The Brexit Botnet and User-Generated Hyperpartisan News

Marco T. Bastos & Dan Mercea

(City, University of London)

Abstract

In this paper we uncover a network of Twitterbots comprising 13,493 accounts that tweeted the U.K. E.U. membership referendum, only to disappear from Twitter shortly after the ballot. We compare active users to this set of political bots with respect to temporal tweeting behavior, the size and speed of retweet cascades, and the composition of their retweet cascades (user-to-bot vs. bot-to-bot) to evidence strategies for bot deployment. Our results move forward the analysis of political bots by showing that Twitterbots can be effective at rapidly generating small to medium-sized cascades; that the retweeted content comprises user-generated hyperpartisan news, which is not strictly fake news, but whose shelf life is remarkably short; and, finally, that a botnet may be organized in specialized tiers or clusters dedicated to replicating either active users or content generated by other bots.

Keywords: Brexit, Twitter, Fake news, Sockpuppets, Retweets cascades, Political bots

Introduction

A sockpuppet account is a false online identity used to voice opinions and manipulate public opinion while pretending to be another person. The term draws from the manipulation of hand puppets using a sock and refers to the remote management of online identities to spread misinformation, promote the work of one individual, endorse a given opinion, target individuals, and challenge a community of users (Zheng, Lai, Chow, Hui, & Yiu, 2011). Sockpuppet accounts are often automatic posting protocols (i.e., bots) operating under a fictitious identity and as such they breach the Terms of Service of social networking sites like Facebook and Twitter. The administration and deployment of bots and sockpuppet accounts are largely centralized and rely on trivial computing routines that allow users and organizations to control substantial subcommunities across any given social media platform (Kumar, Cheng, Leskovec, & Subrahmanian, 2017).

Concerns about the activity of bots and sockpuppets in the context of the U.K. E.U. Referendum were articulated in the press (Silva, 2016) and academia (Shorey & Howard, 2016), with researchers cautioning against the automation of political communication and the possible distortion of vital processes at the heart of contemporary liberal democracies, chief among which are competitive elections (Woolley & Howard, 2016). The scale of bot deployment and its effect on information diffusion are topical concerns (Bessi & Ferrara, 2016), with previous research reporting that bots are often deployed in contexts of polarized political discussion (Ferrara, Varol, Davis, Menczer, & Flammini, 2016). We seek to contribute to this growing body of scholarship by scrutinizing a large network of bots that operated during the Brexit debate. We explored the tactics employed by bot masters deciding which tweets are retweeted and by which subgroup of accounts linked to the botnet.

The referendum on Britain's membership of the European Union, specifically, was the fruit of more than four decades of efforts to extricate the country from the E.U. by political actors perennially suspicious of the supranational organization (Becker, Fetzer, & Novy, 2016). The vote was portrayed as a milestone in the political life of the country (Asthana, Quinn, & Mason, 2016) that opened up fault lines largely at odds with the traditional alignment of British political parties (Becker, et al., 2016). The observed political realignment foregrounds a socio-cultural cleavage between young and well-educated sections of the population who embrace progressive post-materialist values of equality, human rights, environmental protection, and a greater tolerance of immigrants; and on the other hand, an older, less educated demographic who witnessed both a decline in its material conditions and a gradual erosion of traditional values associated with industrial societies (Inglehart & Norris, 2017).

The political realignment and dealignment revealed by the U.K. E.U. Membership Referendum throws deeply engrained ideological leanings into question (Kriesi & Frey, 2008). It also feeds into a context of polarization, alternative media, and hyperpartisanship consistent with emerging patterns of news consumption (Bastos, 2016; Starbird, 2017). While the young and well-educated are significantly more likely to access news via social media (Ofcom, 2017), an older and less educated readership is traditionally associated with tabloids, which account for a substantive portion of the British press (Boykoff, 2008) and are making inroads on social media platforms (Newman, Richard Fletcher, Levy, & Nielsen, 2016). The partisan realignment observed in Britain, and the context of a highly polarized political climate underpinning the Referendum campaign, offered fertile ground for political actors interested in developing and deploying bots.

Canvassers have strategically explored these cleavages. The successful Vote Leave campaign was repeatedly and intensely chastised by policymakers and media pundits for stoking anxiety about immigration by making misleading pronouncements about Turkey's future E.U. membership. Equally, it was criticized for its disingenuous pledge to boost the National Health Service (NHS), a weakened pillar of the industrial welfare state in the U.K., by redirecting Great Britain's E.U. membership contribution into the service (Doherty, 2016; Swinford, 2016). Despite these embattled claims, Vote Leave canvassers effectively relied on data analytics (Cummings, 2016) to capitalize on this ostensible tension between the so-called circles of hard-working families and progressive elites, with a later analysis finding social media activity to be a positive predictor of the outcome of the vote (Celli, Stepanov, Poesio, & Riccardi, 2016).

