International Journal of Communication 15(2021), 814–839 1932–8036/20210005

Not the Bots You Are Looking For: Patterns and Effects of Orchestrated Interventions in the U.S. and German Elections

OLGA BOICHAK
University of Sydney, Australia

JEFF HEMSLEY
Syracuse University, USA

SAM JACKSON
University at Albany, USA

REBEKAH TROMBLE
George Washington University, USA

SIKANA TANUPABRUNGSUN
Microsoft, USA

Zooming in on automated and semiautomated social actors created to influence public opinion on social media, we employ a novel analytic approach to identify patterns of inauthentic behavior across election campaigns on Twitter. Comparing two recent national election campaigns, the 2016 U.S. presidential election and the 2017 German federal election, we analyze patterns and effects of orchestrated intervention in political discourse on Twitter. Focusing on two main aspects of information flows—scale and range—we find that orchestrated interventions help amplify, but not diffuse, the candidates' messages, mostly failing to reach new audiences in the process. This study adds an information diffusion perspective to a growing body of literature on computational propaganda, showing that although false amplification is quite effective in increasing the scale of information events, in most cases the information fails to reach new depths.

Keywords: computational propaganda, bots, information diffusion, elections, Twitter, U.S., Germany

Olga Boichak: olga.boichak@sydney.edu.au
Jeff Hemsley: jjhemsle@syr.edu
Sam Jackson: sdjackson@albany.edu
Rebekah Tromble: rtromble@email.gwu.edu
Sikana Tanupabrungsun: sikana.tanu@gmail.com
Date submitted: 2020-03-16

In the past few years, major national elections around the world saw external actors seeking to influence the course of election campaigning online (e.g., Calabresi, 2017; Greenberg, 2017). Several studies provide evidence of external actors meddling in the U.S. presidential election campaign through media; rather than pursuing particular electoral outcomes, their alleged goal was to disrupt political deliberation and decrease public trust in democratic institutions (Kollanyi, Howard, & Woolley, 2016; Pazzanese, 2017). Research has found that as much as a fifth of all traffic related to the 2016 U.S. presidential election on Twitter was orchestrated (Bessi & Ferrara, 2016). The Brexit referendum in 2016 and the French presidential election in 2017 both saw large amounts of misinformation shared via social media (Greenberg, 2017). Orchestrated campaigns have also been deployed to counteract criticism in Mexico during elections (Salge & Karahanna, 2016), and have been used by oppositional parties in Venezuela to attack the regime and spread misinformation (Forelle, Howard, Monroy-Hernández, & Savage, 2015). One of the main tools for interference that researchers have pointed to is the use of automated and semiautomated accounts (i.e., *bots*; Ferrara, 2017; Howard & Kollanyi, 2016; Kollanyi et al., 2016).

In online spaces, orchestrated political activity is rarely fully automated—a *bot* is an ambiguous concept that may encompass a range of structural and functional forms, and serve various purposes (Gorwa & Guilbeault, 2018). Though most social bots share an alleged aim to impersonate human users (Abokhodair, Yoo, & McDonald, 2015; Anderson et al., 2017), the accounts may be run by a human, an algorithm, or various combinations of both (Anderson et al., 2017). Bots are also known to fulfill a range of information diffusion functions, from search engine optimization and information retrieval (Gorwa & Guilbeault, 2018) to aggregation and/or amplification of political messages that might or might not be part of strategic disinformation campaigns (Boichak, Jackson, Hemsley, & Tanupabrungsun, 2018; Ferrara, 2017; Giglietto, Iannelli, Rossi, & Valeriani, 2016). Recognizing the structural and functional diversity of actors under this umbrella term, we chose to focus this article on inauthentic, "nonorganic" behavior patterns, which may include, but are not limited to, automated and semiautomated actors also known as bots. In doing so, we assume the participants of orchestrated interventions share a motive to alter, disrupt, or otherwise intervene in the course of election campaigns on social media, regardless of their structural or functional forms.

This study's main contributions are twofold. First, we employ a novel analytic approach—operationalizing information flows as retweet events (RTEs)—to identify patterns of orchestrated behavior across election campaigns on Twitter. By collecting large amounts of data during the campaign period, using established methods for estimating the likelihood of a particular account engaging in suspicious activity (Davis, Varol, Ferrara, Flammini, & Menczer, 2016), and focusing on high-volume flows of information associated with specific candidates, we can retrospectively look at how various kinds of accounts—including social bots (i.e., accounts with a high probability of being automated or otherwise inauthentic accounts)—engaged with messages posted by these candidates on Twitter. Second, our research is among the first (to our knowledge) comparative analyses of orchestrated activities around election campaigns in different countries at scale. Offering a comparative analysis of two recent national election campaigns with alleged orchestrated interference, we seek to compare the patterns and effects of orchestrated intervention in two countries whose election outcomes carry particular international geopolitical significance and gain a broader perspective on the potential capabilities and aims of orchestrated interference campaigns in online communication.

Below, we begin by analyzing the contexts of orchestrated interventions in election campaigns in the United States and Germany. We then present a literature review to unpack the three concepts central to our study of information flows: *scale*, *range,* and *speed*. *Scale* examines the number of times a given political candidate's tweet is retweeted; *range* (sometimes referred to as "depth") considers whether high-volume retweeting helps candidates increase the reach of their messages across new networks; and *speed* reflects the temporality of information flows (Yang & Counts, 2010). Questioning the assumptions made in the literature on computational propaganda, we present our two research questions:

RQ1: Scale—What were the temporal patterns of amplification of candidates' messages in the United States and the Federal Republic of Germany?

RQ2: Range—Did social bots have an effect on candidates' followership on Twitter in either national election?

We then turn to explaining our innovative approach to the detection of orchestrated activity using information flow signatures and present our findings.

Context: Social Bots on Twitter. Evidence from the United States and Germany

In the aftermath of the 2016 presidential election in the United States, observers began wondering whether other major democratic elections would face similar types of interference (e.g., Friedman, 2017). The federal elections in Germany, held in 2017, became a particular focus of attention, as commentators began arguing that Germany, under Angela Merkel, was the country best positioned to push back against interference; some even wondered if Merkel was "the leader of the free world," given her willingness to proactively push against electoral interference from Russia (Noack, 2016; Smale & Erlanger, 2018; Stelzenmüller, 2017). This led observers to anticipate that Germany might face some of the types of interference experienced by the United States (and to a lesser extent, the UK in the 2016 Brexit referendum [Howard & Kollanyi, 2016]).

Twitter played an important role in the German federal election of 2017: Despite higher levels of professional news consumption in Germany compared with the EU average, studies show this online media platform becoming a highly visible place for political discussions (Majó-Vázquez, Nurse, Simon, & Kleis Nielsen, 2017; Neudert, Kollanyi, & Howard, 2017). Though only 11% of the German public use Twitter, it has been widely adopted among the so-called legacy media organizations and public interest groups, and has been shown to drive political participation, especially among youth (Majó-Vázquez et al., 2017). In the United States, where 69% of the public believe Twitter to be an important platform for political activism (Anderson, Toor, Rainie, & Smith, 2018), 39% of Twitter users with public accounts have discussed national politics, although, similar to Germany, political discourse remains driven by a smaller share of accounts and constitutes a relatively small volume of overall conversation on Twitter (Pew Research Center, 2019).

The differences between the electoral systems in the United States and Germany, particularly the electoral college in the U.S. presidential election versus the personalized proportional representation system to the German Bundestag, might predicate different strategies for orchestrated interference. External actors

seeking to interfere with the U.S. presidential election may target conversations around the most prominent candidates. In Germany, external actors may target a wider range of political conversations in an attempt to influence the number and proportion of seats in the newly elected Bundestag, which will subsequently get to elect the chancellor. Despite these differences, just as with the U.S. presidential election, the German party leaders are the most prominent figures in the election campaigns who engage in personalized campaigning and receive outsized media attention, both online and offline (Enli & Skogerbø, 2013). Certain activist groups, including those supporting a national-conservative stance, became a subject of disproportionate user engagement on Twitter (in comparison with traditional matters of voter support), which was consistent with the information diffusion patterns in the United States from a year before (Neudert et al., 2017). This motivates our decision to focus on online political discussion involving selected candidates in both cases: the United States and Germany. In next section, we turn to examine the three dimensions of information diffusion: speed, scale, and range, and each of their roles in online political communication.

Background: Scale, Range, and Speed in Online Political Communication

Twitter is a strategic component of political campaigning, used by candidates to broadcast messages directly to the public (Bruns & Highfield, 2013). From an information theory standpoint, messages reach their recipients (in this case, the potential voters) through information cascades that propagate among user networks (Yang & Counts, 2010). Studies of virality have established that information flows can be suppressed, or promoted, by actors who occupy key positions in networks—these actors are known as network gatekeepers (Nahon & Hemsley, 2013). Each Twitter user acts as a gatekeeper when they decide to share a political message among their networks—facilitating downstream information flows that bridge multiple networks together (Nahon & Hemsley, 2013).

