% of the sample who retweeted @TEN_GOP and contributed to

Table 2 Retweeting of IRA accounts by six categories of non-IRA Twitter accounts

	Conservative Accounts						Liberal Accounts					
	@TEN_GOP			@Pamela_Moore13			@Crystal1Johnson			@gloed_up		
	percentage_ users	percentage_ retweets	average_ retweets_ per_user	percentage_ users	percentage_ retweets	average_ retweets_ per_user	percentage_ users	percentage_ retweets	average_ retweets_ per_user	percentage_ users	percentage_ retweets	average_ retweets_ per_user
Trump supporters	13.26	9.74	9	11.03	8.37	4	3.33	1.67	1	2.82	2.20	1
Far right conservatives	41.89	55.51	15	49.39	58.22	6	5.29	3.47	1	4.23	3.30	1
Alt-right	30.51	26.96	10	29	25.88	4	5.88	3.42	1	4.93	3.85	1
Mainstream conservatives	10.57	6.94	8	8.28	6.61	4	6.86	5.09	2	7.04	6.04	1
Mainstream politics	2.43	0.50	2	1.40	0.46	2	19.22	15.59	2	19.72	18.13	1
Liberals	1.33	0.35	3	0.90	0.46	3	59.41	70.75	3	61.27	66.48	1
Total	100	100	\	100	100	\	100	100	\	100	100	\

over 80% of the retweets. If we consider the relatively small number of these accounts within our sample, their hyperactivity is even more evident: the far-right conservatives and the alt-right retweeted @TEN_GOP 15 and 10 times per user, respectively. A similar yet less extreme pattern is visible for @Pamela_Moore13.

Accounts that retweeted @Crystal1Johnson and @gloed_up seem to be the partisan mirror opposite of those who retweeted @TEN_GOP and @Pamela_Moore13, although overall there were far fewer accounts retweeting @Crystal1Johnson and @gloed_up, a product of both those accounts' more modest success and this sample's bias toward conservatives. Nevertheless, liberals accounted for about 60% of the sample who retweeted @Crystal1Johnson and @gloed_up and contributed to around 70% of the retweets.

The partisan retweeting of the four IRA accounts' tweets is clear. IRA messages appealed to partisan audiences, who were galvanized into retweeting IRA tweets. The result is not just that these four IRA accounts gained more retweets and message exposure. The frantic retweeting of their tweets directly contributed to IRA accounts' followings. Given that IRA accounts received the most retweets from polarized partisan enclaves, it can be concluded that partisans are vulnerable to both consuming *and disseminating* political disinformation, particularly when disinformation is highly opinionated and emotionally provocative (see Appendix SIV for the analysis of IRA tweets).

Discussion

Our story of how the four most successful IRA Twitter accounts built their followings reveals several dynamics at work in the polarized, asymmetrical, attention-driven, and disinformation-laden political communication system. First, we confirmed previous work demonstrating the significance of networks effects internal to social media in driving follower gain. It was mainly the bottom-up network effect where the accumulation of massive retweets powerfully increased followers. However, for one account, the top-down network effect was also at work; notably, being retweeted by verified users provided an additional boost to @TEN_GOP's follower gain. With their messages widely retweeted, these IRA accounts transformed short-term communication success into long-term assets: engagement-worthiness and newsworthiness manifested in the metric of follower count.

In addition to retweeting being immensely effective, our results demonstrate that much of this retweeting behavior occurred within rather extreme ideological enclaves, which is consistent with [Bail et al.'s \(2020\)](#) finding. The impression of disinformation actors "fooling" users at random does not appear to be an accurate one. As shown in the analysis of the retweeting patterns of politically engaged Twitter users, all four IRA accounts were primarily retweeted by ideologically consistent clusters. Embedded in ideological enclaves, partisan Twitter users were likely to engage with partisan messages aligning with their viewpoints. As partisan audience became interpreters and amplifiers of disinformation on social media, they directly

contributed to IRA accounts' accumulation of followers. Partisans' selective interaction with the four IRA accounts also shows their susceptibility to deceptive political practices, especially with strategic attempts of the Russian government to influence the U.S. election outcome.