In what follows, we first ground this study in the literature on political bots—i.e., bots deployed in a context of polarized political debate. We subsequently show that the group of Twitter accounts scrutinized in this paper often tweeted in a concerted fashion and could well be described as a botnet or a supervised network of zombie agents—be they internet-connected devices (Paulson, 2006) or social media accounts (Abokhodair, Yoo, & McDonald, 2015). As witnessed in other cases, this pool of accounts was used to automatically replicate posts on Twitter (Woolley, 2016). In contrast to previous research, we identified how the botnet is subdivided into subnetworks dedicated to retweeting content generated either by bots or humans, thereby engineering different retweet cascades. As we show below, the sophistication of the operation deviates considerably from traditional Twitterbots. Common to most accounts in this botnet was the curated replication of content that was both user-generated and a reproduction of tabloid journalism. Another important marker of this group was the overwhelming prominence of content associated with or authored by user accounts affiliated with the Vote Leave campaign.

The overall tone of the messages was much in line with the context of disaffection with immigration and the cultural backlash spearheaded by older, traditional, and less educated readership of tabloids (Boykoff, 2008). This cultural backlash was strategically leveraged and maximized by populist parties and leaders in order to promote "traditional cultural values and emphasize nationalistic and xenophobia appeals, rejecting outsiders and upholding old-fashioned gender roles" (Inglehart & Norris, 2016, p. 30). Our analysis makes no claim as to the veracity (or lack thereof) of the facts reported by the large network of users tweeting the Vote Leave campaign; instead, it seeks to shed light on this unusual user base tweeting the referendum by sourcing hyperpartisan content which is unlikely to fit a normative definition of fake news, but that is likely to have played a role in the emerging and loosely defined fake news ecosystem (Benkler, Faris, Roberts, & Zuckerman, 2017).

Twitterbots in Political Campaigns

The literature investigating bot activity is concerned with the imitation of human activity on social media by computer scripts (Bessi & Ferrara, 2016). These algorithms, also referred to as "social bots," have been shown to approximate (Woolley & Howard, 2016) and upscale human conduct (Bessi & Ferrara, 2016), often influencing communication exchanges on polarizing topics (Howard & Kollanyi, 2016). Social bots can be deployed in a wide variety of contexts and constitute a growing subfield of communication and political science research, which cautions against their detrimental impact on electoral politics, policy discussions, and deliberation of contentious issues. Indeed, prominent political events such as the referendum on the U.K.'s Membership of the European Union or the 2016 U.S. Presidential Elections were shown to have

been susceptible to such automated interference especially on Twitter (Bessi & Ferrara, 2016; Howard & Kollanyi, 2016).

Significant efforts have been made to detect patterns of activity that pertain to automation. Evidence to this effect points to the generation and republication of high volumes of partisan content with retweets—the practice of republishing a message already in circulation (Murthy & Dawsonera, 2013)—to boost the visibility of said content (Murthy et al., 2016); or alternatively, to corrupt communication (Woolley, 2016), particularly so as to create "a false sense of group consensus about a particular idea" (Ratkiewicz, Conover, Meiss, Gonçalves, Flammini, et al., 2011, p. 299). Another marker of account automation is the lack of detailed information about the user and the absence of geolocational metadata (Bessi & Ferrara, 2016) that could allow detection by users or social networking sites (Hwang, Pearce, & Nanis, 2012). Yet, bots can occupy an influential position in communication networks, often appearing at the center of highly connected network subgraphs (Bessi & Ferrara, 2016) in which information diffusion is centralized (Ratkiewicz, Conover, Meiss, Gonçalves, Flammini, et al., 2011).

Howard and Kollanyi (2016) approximated the use of political bots during the Brexit referendum by considering users that were extraordinarily active in the Twitter feed associated with the plebiscite. The authors reported that such users accounted for 32% of all Twitter traffic about Brexit. While acknowledging that there is no definitive way of telling which accounts are actual bots, they inferred that the top ten accounts producing the highest volume of messages (north of 350 tweets) were likely automated. Indeed, other research has described bot activity levels as incessant which on Twitter, specifically, translated into "excessive amounts of tweets" (Bessi & Ferrara, 2016). Nonetheless, user activity alone has been shown to be an unreliable metric to determine the presence of bots, as prolific Twitter posters can tweet abundantly by

taking turns managing Twitter accounts and pushing several hundred tweets a day with little to no automation (Bastos & Mercea, 2016; Mercea & Bastos, 2016).