Information diffusion on online platforms can be characterized by three key dimensions: *speed*, which reflects temporality of information events; *scale*, which speaks to the visibility of a message on a platform through popularity metrics, such as "likes" and "retweets"; and *range*, which denotes the depth of diffusion once the message gets propagated through user networks, reaching new audiences in the process (Yang & Counts, 2010, pp. 356–357). From the structural viewpoint, two processes are at play with regard to range: *broadcast*, whereby a single influential node passes information onto a large audience who then adopt it, and *virality*, in which each node passes information to a few others who then diffuse it among their networks in a series of multilevel bursts (Goel, Anderson, Hofman, & Watts, 2016). When an information event is organic (i.e., not influenced by orchestrated activities), scale and range are correlated—when Twitter users retweet messages, they boost the visibility of the message and diffuse them through their networks, thus allowing the candidate's messages to reach a larger audience and gain new followers (Conway, Kenski, & Wang, 2013; Hemsley, 2016). When candidates gain new followers, their future messages can potentially reach even larger audiences, which adds both to the scale, as well as the range of information events (Hemsley, 2016). In this way, candidates can grow their audiences in waves that propagate their messages into new, distant networks. Speed is another parameter that characterizes the virality of the events, which might otherwise differ in scale and range. In this study, we consider speed in a combination with scale to detect orchestrated intervention in information events.

Though it is well established that bots (falsely) amplify political messages in an attempt to increase scale and create the "megaphone effect" (Woolley & Guilbeault, 2019, p. 193), we find this research to be scarce when looking specifically at the role of automated, semiautomated, and otherwise nonorganic accounts in the range of information propagation among new audiences. Studies in computational propaganda assume that by amplifying a candidate's message, social bots also "spread" it on the platform, assuming they have an audience to spread it to. In this study, we use empirical data to investigate this assumption. Since it is established that the number of retweets is related to how many new followers a user can get (Conway et al., 2013; Hemsley, 2016), and that more followers translates into a larger audience for later posts (Hemsley, 2016), we wonder whether social bots were as effective in amplifying political messages (i.e., increasing their scale) as they are at diffusing them (i.e., bringing new followers to candidates through their retweets). We consider these questions with regard to the speed of information events, which informs our data analysis models presented below.

Method

Information Flow Events

Detection of nonorganic information diffusion on Twitter is a potentially useful approach when looking for orchestrated political interventions. A growing body of literature is dedicated to identifying bots on Twitter: Dickerson, Kagan, and Subrahmanian (2014) were successful in identifying social bots using sentiment analysis on tweets drawn from the 2014 Indian election, and Wang (2010) used machine learning with spam bots to find that including detection features such as the numbers of duplicate tweets, URLs and replies or @mentions improved the models, as did the numbers of followers. Working to classify human, bot, and cyborg accounts on Twitter, Chu, Gianvecchio, Wang, and Jajodia (2012) found that account properties, tweeting behavior, and the content of the tweet could all be used to predict inauthentic behavior. Finally, moving more toward identifying groups of accounts working in concert, Chavoshi, Hamooni, and Mueen (2016) found that orchestration could be identified by looking at time intervals between posting of the same or similar material. Following this, we wonder whether there are other ways to detect the presence of social bots in political information events. This knowledge could be a useful way to identify attempts to amplify candidates' voices with the help of orchestrated intervention.

Data

In this work, we use data drawn from Twitter—specifically, the tweets posted by candidates during the 2016 U.S. election and the 2017 German election, as well as their retweets. As described below, we used Botometer (Varol, Ferrara, Davis, Menczer, & Flammini, 2017) in a somewhat unconventional manner—not only to generate "botness" scores for accounts that retweeted the candidates' tweets in our sample, indicating the probability of the accounts being inauthentic, but more importantly, to discover other properties of those accounts, such as instances when a user had deleted all of their tweets or used the protected tweets feature in the aftermath of the election. Choosing information events as our units of analysis, we are less interested in the structural and functional properties of individual accounts, and more interested in revealing patterns of their behavior in aggregate. As we explain in detail below, we generated some analytical plots to visualize the patterns of information diffusion in either election. We then use two

ordinary least squares (OLS) regression models for each election to compare the scale and range of diffusion of political messages.

Our data consist of Tweets collected during the 2016 U.S. presidential election and the 2017 German federal election. Tweets were collected using STACK (Hemsley, Ceskavich, & Tanupabrungsun, 2014), an open-source tool that collects data using Twitter's streaming API,[1] which supports the collection of tweets using a given set of search terms (the candidates' Twitter handles).[2] We used the "follow" parameter[3] for the API, which returns tweets created or retweeted by the user, as well as replies and retweets to/of any tweet created by the user. For the U.S. election, we collected candidates' tweets from August 1, 2016, through November 11, 2016. For the German election, we began collecting tweets from all candidates on August 12, 2017, and stopped on September 25, 2017 (the day after the election). In total, we collected slightly over 40,000,000 and 570,000 tweets for the U.S. and German elections, respectively.

With our interest in the role of social bots in information diffusion, we operationalize information flow cascades on Twitter as a retweet event (RTE). The concept of an RTE has been fruitfully used in other diffusion studies comparing aspects of information flows (Hemsley, 2016). An RTE is a group of tweets that includes a tweet and all retweets of that tweet. As an example, if the candidate Donald Trump posted a tweet that users retweeted 100 times, the corresponding RTE would constitute 101 tweets: Trump's initial tweet and all subsequent retweets of that tweet by the users. Grouping retweets together with their origin tweet allows us to compare the patterns of information flows, as well as to aggregate data about the individual retweets up to the RTE level (see Figure 1). Doing this with many RTEs allows us to compare information diffusion patterns. An RTE is an imperfect measure of an information flow since it does not include manual retweets, quotes, or user engagement metrics—the total number of people who may have read or otherwise interacted with the original tweet. However, an RTE is a useful measure of the number of times users attempted to increase the audience for a tweet by sharing it. In total, we have 26,396 RTEs for the U.S. election and 21,297 RTEs for the German election. In the U.S. sample, 196,375 accounts participated in the RTEs, and 15,352 accounts participated in the German RTEs. These numbers suggest that the proportion of individuals participating in Twitter discussions in Germany (compared with the general population of voters) was lower, although this audience demonstrated comparatively high engagement with the online content shared by the candidates.

[1] https://developer.twitter.com/en/docs/tweets/filter-realtime/overview

[2] For the German election, we also collected the political party accounts, but we did not analyze them for this work.

[3] https://developer.twitter.com/en/docs/tweets/filter-realtime/guides/basic-stream-parameters.html

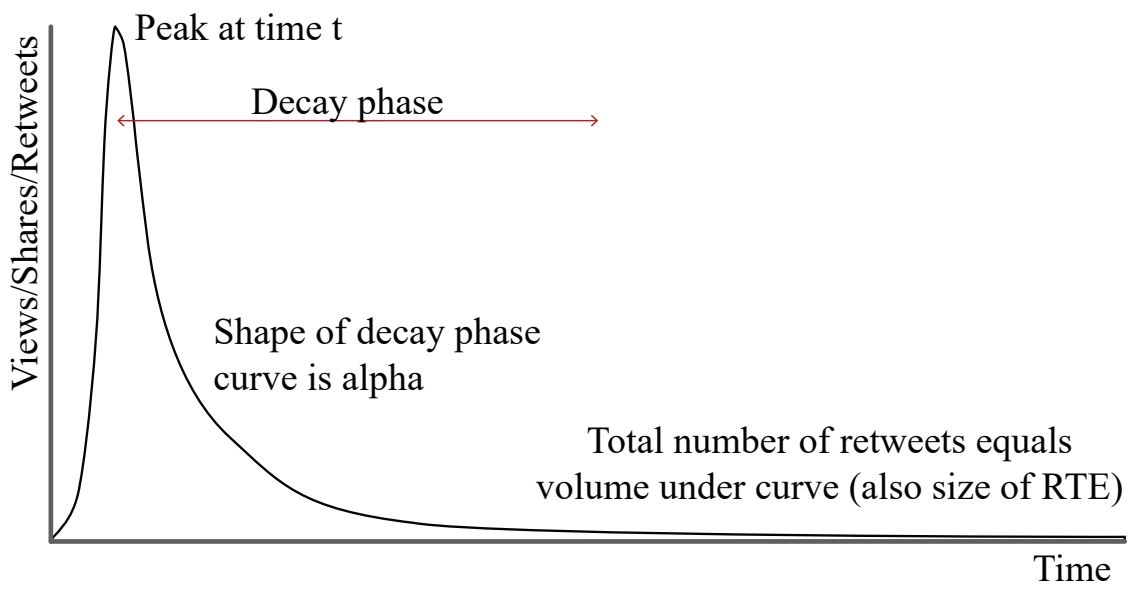

Figure 1. Information flow event signature model.