Second, we demonstrate that amplification by mainstream and like-minded hyperpartisan news media were effective mechanisms for conservative IRA accounts to gain Twitter followers. Hyperpartisan news media and mainstream news media might have brought IRA accounts to the attention of more potential followers, who would follow them due to mere exposure and/or the credential and follow-worthiness associated with appearing in news. It is noteworthy that the effect of hyperconservative media is mediated by mainstream media, speaking to mainstream media's direct and indirect amplification of IRA accounts and their role in exacerbating or curtailing the influence of disinformation. Here, the results are an echo of [Benkler et al.'s \(2018\)](#) notes on the relationship between Breitbart and centrist news organizations in their coverage of scandals related to the Clinton Foundation.

News media's embedding of IRA tweets accentuates the problem of journalism under the attention economy, where content thought to engage audience and command attention is highly sought after ([Karpf, 2016](#)). Embedding Twitter contents that attract eyeballs or confirm preferences, media are likely to be rewarded with traffic and revenue. This is especially true with online partisan media that operate under low budgets and need to build a loyal audience ([Hamilton, 2004](#)). Mainstream, more reliably journalistic media are not free from the pressure to deliver arresting content quickly and frequently. This is not to say that news media feed their audience with outrageously false content, flagrantly violating traditional news values. Instead, media tend to use IRA tweets for two purposes—as sources or providers of evidence, often through visual content, and as part of a social media proxy for public opinion, colloquially known as “vox populi” quotes ([Lukito et al., 2020](#)). For news media, picking up such messages seems not only economically viable but also adherent to newsworthiness. As such, our results show one way in which such an economy of news media opens a door to manipulation and exploitation from other actors in the system, who “create, tap, or steer information flows in ways that suit their goals” ([Chadwick, 2017](#), p. 4).

Both sets of processes—retweeting by partisans on Twitter and amplification by partisan and mainstream news media—reveal much about networks and audiences coalescing on social media through a coordination between social and news media. By energizing ideological enclaves on Twitter, IRA accounts manage to “prove” their status and prominence ([Marwick, 2013](#)), attracting attention and coverage from wider circles of news media, which seek news hooks. By appearing in news media, partisan, and mainstream alike, IRA accounts gain credentials that might attract more followers. This result aligns with existing research demonstrating that the accumulation of social media audience is an iterative process involving news media exposure and social media engagement ([Wells et al., 2020](#)). This speaks to the fluidity of attention and information flows in a highly

interconnected hybrid media system. Any actor aiming to accumulate audience needs to tap into such fluidity.

Furthermore, our results paint a picture of IRA accounts taking advantage of the dynamics of and fractures in the American political communication system to gain followers on Twitter, and in the process shed new light on the contours of our contemporary “disinformation order” (Bennett & Livingston, 2018). The strategies, which allow regular Americans or political groups to gain followings, can be manipulated to create a disinformation order. Partisan actors on social media helped build disinformation actors’ influence by amplifying the disinformation produced by IRA accounts through retweeting. News media, driven by the attention-seeking impulse, did so by embedding IRA tweets directly in news. Disinformation actors took advantage of these weaknesses to insert themselves into partisan networks and news media discourses, and to develop unsuspecting followings. This practice aligns with previous research on the partisan nature of IRA-produced content (Farkas & Bastos, 2018) and the IRA’s attempts to engage with American media outlets (Lukito *et al.*, 2020).

However, the mechanisms for accumulation of followers are not symmetrical for the conservative and liberal IRA accounts, echoing Benkler *et al.*’s (2018) argument about the asymmetry between the partisan media ecosystems and their differing susceptibility to disinformation (see also Badawy *et al.*, 2019). While the liberal IRA accounts accumulated followers through activating partisan networks on social media, the conservative IRA accounts amassed followers through activating the entire conservative media ecosystem, where the conservative networks and conservative pundits on Twitter as well as the hyperconservative media form a “propaganda feedback loop” that amplified conservative IRA accounts and tweets. The analysis of tweets from conservative IRA accounts like @TEN_GOP and @Pamela_Moore13 shows that their messaging aligned with issues on Trump’s agenda, attacks of Hillary Clinton, and suppression of Democratic voters (Jamieson, 2020; Appendix SIV). As such propagandic content appealed to the conservatives, it reverberated through the conservative media ecosystem.