Secondly, previous research has identified political bots to be tweeting at a rate of seven tweets per minute or 929 tweets in 138 minutes (Metaxas & Mustafaraj, 2010). In that reported instance, a small botnet comprising nine Twitterbots was set up in thirteen minutes to target accounts of interest by virtue of their previously expressed concern with the 2009 U.S. Massachusetts senate race. Those bots succeeded at starting cascades retweeted by posters whose political alignment resonated with the content of the message (Metaxas & Mustafaraj, 2010). Other studies looking into the same senatorial elections have shown that botnets can raise retweeted URLs to the top of Google search results (Ratkiewicz, Conover, Meiss, Gonçalves, Patil, et al., 2011).

However, the investigation by Bessi and Ferrara (2016) into the 2016 U.S. Presidential Elections found that overall humans posted more tweets than bots within the period under study. Furthermore, bots seemed inept at interacting with humans, replying (i.e. by quoting one's Twitter handle preceded by the @ character) primarily to other bots. In their turn, humans were replying to humans more than to bots, another marker that humans and bots operate in largely disconnected subgraphs. These results are somewhat at odds with the political botnet studied by Metaxas and Mustafaraj (2010), which directed replies at recipients purposefully selected for their partisan interest in the Massachusetts elections, a quarter of whom went on to retweet the automated message they received (Mustafaraj & Metaxas, 2017). The variability of @-mentioning and retweeting practices indicate that bot masters are likely implementing a range of different strategies depending on the political objective set for the botnet.

Notwithstanding, the bots recorded in the 2016 U.S. Presidential Elections were effective information disseminators. They were just as apt as humans at retweeting, republishing a similar volume of content to humans. Similarly, in the case of the E.U. referendum, the most prolific accounts did "not generate new content but simply retweeted content from other users" (Howard & Kollanyi, 2016). While the latter scholars conceded that human agents could achieve similar levels of activity if they confined themselves exclusively to retweeting, Bessi and Ferrara (2016) cautioned that bots could have a debilitating effect on human communication because of their noted capability to disseminate content among human users.

Finally, Howard and Kollanyi (2016) claim that in the E.U. referendum bots were designed to take sides in the debate about the U.K.'s membership of the European Union. Similarly, Bessi and Ferrara (2016) determined that bots produce systematically more positive content in support of a candidate, a fact they submitted can distort perceptions of support for that candidate. Similarly, Woolley (2016) posited that accounts exhibiting bot activity feature prominently in Twitter "bombs:" streams of tweets that flood into hashtags used by opponents and are retweeted by bots so as to disrupt the communication and organization of the opposite side. In summary, the literature on political bots offers a growing catalogue of metrics for pinpointing political bots (Abokhodair, et al., 2015; Ferrara, et al., 2016; Ratkiewicz, Conover, Meiss, Gonçalves, Patil, et al., 2011; Ratkiewicz, Conover, Meiss, Gonçalves, Flammini, et al., 2011), but little is known about the actors controlling these bots, the decisions on which tweets are to be retweeted, and the type of content relayed by such accounts (Nied, Stewart, Spiro, & Starbird, 2017).

Research Objectives

With this study we seek to identify a large network of bots that tweeted the Brexit debate and the type of content relayed by these accounts. We explore bot activity with insights into the prevalence of hyperpartisan and polarizing content (Benkler, et al., 2017), which constitutes our first Research Objective (RO1). To this end, we began with an inspection of the webpages attached to tweets to identify the domain name of websites sourcing information to bots. Next we hypothesized that bot activity would be marked by a high-volume posting signature followed by a drop in activity levels characteristic of the lifecycle of bots (RO2). To this end, we conducted a time series analysis and modelled the mean cascade time of bots and active users in the Twitter referendum data, thus distinguishing seasonal patterns and the posting behavior of political bots.

Research Objectives RO3 and RO4 probe the impact of bot communication on the Brexit debate. We inspected our dataset to determine whether bots could generate greater message cascades than active users (RO3). In close connection, we calculated the maximum, minimum, and mean cascade time to ascertain if bots triggered faster cascades than active users in the network (RO4). We thus scrutinized the impact of bots as an upshot of the intensity, reach, and speed of their activity, in addition to examining their network influence and the information dissemination patterns that characterized their actions during the last month of the E.U. referendum campaign. We relied on such metrics to contrast the activity patterns of bots with regular users, as well as of different types of bots operating within the same botnet.