Given the differences in the scale of social media activity between the U.S. election and the German election, we used different sampling strategies for the two elections. We selected a sample size of 100 RTEs for both elections, which, according to a power analysis for regression models with 10 predictors (see below for model specifications), is more than sufficient for 95% confidence levels (Ott & Longnecker, 2010). As can be seen from Figure 2, for the U.S. election, the sizes of RTEs (i.e., the number of retweets in the RTE) are nearly power-law distributed, with the majority of RTEs being over size 100. Because of this distribution, we opted to use a stratified sample. We defined four categories based on the RTE's position in the distribution: low (up to 80th percentile), medium (80–90th percentile), medium/high (90–99th percentile), and high (99–100th, or the top 1 percentile). Then, we built a sample, selecting random RTEs matching the distribution, with 25 RTE coming from each of the U.S. presidential candidates representing the four main political parties: Donald Trump (Republican), Hillary Clinton (Democrat), Jill Stein (Green), and Gary Johnson (Libertarian).

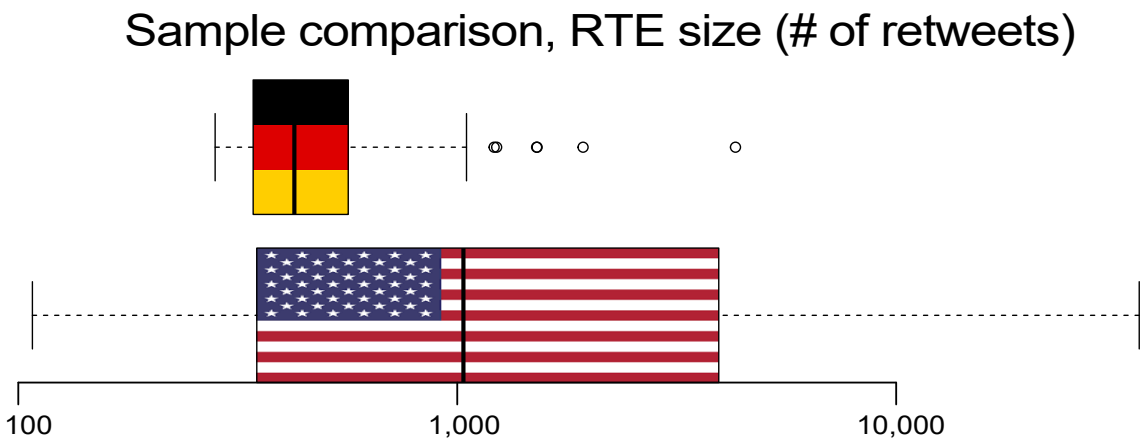

Figure 2. Distributions of sizes of retweet events (RTEs) for both samples.[4]

The scale of social media activity, as well as the distribution of sizes of RTEs, was much smaller for the German election, which meant that this stratified sampling strategy was not appropriate. Instead, our sample consists of the 100 largest RTEs associated with tweets from five leading politicians from five of the six parties that won seats in Germany's parliament (Angela Merkel did not have a Twitter account during the 2017 elections). These candidates were Martin Schulz (SPD), Alice Weidel (AfD), Christian Lindner (FDP), Sahra Wagenknecht (Left), and Cem Özdemir (Green). Dietmar Bartsch (Left) and Katrin Göring-Eckardt (Green), both of whom are prominent political figures in their respective political parties, did not receive much amplification on Twitter: None of their RTEs were in the 100 largest in the German data set, and thus were not included in our sample.

Botometer

As part of our effort to understand the role of various accounts in diffusion of political information, we used Botometer, a feature-based system that assigns scores to individual Twitter accounts indicating the probability that an account is automated (Varol et al., 2017). In the process of classification, Botometer extracts and considers more than a thousand account features across six different categories: user-based features, such as number of tweets and followers; behavior features, such as retweeting and being retweeted, mentioning and being mentioned; network features, such as the positionality of the user in retweet and mention networks, and measures of their popularity; content and language features; and finally, a range of sentiment features of the users' posts (Varol et al., 2017). Because of the set of markers used by this classification algorithm, we believe some of the accounts classified as "bots" might not necessarily be automated—human accounts known as "trolls" would also fall into this category: Their accounts would be recent, have suspicious features, and/or exhibit nonorganic behavior. Yet, given that our interest for this article lies in detecting any kind of orchestrated activity aimed at disrupting or altering the flows of political

[4] RTE size has been transformed with a log 10 to better show distributions. The *x*-axis has been adjusted appropriately.

information, and not in detecting the properties of individual accounts, we obtained Botometer scores for all users who participated in the RTEs in our data set, excluding users whose account was deleted, users who made their account private, and users who deleted all of their tweets. Importantly, Botometer is a time-sensitive tool—individual scores might change depending on most recent online behavior of the user. In the U.S. election, which happened first, we found that by the time we obtained Botometer scores (eight months following the election, July 16, 2017), many of the accounts that were active around the election had been deleted or suspended (22,281; 11.3%), deleted all of their tweets (697; 0.4%), or switched to a protected status (18,393; 9.3%). For the remaining users, "botness" scores appeared to be normally distributed, with a slight skew to the left (see Figure 3).[5]

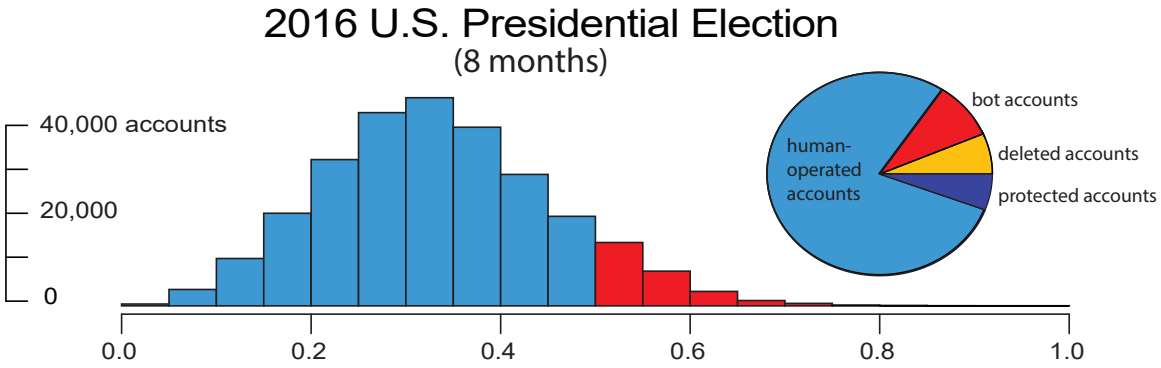

Figure 3. Distribution of accounts identified by Botometer as "bots" (red), protected accounts (blue), and deleted accounts/tweets (yellow), in the U.S. sample.[6]

Traditionally, accounts with a Botometer score of 50% and higher raise suspicion of being automated (Woolley & Guilbeault, 2019). What makes our analysis different is that we also include accounts that switched to private status, were deleted or suspended, or deleted all of their tweets because of their large volume in our sample. As Figure 3 shows, the distribution of accounts broadly categorized as "suspicious" among the U.S. users in our sample is consistent with prior findings of about a quarter of all political discourse being run by inauthentic and/or nonhuman agents (Bessi & Ferrara, 2016). Importantly, researchers traditionally avoided the two other categories of suspicious users—those who deleted their accounts immediately following the election, and those who switched their accounts to a "protected" status following the election. Prior research on the U.S. presidential election (Boichak et al., 2018), however,

[5] One potential limitation of this approach is the possible changes to the proprietary bot detection algorithm by Botometer over time, which might have influenced the reliability of results across the two samples. However, as the two election campaigns happened at different points in time, it was not possible to obtain perfectly reproducible results for each sample. To mitigate this limitation and ensure our results had an accurate basis of comparison across the two samples, we reran the Botometer scores for each of the accounts in the German sample (see Figure 4). Though the bot score is only one of our measures of interest, the other measures, including the deleted and protected account status, have been consistent over time and are therefore comparable.

[6] These data were collected eight months following the U.S. election.

suggests that these categories of accounts might have distinct behavior patterns and serve a range of specific functions. Given that the deleted and protected accounts jointly comprise more than 20% of our U.S. sample, we made a decision to include them in our analysis as a separate category. As we explain below, this decision informed our strategy of analyzing the accounts from the German sample.

For the German sample, we initially obtained Botometer scores in near real time (with a small lag based on rate limits for the Botometer API). The distribution of scores is represented at the top in Figure 4.