We recognize one potential limitation in our analysis: we were not able to separate organic mechanisms, such as news media amplification and network dynamics, from internally-coordinated mechanisms: the IRA campaign might have IRA accounts follow and/or retweet each other to appear more authentic, engagement-, and news-worthy in the eyes of genuine Twitter users, including journalists. One finding alleviates this concern—we tracked IRA accounts in the Twitter datasets and found scant evidence for other IRA accounts retweeting the four successful IRA accounts (Appendices SV and SVI).

Another potential limitation might be the selection of accounts. We examined four out of the several thousand IRA accounts. There might be other dynamics that we are not able to capture using this quantitative case study strategy (see Appendix SVII for additional analysis). However, since these accounts were the most successful ones (in terms of the frequency of being retweeted, follower base, and media

exposure) and provided rich data, they merit more attention than other less successful accounts, fitting the purpose and scope of this study.

Conclusion

By taking a media system approach to study network/audience formation and examining it in the context of a major disinformation campaign, our study highlights a number of challenges facing political communication and society at large: the profound interconnectedness of the hybrid and asymmetric media system, the deep political polarization, the attention-focused media logics, and the disinformation that can flourish in such an environment.

Importantly, our findings emphasize that “disinformation” is not a phenomenon that is easily separated—conceptually or practically—from “legitimate” political expression in the polarized and fragmented media environment in which it now flourishes. Strongly opinionated individuals are susceptible to confirmation bias and tend to engage with proattitudinal information. Profiting off such partisan inclinations, news media serve up ideologically palatable and attention arresting content, while enlisting social media as a cost-effective information source for content production and a convenient platform for content distribution and diffusion (Munger, 2020). All these factors contribute to a fractured media system and polarized publics, which opens up ample opportunities for bad actors to exploit the system vulnerabilities and the divided publics. But as our results show, the level of susceptibility to disinformation actors and their influence differs across the political left and right—the structure of the conservative media ecosystem sets it up readily for identity confirming disinformation to reverberate.

Ultimately, political disinformation and its underlying social and political contexts only reinforce each other. While some scholars propose fact checks and platform moderation as the solutions to information disorder (Lazer *et al.*, 2018), this study shows that disinformation is not so much a purely informational or technological problem. It is a social and political problem—the success of disinformation actors to spread messages and gain influence manifests the deep political division in American society, the fractured political communication system driven by attention and profit, and the particular susceptibility of the conservative media ecosystem to disinformation. Flagging false information or increasing media literacy for news consumers, thus, only provide part of the answer. Similarly, raising the awareness of individual journalists about potential disinformation operations is not enough, not only because such activities cannot be accurately detected without the participation of social media platforms, but also due to attention and profit incentives influencing journalistic routines, particularly of partisan media.

Furthermore, this study demonstrates the fluidity and connectivity of the hybrid media system and suggests that to assemble social media audience, one has to rely on both social media network dynamics and news media amplification. This, however, points to some disheartening consequences. As the assemblage of social

media following might rely on stirring up the passions of ideological enclaves and producing arresting content, those seeking to gain a diverse coalition of audience and to speak across partisan divisions without the temptation to instigate or provoke may find themselves unfollow-worthy. Such a tendency might further aggravate political polarization and lead to segmentation of audiences.

Overall, by situating the Russian IRA's disinformation campaign in the American political communication system, this article points to the susceptibility of polarized publics and attention-driven media to disinformation operations as well as the large social and political contexts of political disinformation. Our empirical contributions present a portrayal of disinformation as a product of not only foreign interlopers, but also the context of decentralized, fragmented, and asymmetrically polarized media, strongly opinionated, and confirmation seeking audiences, and an attention economy optimized for engagement rather than accuracy.