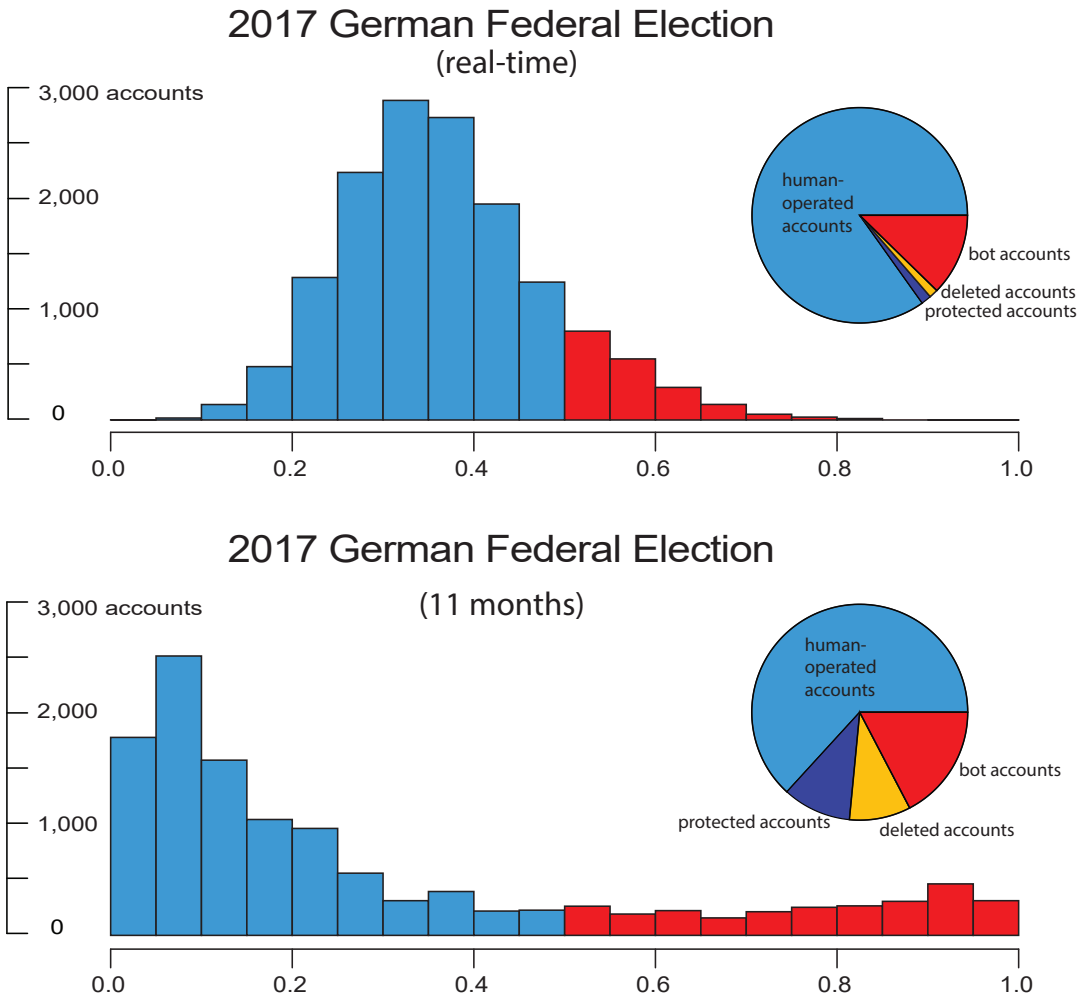

Figure 4. Distribution of accounts identified by Botometer as "bots" (red), protected accounts (blue), and deleted accounts/tweets (yellow), in the German sample.[7]

[7] Top: Data collected during election campaigning (date). Bottom: Eleven months following the election.

Of the 15,352 accounts from our data set, at the time of data collection, 13,026 were initially categorized by Botometer as human (85%), 1,889 as bots (12%), 205 had deleted their accounts (1.3%), and 232 accounts had a protected status (1.5%). Yet, when we reran the same sample through Botometer 11 months following the election (see Figure 4, bottom), the results were drastically different: Only 9,704 accounts were categorized as human (63%), whereas 2,662 accounts were identified as bots (17%). The number of protected accounts, however, saw a six-fold increase to 1,575 (10.2%), and 1,411 accounts have been deleted (9.2%), which is seven times as many as before. One plausible explanation for this sharp increase in bot score: Given that Twitter drove engagement around election among youth (Majó-Vázquez et al., 2017), it might be the case that many users were less active on Twitter when the election was over, which might have falsely increased the probability of these accounts being identified as a bot. As explained above, we used the second set of scores for analysis to ensure an accurate basis for comparison between the U.S. and German data.

Scale and Range

To answer RQ1 and discover what types of Twitter accounts drove the scale of candidates' RTEs, we ran an OLS regression (1) with RTE size (number of retweets) as the dependent variable (see Table 1). We included the number of candidate's followers at the start of an RTE as our control variable because this relationship is well established (Suh, Hong, Pirolli, & Chi, 2010). We also controlled for duplicate users, as we noticed that some accounts participated in an RTE more than once. Because of the nature of the dependent variable (RTE sizes follow a power-law distribution, which was used to stratify our sample), we decided to subject it to a transformation using a natural logarithm, turning Model 1 into a log-linear model that can be represented as follows:

$$ln(RTE\ size_i) = \alpha + \beta_1 mean\ bot\ score_i + \beta_2 protected\ account\ ratio_i + \beta_3 deletedaccount\ ratio_i + \beta_4 duplicate\ users_i + \beta_5 candidate's\ followers_i + \varepsilon_i\ , \quad (1)$$

in which mean bot score is the average bot score of all accounts that participated in an RTE, protected/deleted account ratios stand for the proportions of the respective accounts in an RTE, and the number of candidate's followers is a control variable. We checked for multicollinearity by calculating a variance inflation factor for each variable, all of which were below the recommended threshold of four. We also made a series of postestimation descriptive plots and tests to check the OLS regression assumptions—linearity, normality, exogeneity, as well as residual (error) distribution. We removed one outlier observation from the U.S. data set and two from the German data set, as they showed a large leverage effect. We ended up analyzing 99 RTEs in the U.S. data, and 98 RTEs in the German data. To ensure comparability of effect sizes, we include standardized coefficients (beta).

Table 1. Scale: Patterns of Information Amplification in the United States and Germany (Model 1).

DV: ln(RTE size)	United States			Germany		
IVs:	Coef.	*SE*	Coef. (beta)	Coef.	*SE*	Coef. (beta)
Intercept	9.97***	0.80	–	5.71***	0.29	–
Mean bot score	−11.45***	1.72	−0.26***	1.27	1.03	0.10
Protected account ratio	12.33**	4.61	0.11**	−0.47	0.52	−0.10
Deleted account ratio	−8.37*	3.99	−0.10*	−2.58***	0.69	−0.33***
Number of duplicate users	0.01***	0.002	0.38***	0.003***	0.000	0.98***
Number of cand. followers	0.000***	0.000	0.57***	0.000*	0.000	0.25*
Multiple *R*-squared		0.89			0.58	
Adjusted *R*-squared		0.89			0.56	
F statistic		152.4 on 5 and 93 *df*			25.72 on 5 and 92 *df*	

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

The standardized coefficients show similarities, as well as differences, in the scale of amplification of candidates' messages in the United States and Germany. Expectantly, the number of candidate followers was a significant predictor of RTE size across both elections. We see that bots' accounts, identified by Botometer, were significant predictors of RTE size in the U.S. election, which was not the case in Germany. Accounts with a protected status amplified the scale of RTEs in the U.S. election, yet did not seem to have a similar effect in the German one. Accounts that were subsequently deleted had a significant negative impact on the scale of RTE in both elections. Interestingly, the number of duplicate users was a significant predictor of RTE scale in both elections, which suggests that some users may have retweeted the same message more than once—a behavior pattern found in other contexts, such as marketing (Ghosh, Surachawala, & Lerman, 2011).

In Figure 5, we investigate duplicate retweeting of the same message as a basic amplification strategy. Note that each point is a RTE in our set, with the *y*-axis indicating how many of the retweets in that event were done by users who retweeted that event more than once. The points are sorted based on the RTE size, with the largest on the left. We see that every tweet from Donald Trump in our U.S. sample has been retweeted multiple times by some accounts. At least five of Hillary Clinton's tweets from our U.S. sample involved large amounts of duplicate retweeters, followed by a few where the numbers of duplicate retweeters were fewer than 25. Since the plot is sorted by RTE size, we can see a positive relationship between the number of duplicate tweets and the number of times a tweet was retweeted. Though this behavior is not surprising per se, it does appear to be an amplification strategy by those following the candidates.

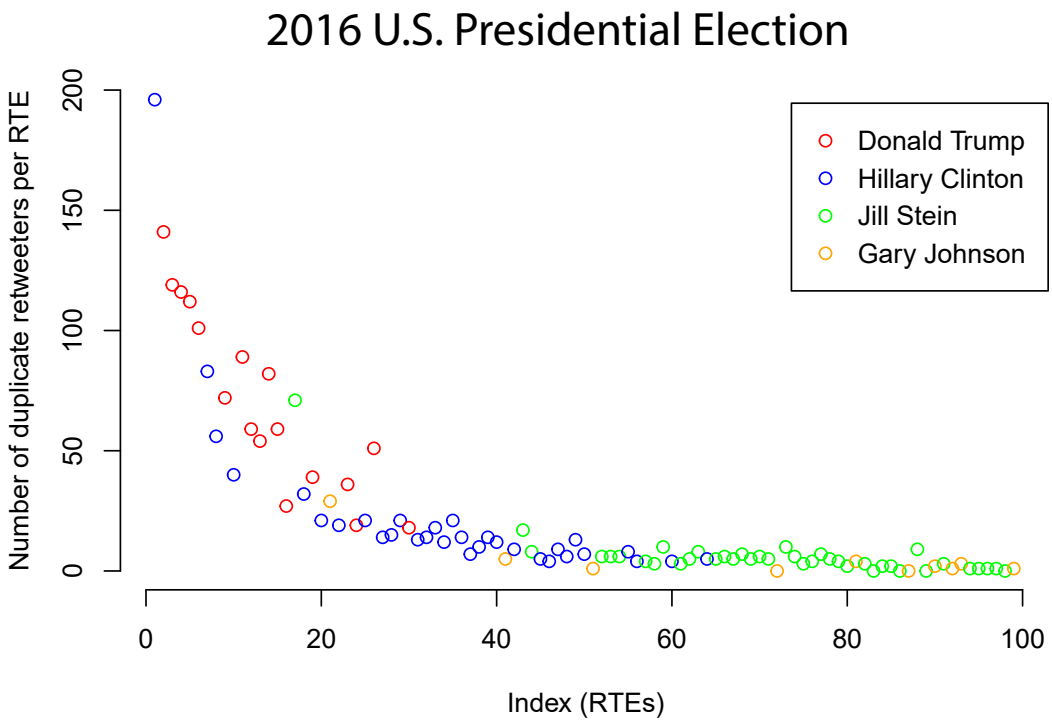

Figure 5. Amplification patterns in the U.S. election—duplicate retweeters.

In the German sample, we find that this pattern persists in the case of two candidates: Alice Weidel and Sahra Wagenknecht, whose every tweet was amplified by duplicate retweeters (see Figure 6).

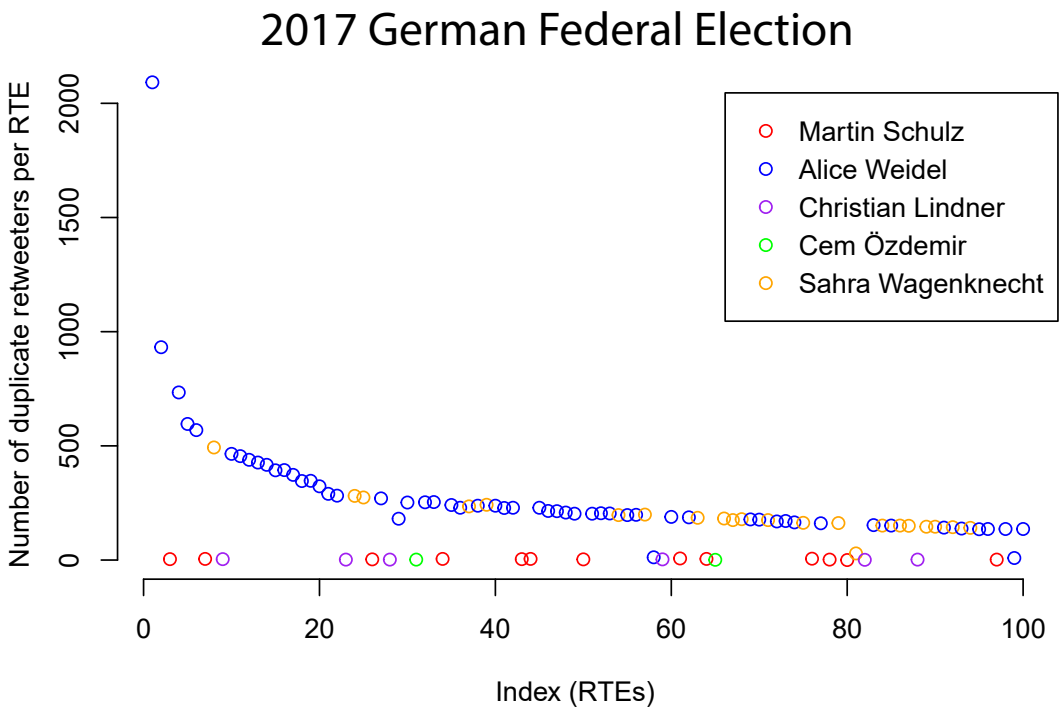

Figure 6. Amplification patterns in the German election—duplicate retweeters.

To answer RQ2, we ran a second OLS regression (2) with candidates' followership as the dependent variable (2). Candidate's followership was part of metadata returned by Twitter for each retweet of a candidate's message—therefore, we could observe its change over the course of an RTE. We obtain change in followership over the course of each RTE by subtracting the number of candidate's followers at the end of the RTE from the number of followers at the beginning. We include the mean bot score of an RTE, the ratios of protected and deleted accounts, and the number of candidate's followers at the beginning of each RTE as our independent variables. Model 2 could be represented as follows:

$$\Delta followers_i = \alpha + \beta_1 mean\ bot\ score_i + \beta_2 protected\ account\ ratio_i + \beta_3 deleted\ account\ ratio_i + \beta_4 candidate's\ followers_i + \varepsilon_i. \quad (2)$$

Similar to Model 1 (see Table 1), we checked the variance inflation factor and made descriptive plots to verify regression assumptions. We found all assumptions to hold. We also dropped the same outliers (one in the U.S. data set and two in the German data set) for consistency.

Table 2. Patterns of Information Diffusion, United States and Germany (Model 2).

DV: Change in followers	United States			Germany		
IVs:	Coef.	*SE*	Coef. (beta)	Coef.	*SE*	Coef. (beta)
Intercept	−4498432.60	2885593.43	–	6,625.53	7,694.77	–
mean bot score	13935610.46*	6173223.93	0.14*	22,423.62	27,004.43	0.07
Protected account ratio	44911974.76**	16379096.14	0.18**	−29,158.53*	13.390.39	−0.23*
Deleted account ratio	−22120797.86	14453224.91	−0.12	−25.069,06	17,419.26	−0.13
Candidate followers	0.36***	0.05	0.60***	0.03***	0.006	0.51***
Multiple *R*-squared	0.70			0.54		
Adjusted *R*-squared	0.69			0.52		
F statistic	54.79 on 4 and 94 *df*			26.89 on 4 and 93 *df*		

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

We see that protected accounts had a different relationship with the change in followers in the two elections. Though they were strongly associated with gaining new followers in the U.S. election, they had an inverse relationship in the German election. In both election campaigns, the number of the candidate's followers was a significant predictor of gaining new followers, which speaks to the range (depth) of information events. In the U.S. election, mean bot score was also positively associated with gaining new followers.

Next, we compare the rate with which the accounts tweeted over the course of the election campaign, plotted by the numbers of their followers, to compare behavior patterns among the two samples (see Figures 7 and 8). We see that, compared with the United States, all types of suspicious accounts in the German election tweeted less and had smaller networks of followers. Accounts assumed to be social bots in the U.S. sample, on the other hand, had comparatively higher numbers of followers than the alleged authentic user accounts, which might suggest they could be part of botnets—networks of inauthentic accounts connected to each other (Abokhodair et al., 2015; Woolley & Guilbeault, 2019).

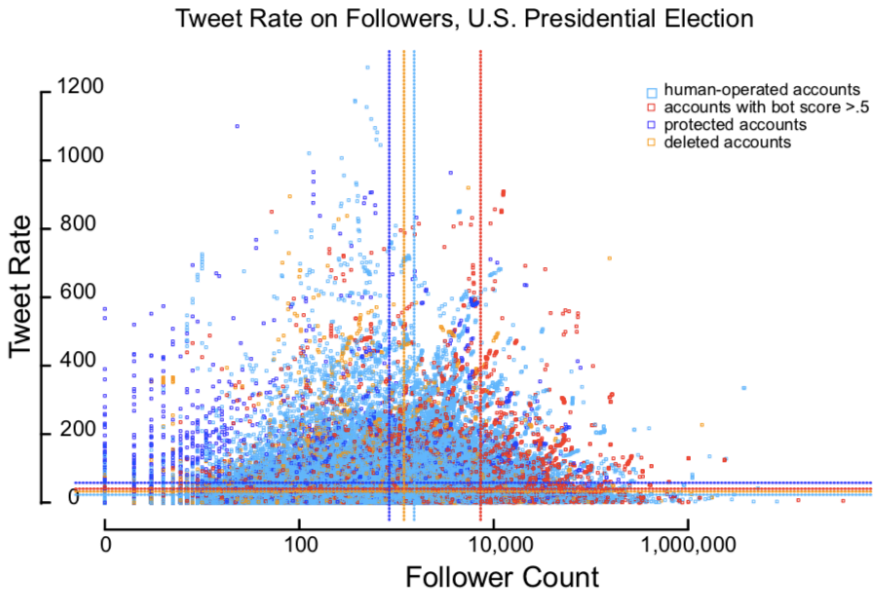

Figure 7. The distribution of tweet rate by follower count in the U.S. election.[8]