Acknowledgments

The authors are grateful for helpful feedback from the reviewers and the editor. We also appreciate the resource support from the Mass Communication Research Center at UW-Madison.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Supporting Information

Additional Supporting Information may be found in the online version of this article.

Please note: Oxford University Press is not responsible for the content or functionality of any [supplementary materials](#) supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

Notes

1. A Slate journalist who reflected on how @TEN_GOP wound up in one of her articles specifically referenced the account's high follower number as a measure of popularity and credibility that led her to quote its tweet <https://slate.com/technology/2018/03/how-a-russian-ten-gop-tweet-wound-up-in-slate.html>.
2. The accounts came from House Exhibit B in Permanent Select Committee on Intelligence Minority (https://democrats-intelligence.house.gov/uploadedfiles/exhibit_b.pdf; see full memo here: <https://democrats-intelligence.house.gov/uploadedfiles>). One key characteristic of the 1% data is that the follower count of a handle was retrieved around the time of a given post (as tweets were downloaded on a daily basis). However, Twitter's recently

- released IRA data only provide a static number of follower count at around the revelation of IRA operation, which makes our analysis IMPOSSIBLE to do (Appendix SVIII).
3. Ideally, we would construct a random sample from all those who retweeted IRA account in the 1% Twitter data, download their social network information, and cluster them. However, this process would be computationally costly and time-consuming, which is a project on its own and out of the scope of this article. Hence, we rely on this convenient sample where accounts retweeting the four IRA accounts were classified with relatively high accuracy.
 4. There can be some or significant overlap between IRA-retweeting accounts' followers. However, the number of impressions should be positively correlated with the likelihood to follow. However, this measure is not able to distinguish an IRA account's message spreading to a wide and thin audience from it reverberating through a dense community.
 5. The time series of daily follower counts were first differenced for all the four accounts (which became "flatter"). Therefore, this number does not represent the increase in raw follower count, which could be higher.

References

- Aral, S., Muchnik, L., & Sundararajan, A. (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(51), 21544–21549.
- Badawy, A., Addawood, A., Lerman, K., & Ferrara, E. (2019). Characterizing the 2016 Russian IRA influence campaign. *Social Network Analysis and Mining*, *9*(1), 31.
- Badawy, A., Ferrara, E., & Lerman, K. (2018, August). Analyzing the digital traces of political manipulation: The 2016 Russian interference Twitter campaign. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)* (pp. 258–265). IEEE. Retrieved from <https://doi.org/10.1109/ASONAM.2018.8508646>
- Bail, C. A., Guay, B., Maloney, E., Combs, A., Hillygus, D. S., Merhout, F., . . . Volfovsky, A. (2020). Assessing the Russian Internet Research Agency's impact on the political attitudes and behaviors of American Twitter users in late 2017. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(1), 243–250.
- Bakshy, E., Rosenn, I., Marlow, C., & Adamic, L. (2012, April). The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web* (pp. 519–528). ACM. Retrieved from <https://doi.org/10.1145/2187836.2187907>
- Baum, M. A., & Groeling, T. (2008). New media and the polarization of American political discourse. *Political Communication*, *25*(4), 345–365. doi:10.1080/10584600802426965
- Beckers, K., & Harder, R. A. (2016). "Twitter Just Exploded" Social media as alternative vox pop. *Digital Journalism*, *4*(7), 910–920. doi:10.1080/21670811.2016.1161493
- Benkler, Y., Faris, R., & Roberts, H. (2018). *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. New York, NY: Oxford University Press.
- Bennett, W. L., & Livingston, S. (2018). The disinformation order: Disruptive communication and the decline of democratic institutions. *European Journal of Communication*, *33*(2), 122–139. <https://doi.org/10.1177/0267323118760317>
- Bennett, W. L., & Segerberg, A. (2013). *The logic of connective action: Digital media and the personalization of contentious politics*. Cambridge, NY: Cambridge University Press.
- Bode, L., & Vraga, E. (2018). Studying politics across media. *Political Communication*, *35*(1), 1–7.