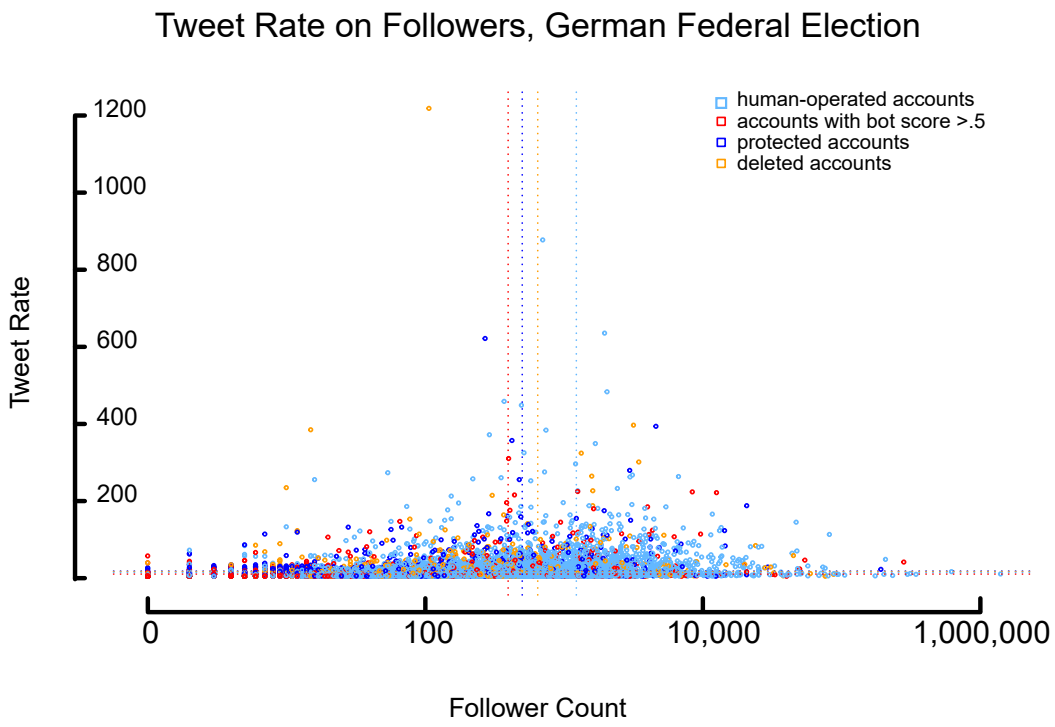

Figure 8. The distribution of tweet rate by follower count in the German election.8

Findings

As the above analysis shows, a few distinct information diffusion patterns persist across election campaigns and point to some important regularities in orchestrated political intervention.

Scale

Our first OLS regression model (see Table 1) illuminates a simple, heavy-handed amplification technique used across both elections—duplicate retweeting. Duplicate retweeting happens when a Twitter account participates in an RTE more than once, essentially retweeting the same candidate's message multiple times, which leads to the increase of scale of the respective RTE. We see that, in the U.S. data set, Donald Trump's tweets heavily benefited from duplicate retweeting—in some RTEs, up to 150 accounts have retweeted the same message multiple times. Curiously, this is the case with all of Donald Trump's RTEs in our sample, which suggests there is an amplification pattern (see Figure 5). One of Hillary Clinton's messages in our sample was heavily amplified using the same method, and a couple of other RTEs looked

[8] Dotted lines indicate mean values for each type of account. Note that plot *x*-axis is a log-base 10.

like they have been assisted using this method. As we see from Figure 7, many accounts with a protected status exhibited abnormally high tweet rates—the status makes it difficult to see which candidate's messages they were retweeting. Yet it is possible that they were retweeting the content more than once to achieve the "megaphone effect" for candidate's messages (Woolley & Guilbeault, 2019, p. 193). In support of this finding, we see that protected accounts were a significant predictor of the scale of RTEs in the U.S. election (see Table 1).

In the German election, we observe a similar pattern of duplicate retweeting to amplify candidates' messages. Although the influence of protected accounts on RTE size is insignificant in Model 1 (RTEs from the German sample were generally much smaller), it is clear that duplicate retweeters were a powerful driver of RTE scale. In the previous version of Model 1, before we made a decision to control for duplicate accounts, protected accounts were a significant predictor of RTE size (we will take a closer look at those accounts and their behavior below). As seen from Figure 6, all of Alice Weidel's and Sahra Wagenknecht's tweets from our sample were amplified by duplicate tweeters. In fact, this false amplification was the reason why Alice Weidel's RTEs were much larger than those of the other candidates. Unlike in the U.S. election, in Figure 7 we see that both the protected accounts and the accounts identified as "bots" by Botometer from the German sample have very low followership, but comparatively high tweet rates. Once again, we might speculate some of these accounts might have been involved in duplicate retweeting. These findings suggest that protected accounts were part of a larger orchestrated endeavor, such as botnets that have been previously mentioned in this article. Without speculating whether these accounts were being controlled by a human or an algorithm, we conclude that accounts with a protected status were important players in the amplification game in both elections.

We also find, in both elections, that the deleted accounts had a significant inverse relationship with the scale of RTEs; though we do not know whether these accounts have been deleted or suspended by Twitter as possible bot accounts, we can speculate that they might have been used to amplify tweets that have received insufficient attention from the candidates' human followers. In Figures 7 and 8, we see that the deleted accounts exhibited no particular pattern of activity and did not have consistent numbers of followers in both elections. Finally, consistent with previous findings (Conway et al., 2013; Hemsley, 2016), the number of followers was found to be a significant predictor of RTE scale.

Range

In our second OLS regression model (see Table 2), we see that protected accounts were strongly associated with gaining new followers only in the U.S. election. Mean bot score was also positively associated with gaining new followers, as was the size of the candidate's audience at the time of the tweet (Conway et al., 2013; Hemsley, 2016). Figure 7 helps us interpret those findings: We see that in the U.S. sample, "bot" accounts and some protected accounts had very high followership, which might indicate one of two scenarios: Either the bots were following each other, in which case they would make a botnet (Abokhodair et al., 2015; Woolley & Guilbeault, 2019), or, more likely, human accounts were following accounts thought to be social bots, and decided to follow the candidate after seeing their tweet (see, e.g., Timberg & Harris, 2018). In light of an alleged impact of an orchestrated intervention by the Russian Internet Research Agency (Timberg & Harris, 2018), positing that human users would follow the accounts that impersonated U.S.

citizens, is not a far-fetched assumption. These findings suggest that, in the U.S. election, accounts broadly categorized as social bots were not only amplifying, but also diffusing candidates' messages, helping them reach new audiences.

We do not find this to be true in the German election: Though candidate's followers were a significant predictor of engaging new followers, we find that the influence of "bot" accounts on the range of diffusion is insignificant, as protected accounts have a significant inverse relationship with the number of followers. Figure 8 helps explain this finding: On the plot, we see that some of the accounts categorized as "bots," as well as some protected accounts, had zero followers. This highlights an important difference between the U.S. and German election: Though we assume that social bots might have had human followers in the United States, in Germany, they on average did not, and there is very little evidence of botnets (i.e., bots were not following each other). Alternatively, this might have been an audience effect, as Twitter was less prominent in the German election compared with the United States. The influence of deleted or suspended accounts was not a statistically significant predictor of information diffusion across both samples.

A comparative analysis of RTE signatures also yields interesting findings about differences among candidates. In the U.S. sample, Donald Trump (Republican) appeared to have benefited most from amplification and was also disproportionately likely to gain new followers in the course of the election campaign. For instance, one retweet event (see Figure 9) had a size of 28,010 and a mean bot score of 0.41. Among the accounts that retweeted this message, 5,702 were created in a three-day window: between December 31, 2014, and January 2, 2015. Moreover, the RTE had 4,543 protected accounts (16%), 1,759 of which were created on January 19, 2015. Finally, the RTE had 1,223 deleted accounts (4.3%). In other words, this, and many other of Donald Trump's RTEs, appear to be strongly driven by suspicious inauthentic accounts. Hillary Clinton (Democrat) and Jill Stein (Green) demonstrated medium levels of orchestrated, nonorganic engagement—a few of their tweets from our data set appear to have been falsely amplified. Gary Johnson (Libertarian) had rather low levels of both organic and orchestrated engagement with his messages.

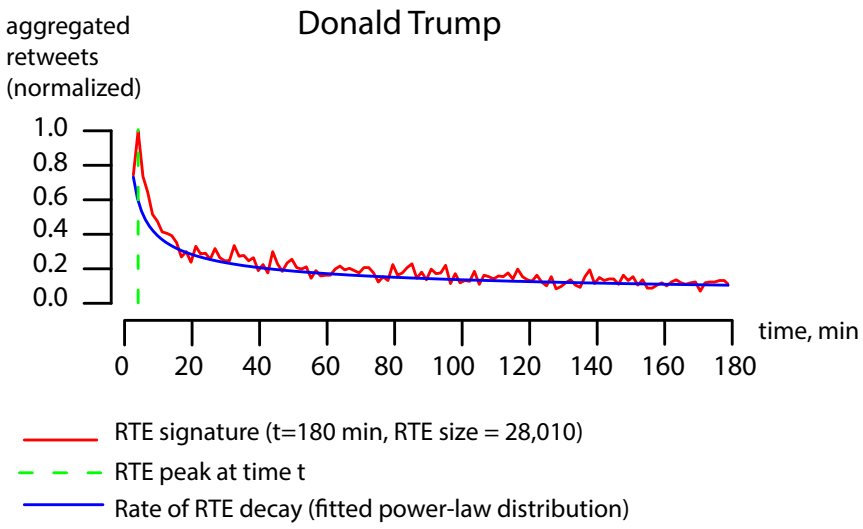

Figure 9. One of Donald Trump's tweets that "appears" organic, but is in fact largely driven by suspicious inauthentic accounts.

When examining RTE signatures, more gradual declines in retweets after the peak have been shown to mean that the information flows from user to user to more distant parts of the network, whereas sharper declines indicates that the information did not flow past one's own followers, speaking to its "broadcast" structure and limited range (Hemsley, 2016). Using this knowledge, we can gauge the *range* of RTEs (i.e., the degree to which an RTE likely reached new audiences). In the German federal election, Martin Schulz (SPD) appeared to have had more RTEs that reached far-off audiences, which gained him comparatively high numbers of followers throughout the campaign. Dr. Alice Weidel (AfD) had been the most amplified candidate in our data set—one of her viral tweets (see Figure 10) received 1,212 retweets from only 266 unique users; that RTE had a mean bot score of 0.44 and 70.42% of the accounts were protected, some of which have Russian names or handles. Most of the accounts that retweeted this message had been created between August and September of 2017, and 11.93% have since been deleted. Christian Lindner (FDP) had the highest diffusion of tweets to distant parts of the network, which seemed mostly organic (low numbers of suspicious accounts), and helped the candidate gain followers in the course of the campaign. Cem Özdemir's (Green) RTEs exhibited low levels of organic engagement, and Sahra Wagenknecht's (Left) RTE signatures showed medium levels of predominantly organic engagement, with occasional orchestrated amplification.

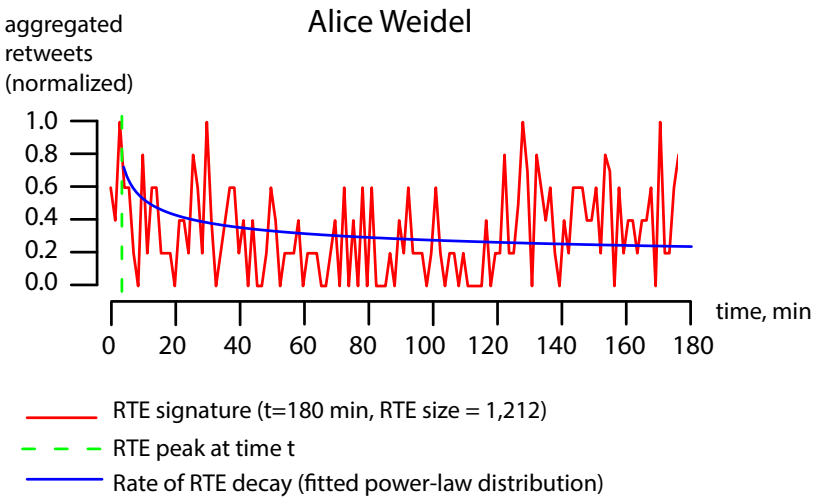

Figure 10. One of Alice Weidel's tweets exhibits a clear pattern of orchestrated amplification.

Importantly, the RTE signatures shown above are somewhat extreme examples and serve to illustrate our method, rather than to make definitive conclusions about the use of social bots by (or on behalf of) any of these candidates. Though we may not know the motive behind the deployment of those accounts (whether to boost the candidate's visibility, or to discredit the candidate), our method nonetheless helped detect some instances of nonorganic interference in political campaigns.

Discussion and Conclusion

Our study provides insight on the range of effects of orchestrated activity in different political contexts—directly on the scale and range of information diffusion online, and indirectly on public opinion and electoral outcomes. Distinguishing and critically interrogating the patterns of behavior that included inauthentic Twitter accounts, otherwise known as "social bots," is crucial for understanding the scale and range of orchestrated campaigns of political manipulation across political contexts. Most importantly, finding similar patterns across the U.S. and the German elections helps ground claims about suspicious activity aimed at interfering in elections. Though we have little evidence to conclusively link this activity to external actors, finding similar patterns across election campaigns is a necessary step to detecting orchestrated political intervention.

Our main contribution to the literature on computational propaganda lies in distinguishing organic and orchestrated information diffusion patterns—through the concepts of scale, range, and speed. The findings bring us to the following conclusion: Although social bots do increase the scale of RTEs, they rarely diffuse information to new audiences. We find this pattern to be a basic feature of orchestrated campaigns—whereas in organic information flows scale and range go hand-in-hand, social bots only mimic social spread of information, but, in most cases, are not capable of driving user engagement around the content. This disconnect between scale and range in information events is a crucial indicator of

orchestrated (by algorithmic or otherwise inauthentic agents) activity. Yet, as we see in the U.S. election, social bots can successfully impersonate human users and become influential actors in reaching new audiences. These findings are consistent with the nature of virality, which has the power to reproduce, but also transform, social norms and institutions (Goel et al., 2016; Nahon & Hemsley, 2013)—including those around election campaigns. It appears that, from an information diffusion standpoint, some bots have been successful at transforming the flows of political communication leading up to elections, which is consistent with prior research in this context (Woolley & Guilbeault, 2019).

We also hope that our novel analytic approach—using RTEs as a unit of analysis in detecting orchestrated interventions—will contribute to the body of literature on detecting computational propaganda and mitigating its effects. As we demonstrate above, visualizing RTE signatures may not always help detect nonorganic, orchestrated intervention, yet using RTEs as a unit of analysis appears beneficial to detect such intervention by using computational tools (such as looking at the proportion of deleted accounts, or generating a timeline of account creation). Furthermore, we find that it is very difficult to detect orchestrated intervention in real time. Usually, such intervention becomes visible in hindsight, months after the election is over and the actors are trying to eliminate traces of their orchestrated activity. In the case of both election campaigns, we see disproportionate numbers of accounts that participated in the political discussion on Twitter resort to two kinds of activities: (1) hide behind a protected status, and (2) delete either the account or the tweets in the account. Users may toggle private status on and off, delete their accounts, or even delete all of their tweets for reasons that have nothing to do with attempts to hide orchestrated intervention or other forms of information manipulation; however, we believe that the instances of this behavior were too common to be idiosyncratic. For these reasons, we conclude that by shifting the focus from individuals to information events, our method, RTE signatures, has demonstrated potential to detect orchestrated intervention in online political campaigns. As we see from this study, social bots and their instigators may use many functional forms and employ diverse strategies to prevent detection.

Finally, our comparative study provides useful factual data on orchestrated campaigns, shedding light on the information flow processes that contributed to a disproportionate prominence of certain candidates in two national elections: the U.S. and the Federal Republic of Germany. Going forward, this will help researchers and policy makers better understand the political consequences of orchestrated intervention on social media and its role in distorting democratic representation.