- Bossio, D. (2017). JOURNALISM AND SOCIAL MEDIA: Practitioners, Organisations and Institutions. In *Journalism and social media* (pp. 133–152). New York, NY: Palgrave Macmillan.
- Bovet, A., & Makse, H. A. (2019). Influence of fake news in Twitter during the 2016 US presidential election. *Nature Communications*, *10*(1), 1–14.
- Box-Steffensmeier, J. M., Freeman, J. R., Hitt, M. P., & Pevehouse, J. C. (2014). *Time series analysis for the social sciences*. Cambridge, NY: Cambridge University Press.
- Boyd, D. (2017, March 2). Google and Facebook can't just make fake news disappear. *Wired*. Retrieved from <https://www.wired.com/2017/03/google-and-facebook-cant-just-make-fake-news-disappear/>
- Burgess, J., & Green, J. (2018). *YouTube: Online video and participatory culture*. Hoboken, NJ: John Wiley & Sons.
- Burt, R. S. (1999). The social capital of opinion leaders. *The Annals of the American Academy of Political and Social Science*, *566*(1), 37–54. <https://doi.org/10.1177/000271629956600104>
- Cha, M., Haddadi, H., Benevenuto, F., & Gummadi, K. P. (2010, May). Measuring user influence in twitter: The million follower fallacy. In Proceedings of the *international AAAI conference on web and social media* (Vol. 4, No. 1).
- Chadwick, A. (2017). *The hybrid media system: Politics and power*. Oxford, UK: Oxford University Press.
- Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., . . . Quattrociocchi, W. (2016). The spreading of misinformation online. *Proceedings of the National Academy of Sciences of the United States of America*, *113*(3), 554–559. <https://doi.org/10.1073/pnas.1517441113>
- Faris, R., Roberts, H., Etling, B., Bourassa, N., Zuckerman, E., & August, B. Y. (2017). *Partisanship, propaganda, and disinformation: Online media and the 2016 US presidential election*. SSRN: Retrieved from <https://ssrn.com/abstract=3019414>
- Farkas, J., & Bastos, M. (2018, July). IRA propaganda on Twitter: Stoking antagonism and tweeting local news. In *Proceedings of the 9th International Conference on Social Media and Society* (pp. 281–285). ACM. <https://doi.org/10.1145/3217804.3217929>
- Freelon, D., & Karpf, D. (2015). Of big birds and bayonets: Hybrid Twitter interactivity in the 2012 presidential debates. *Information, Communication & Society*, *18*(4), 390–406.
- Freelon, D., Marwick, A., & Kreiss, D. (2020). False equivalencies: Online activism from left to right. *Science*, *369*(6508), 1197–1201.
- Freelon, D., McIlwain, C. D., & Clark, M. (2016, February 29). Beyond the hashtags: # Ferguson, # Blacklivesmatter, and the online struggle for offline justice. *Center for Media & Social Impact, American University*, Forthcoming. <http://dx.doi.org/10.2139/ssrn.2747066>
- Fu, J. S. (2016). Leveraging social network analysis for research on journalism in the information age. *Journal of Communication*, *66*(2), 299–313. <https://doi.org/10.1111/jcom.12212>
- Gillespie, T. (2011). Can an algorithm be wrong? Twitter trends, the specter of censorship, and our faith in the algorithms around us. *Culture Digitally*, *19*, 1.
- Goel, S., Anderson, A., Hofman, J., & Watts, D. J. (2016). The structural virality of online diffusion. *Management Science*, *62*(1), 180–196.
- Golder, S. A., & Yardi, S. (2010, August). Structural predictors of tie formation in twitter: Transitivity and mutuality. In *2010 IEEE Second International Conference on Social Computing* (pp. 88–95). IEEE. <https://doi.org/10.1109/SocialCom.2010.22>

- Groshek, J. (2011). Media, instability, and democracy: Examining the Granger-causal relationships of 122 Countries from 1946 to 2003. *Journal of Communication*, 61(6), 1161–1182. <https://doi.org/10.1111/j.1460-2466.2011.01594.x>
- Halberstam, Y., & Knight, B. (2016). Homophily, group size, and the diffusion of political information in social networks: Evidence from Twitter. *Journal of Public Economics*, 143, 73–88. <https://doi.org/10.1016/j.jpubeco.2016.08.011>
- Hamilton, J. (2004). *All the news that's fit to sell: How the market transforms information into news*. Princeton, NJ: Princeton University Press.
- Hjorth, F., & Adler-Nissen, R. (2019). Ideological asymmetry in the reach of Pro-Russian digital disinformation to United States audiences. *Journal of Communication*, 69(2), 168–192.
- Hutto, C. J., Yardi, S., & Gilbert, E. (2013, April). A longitudinal study of follow predictors on twitter. In *Proceedings of the sigchi conference on human factors in computing systems* (pp. 821–830). ACM. <https://doi.org/10.1145/2470654.2470771>
- Jamieson, K. H. (2020). *Cyberwar: how Russian hackers and trolls helped elect a president: what we don't, can't, and do know*. New York, NY: Oxford University Press.
- Jansen, B. J., Zhang, M., Sobel, K., & Chowdury, A. (2009). Twitter power: Tweets as electronic word of mouth. *Journal of the American Society for Information Science and Technology*, 60(11), 2169–2188. <https://doi.org/10.1002/asi.21149>
- Karpf, D. (2016, June 1). The Clickbait Candidate, *The Chronicle of Higher Education*. Retrieved from <http://www.chronicle.com/article/The-Clickbait-Candidate/236815>
- Katz, E., Lazarsfeld, P. F., & Roper, E. (2017). *Personal influence: The part played by people in the flow of mass communications*. New York, NY: Routledge.
- Klinger, U., & Svensson, J. (2015). The emergence of network media logic in political communication: A theoretical approach. *New Media & Society*, 17(8), 1241–1257. <https://doi.org/10.1177/1461444814522952>
- Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., . . . Zittrain, J. L. (2018). The science of fake news. *Science*, 359(6380), 1094–1096. <https://doi.org/10.1126/science.aao2998>
- Linville, D. L., & Warren, P. L. (2020). Troll factories: Manufacturing specialized disinformation on Twitter. *Political Communication*, 37(4), 1–21.
- Lukito, J., Suk, J., Zhang, Y., Doroshenko, L., Kim, S. J., & Wells, C. (2020). The wolves in sheep's clothing: How Russia's internet research agency tweets appeared in US news as Vox Populi. *The International Journal of Press/Politics*, 25(2), 196–216.
- Marwick, A. E. (2013). *Status update: Celebrity, publicity, and branding in the social media age*. New Haven, CO: Yale University Press.
- Marwick, A., & Lewis, R. (2017). *Media manipulation and disinformation online*. New York, NY: Data & Society Research Institute.
- McCaskill, N. (2017, October 20). Trump credits social media for his election. *Politico*. Retrieved from <https://www.politico.com/story/2017/10/20/trump-social-media-election-244009>
- McGregor, S. C. (2019). Social media as public opinion: How journalists use social media to represent public opinion. *Journalism*, <https://doi.org/10.1177/1464884919845458>
- McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1), 415–444.
- Metaxas, P. T., & Mustafaraj, E. (2012). Social media and the elections. *Science*, 338(6106), 472–473. doi:10.1126/science.1230456