References

Abokhodair, N., Yoo, D., & McDonald, D. W. (2015). Dissecting a social botnet: Growth, content and influence in Twitter. *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing—CSCW '15* (pp. 839–851). doi:10.1145/2675133.2675208

Anderson, B., DiResta, R., Little, J., Maiberg, E., Morgan, J., Pasternack, A., & Nimmo, B. (2017, November 2). *The bots that are changing politics*. Retrieved from https://www.vice.com/en/article/mb37k4/twitter-facebook-google-bots-misinformation-changing-politics

Anderson, M., Toor, S., Rainie, L., & Smith, A. (2018). *Activism in the social media age*. Washington, DC: Pew Research Center. Retrieved from http://www.pewinternet.org/2018/07/11/activism-in-the-social-media-age/

Bessi, A., & Ferrara, E. (2016). Social bots distort the 2016 U.S. presidential election online discussion. *First Monday*, *21*(11). doi:10.5210/fm.v21i11.7090

Boichak, O., Jackson, S., Hemsley, J., & Tanupabrungsun, S. (2018). Automated diffusion? Bots and their influence during the 2016 U.S. presidential election. In G. Chowdhury, J. McLeod, V. Gillet, & P. Willett (Eds.), *Transforming digital worlds* (pp. 17–26). Cham, Switzerland: Springer. doi:10.1007/978-3-319-78105-1_3

Bruns, A., & Highfield, T. (2013). Political networks on Twitter. *Information, Communication & Society*, *16*(5), 667–691. doi:10.1080/1369118X.2013.782328

Calabresi, M. (2017, May 18). Inside Russia's social media war on America. *TIME*. Retrieved from https://time.com/4783932/inside-russia-social-media-war-america/

Chavoshi, N., Hamooni, H., & Mueen, A. (2016). Identifying correlated bots in Twitter. In E. Spiro & Y. Y. Ahn (Eds.), *Social informatics* (pp. 14–21). Cham, Switzerland: Springer. doi:10.1007/978-3-319-47874-6_2

Chu, Z., Gianvecchio, S., Wang, H., & Jajodia, S. (2012). Detecting automation of Twitter accounts: Are you a human, bot, or cyborg? *IEEE Transactions on Dependable and Secure Computing*, *9*(6), 811–824. doi:10.1109/TDSC.2012.75

Conway, B. A., Kenski, K., & Wang, D. (2013). Twitter use by presidential primary candidates during the 2012 campaign. *American Behavioral Scientist*, *57*(11), 1596–1610. doi:10.1177/0002764213489014

Davis, C. A., Varol, O., Ferrara, E., Flammini, A., & Menczer, F. (2016). BotOrNot: A system to evaluate social bots. *Proceedings of the 25th International Conference Companion on World Wide Web—WWW '16 Companion* (pp. 273–274). doi:10.1145/2872518.2889302

Dickerson, J. P., Kagan, V., & Subrahmanian, V. S. (2014). Using sentiment to detect bots on Twitter: Are humans more opinionated than bots? *2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2014)* (pp. 620–627). doi:10.1109/ASONAM.2014.6921650

Enli, G. S., & Skogerbø, E. (2013). Personalized campaigns in party-centred politics. *Information, Communication & Society*, *16*(5), 757–774. doi:10.1080/1369118X.2013.782330

Ferrara, E. (2017). Disinformation and social bot operations in the run up to the 2017 French presidential election. *First Monday*, *22*(8). doi:10.5210/fm.v22i8.8005

Forelle, M. C., Howard, P. N., Monroy-Hernández, A., & Savage, S. (2015). *Political bots and the manipulation of public opinion in Venezuela.* Retrieved from https://arxiv.org/abs/1507.07109

Friedman, U. (2017, April 26). Russia's interference in the U.S. election was just the beginning. *The Atlantic.* Retrieved from https://www.theatlantic.com/international/archive/2017/04/russia-election-europe-us/524208/

Ghosh, R., Surachawala, T., & Lerman, K. (2011). *Entropy-based classification of "retweeting" activity on Twitter.* Retrieved from https://arxiv.org/abs/1106.0346

Giglietto, F., Iannelli, L., Rossi, L., & Valeriani, A. (2016). *Fakes, news and the election: A new taxonomy for the study of misleading information within the hybrid media system* (SSRN Scholarly Paper No. ID 2878774). Rochester, NY: Social Science Research Network. Retrieved from https://papers.ssrn.com/abstract=2878774

Goel, S., Anderson, A., Hofman, J., & Watts, D. J. (2016). The structural virality of online diffusion. *Management Science*, *62*(1), 180–196. doi:10.1287/mnsc.2015.2158

Gorwa, R., & Guilbeault, D. (2018). Unpacking the social media bot: A typology to guide research and policy. *Policy & Internet*, *12*(2), 225–248. doi:10.1002/poi3.184

Greenberg, A. (2017, May 9). NSA director confirms that Russia really did hack the French election. *Wired*. Retrieved from https://www.wired.com/2017/05/nsa-director-confirms-russia-hacked-french-election-infrastructure/

Hemsley, J. (2016). Studying the viral growth of a connective action network using information event signatures. *First Monday*, *21*(8). doi:10.5210/fm.v21i8.6650

Hemsley, J., Ceskavich, B., & Tanupabrungsun, S. (2014). STACK (Version 1.0) [Computer software]. Retrieved from https://github.com/bitslabsyr/stack

Howard, P. N., & Kollanyi, B. (2016). *Bots, #StrongerIn, and #Brexit: Computational propaganda during the UK–EU referendum*. Retrieved from http://arxiv.org/abs/1606.06356

Kollanyi, B., Howard, P. N., & Woolley, S. C. (2016). *Bots and automation over Twitter during the First U.S. presidential debate* (Data Memo No. 2016.3). Oxford, UK: Oxford Internet Institute. Retrieved from https://comprop.oii.ox.ac.uk/research/posts/bots-and-automation-over-twitter-during-the-third-u-s-presidential-debate/

Majó-Vázquez, S., Nurse, J. R. C., Simon, F. M., & Kleis Nielsen, R. (2017). *Digital-born and legacy news media on Twitter during the German federal election*. Retrieved from http://reutersinstitute.politics.ox.ac.uk/sites/default/files/2017-10/20171030_RISJ_German_Factsheet_.pdf

Nahon, K., & Hemsley, J. (2013). *Going viral*. Cambridge, UK: Polity.

Neudert, L. M., Kollanyi, B., & Howard, P. N. (2017). *Junk news and bots during the German parliamentary election: What are German voters sharing over Twitter?* (Data Memo No. 2017.7). Oxford, UK: Oxford Internet Institute. Retrieved from http://comprop.oii.ox.ac.uk/wp-content/uploads/sites/89/2017/09/ComProp_GermanElections_Sep2017v5.pdf

Noack, R. (2016, November 21). How Angela Merkel, a conservative, became the "leader of the free world." *The Washington Post*. Retrieved from https://www.washingtonpost.com/news/worldviews/wp/2016/11/21/how-angela-merkel-a-conservative-became-the-leader-of-the-liberal-free-world/

Ott, L., & Longnecker, M. (2010). *An introduction to statistical methods and data analysis* (6th ed.). Boston, MA: Brooks/Cole Cengage Learning.

Pazzanese, C. (2017, May 3). Former officials see Russian effort to disrupt election. *The Harvard Gazette.* Retrieved from http://news.harvard.edu/gazette/story/2017/05/former-officials-see-russian-effort-to-disrupt-election/

Pew Research Center. (2019, October 23). *All the presidents' numbers.* Retrieved from https://www.people-press.org/2004/01/18/all-the-presidents-numbers/

Salge, C., & Karahanna, E. (2016). Protesting corruption on Twitter: Is it a bot or is it a person? *Academy of Management Discoveries*, *4*(1), 32–49. doi:10.5465/amd.2015.0121

Smale, A., & Erlanger, S. (2018, January 20). As Obama exits world stage, Angela Merkel may be the liberal West's last defender. *The New York Times*. Retrieved from https://www.nytimes.com/2016/11/13/world/europe/germany-merkel-trump-election.html

Stelzenmüller, C. (2017, *June 28). The impact of Russian interference on Germany's 2017 elections.* Retrieved from https://www.brookings.edu/testimonies/the-impact-of-russian-interference-on-germanys-2017-elections/

Suh, B., Hong, L., Pirolli, P., & Chi, E. H. (2010). Want to be retweeted? Large scale analytics on factors impacting retweet in Twitter network. *2010 IEEE Second International Conference on Social Computing* (pp. 177–184). doi:10.1109/SocialCom.2010.33

Timberg, C., & Harris, S. (2018, July 20). Russian operatives blasted 18,000 tweets ahead of a huge news day during the 2016 presidential campaign. Did they know what was coming? *The Washington Post*. Retrieved from https://www.washingtonpost.com/technology/2018/07/20/russian-operatives-blasted-tweets-ahead-huge-news-day-during-presidential-campaign-did-they-know-what-was-coming/

Varol, O., Ferrara, E., Davis, C. A., Menczer, F., & Flammini, A. (2017). *Online human-bot interactions: Detection, estimation, and characterization.* Retrieved from http://arxiv.org/abs/1703.03107

Wang, A. H. (2010). Detecting spam bots in online social networking sites: A machine learning approach. *Proceedings of the 24th Annual IFIP WG 11.3 Working Conference on Data and Applications Security and Privacy* (pp. 335–342). doi:10.1007/978-3-642-13739-6_25

Woolley, S. C., & Guilbeault, D. R. (2019). United States: Manufacturing consensus online. In S. C. Woolley & P. N. Howard (Eds.), *Computational propaganda: Political parties, politicians, and political manipulation on social media* (pp. 185–211). New York, NY: Oxford University Press.

Yang, J., & Counts, S. (2010). Predicting the speed, scale, and range of information diffusion in Twitter. *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media* (pp. 355–358). Retrieved from https://www.aaai.org/ocs/index.php/ICWSM/ICWSM10/paper/view/1468