- Munger, K. (2020). All the news that's fit to click: The economics of clickbait media. *Political Communication*, 37(3), 376–397.
- O'Brien, L. (2017, November 01). Twitter ignored this Russia-controlled account during the election. team Trump did not. *Huffington Post*. Retrieved from https://www.huffpost.com/entry/twitter-ignored-this-russia-controlled-account-during-the-election_n_59f9bdcbe4b046017fb010b
- Prior, M. (2007). *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. New York, NY: Cambridge University Press.
- Ratkiewicz, J, Conover, M. D, Meiss, M, Gonçalves, B, Flammini, A., & Menczer, F. M. (2011, July). Detecting and tracking political abuse in social media. In *Fifth international AAAI conference on weblogs and social media*.
- Rojecki, A., & Meraz, S. (2016). Rumors and factitious informational blends: The role of the web in speculative politics. *New Media & Society*, 18(1), 25–43.
- Romm, T., & Molla, R. (2017, November 4). Here's a longer list of news organizations that cited Russia-linked Twitter accounts. *Vox*. Retrieved from <https://www.vox.com/2017/11/4/16606188/twitter-russia-troll-news-citation-list>
- Silverman, C., Lytvynenko, J., Vo, L. T., & Singer-Vine, J. (2017, August 8). Inside the partisan fight for your news feed. *BuzzFeed News*. Retrieved from <https://www.buzzfeednews.com/article/craigsilverman/inside-the-partisan-fight-for-your-news-feed>
- Sobierai, S., & Berry, J. M. (2011). From identity to outrage: Political discourse in blogs, talk radio, and cable news. *Political Communication*, 28(1), 19–41. <https://doi.org/10.1080/10584609.2010.542360>
- Starbird, K., Arif, A., & Wilson, T. (2019). Disinformation as collaborative work: Surfacing the participatory nature of strategic information operations. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), 1–26.
- Susarla, A., Oh, J. H., & Tan, Y. (2012). Social networks and the diffusion of user-generated content: Evidence from YouTube. *Information Systems Research*, 23(1), 23–41.
- Swanson, N. R., & Granger, C. W. (1997). Impulse response functions based on a causal approach to residual orthogonalization in vector autoregressions. *Journal of the American Statistical Association*, 92(437), 357–367. <https://doi.org/10.1080/01621459.1997.10473634>
- Tucker, J., Guess, A., Barberá, P., Vaccari, C., Siegel, A., Sanovich, S., . . . Nyhan, B. (2018, March 19). *Social media, political polarization, and political disinformation: A review of the scientific literature*. William and Flora Hewlett Foundation. <http://dx.doi.org/10.2139/ssrn.3144139>
- Tufekci, Z. (2013). “Not this one” social movements, the attention economy, and microcelebrity networked activism. *American Behavioral Scientist*, 57(7), 848–870.
- Tufekci, Z., & Wilson, C. (2012). Social media and the decision to participate in political protest: Observations from Tahrir Square. *Journal of Communication*, 62(2), 363–379.
- Vanderhill, R. (2013). *Promoting authoritarianism abroad*. Boulder, CO: Lynne Rienner Publishers.
- Vargo, C. J., Guo, L., & Amazeen, M. A. (2018). The agenda-setting power of fake news: A big data analysis of the online media landscape from 2014 to 2016. *New Media & Society*, 20(5), 2028–2049.
- Wardle, C., & Derakhshan, H. (2017). Information disorder: Toward an interdisciplinary framework for research and policy making. *Council of Europe report, DGI (2017)*, 9.
- Watts, D. J., & Dodds, P. S. (2007). Influentials, networks, and public opinion formation. *Journal of Consumer Research*, 34(4), 441–458. <https://doi.org/10.1086/518527>

- Webster, J. G. (2014). *The marketplace of attention: How audiences take shape in a digital age*. Cambridge, MA: MIT Press.
- Wells, C., Zhang, Y., Lukito, J., & Pevehouse, J. C. (2020). Modeling the formation of attentive publics in social media: The case of Donald Trump. *Mass Communication and Society*, 23(2), 181–205.
- Wu, T. (2017). *The attention merchants: The epic scramble to get inside our heads*. New York, NY: Vintage.
- 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), 3161–3182.
- Ziegler, C. E. (2018). International dimensions of electoral processes: Russia, the USA, and the 2016 elections. *International Politics*, 55(5), 557–574.