We subsequently inquired whether the accounts that were swept up in the retweet cascades were also bots themselves. We envisioned that the impact of a botnet may depend on whether it is embedded in a larger network of active users or, alternatively, restricted to a cluster of bots (RO5). Our hypothesis was that the more engagement with human agents the botnet

generates, the more likely it is to widen cascades beyond the botnet. In other words, we would expect botnets to exhibit levels of human curation (Howard & Kollanyi, 2016) that testify to their differentiated optimization and their fundamentally cyborg nature. We adopt the latter term to reflect on the close coupling of human agency and computer scripts characterizing bots that disruptively amplifies human communication (Asenbaum, 2016).

Lastly, we hypothesized that the botnet is subdivided into various subgroups dedicated to retweeting specific accounts, thereby triggering different types of retweet cascades. To this end, we examined retweet activity to distinguish patterns of human and bot activity as well as interactions between them that could evince strategies of bot deployment (RO6). This last Research Objective seeks to establish whether bots are deployed and operate in a concerted fashion; or, alternatively, whether competing strategies are employed to overcome the enduring risk that the botnet will be trapped in so-called "echo chambers," i.e., groups of bots self-referentially communicating with each other.

Methods and Data

We queried the Twitter Streaming API to monitor 39 Twitter hashtags clearly associated with the referendum campaign from April to August 2016 (e.g. #voteleave, #voteremain, #votein, #voteout, #leaveeu, #bremain, #strongerin, #brexit, #euref, etc.). For the purposes of this study, we focus on the two-week period before and after the referendum vote, i.e., June 10 to July 10, 2016. In this interval, we collected approximately 10M tweets associated with the referendum. We subsequently retrieved the profile of over 800K unique users that appeared in our dataset and relied on thresholding and filtering approaches to disentangle real users from bots (Table 2). The combination of methods reported in the literature (Subrahmanian et al., 2016; Varol, Ferrara,

Davis, Menczer, & Flammini, 2017) allowed us to identify a large group of bots whose accounts had been deactivated by the bot master or blocked/removed by Twitter in the aftermath of the referendum. We relied on the implementation of extended regular expression in R (2014) to identify the campaign associated with tweets and the libcurl implementation (Temple Lang, 2016) to retrieve the webpage title of URLs embedded in tweets (when available).

Previous research found the frequentist approach to user activity alone to be an unreliable metric to determine the presence of bots, as prolific Twitter posters can tweet abundantly by taking turns and pushing several hundred tweets a day with little to no automation (Bastos & Mercea, 2016; Mercea & Bastos, 2016). Consequently, in the attempt to differentiate between bots and high-volume posters, we analyzed several metrics of user activity in addition to the temporal posting patterns of potential bots. This composite analysis allowed us to ascertain if their activity endured over time; or conversely, if there was a notable drop in activity levels that might typify what may be described as a "bot lifecycle," in the wake of the E.U. referendum.

The metrics used in this study to identify bot accounts are informed by the relevant literature and include detailed profile information, presence or absence of geographical metadata (or propensity to post using web clients), retweet to tweets ratio, @-mention to tweet ratio, activity level, followers to followees ratio, account creation date, and absence of known words in the username (Table 2). Positive predictors of bot activity are shown in Table 2 and include tweets to user (tw2user), mean tweet to retweet (tw2rtMean), common words in the username (commonWords), use of web interface to relay content (webClient), ratio of outbound to inbound @-mentions (mentionOut2In), ratio of inbound to outbound retweets (retweetIn2Out), account creation date (newAccount), retweet reciprocity (rtReciprocity), and retweet cascade mean time (ccdMeanTime). For the purposes of this study, we contrast retweeting behavior observed in this

group against the larger set of accounts we refer to as active users (as opposed to deactivated or recycled users, defined underneath). Retweet and @-mention behavior are defined as $A \rightarrow B$ when B retweets A and $A \rightarrow B$ when A mentions B (thus following the directionality of the information flow). While previous studies have explored Twitter cascades by tracking the diffusion of URLs (Bakshy, Hofman, Mason, & Watts, 2011) and hashtags (González-Bailón, Borge-Holthoefer, Rivero, & Moreno, 2011), we rely on retweets to inspect user-to-bot and bot-to-bot cascade composition.

Unfortunately, it is not currently possible to rebuild every step of the retweet cascade, as each retweet includes only a reference to the original message, so that if user C retweets user B who has previously retweeted user A, we can only establish that user A was retweeted by user C, with the intermediary steps of the cascade remaining unknown. As such, we cannot account for independent entry points that might have influenced the cascade (Cheng, Adamic, Dow, Kleinberg, & Leskovec, 2014). However, given that each retweet includes a unique identifier arranged chronologically from the original tweet to the most recent retweet, we can rebuild cascades from the seed message to the retweets that have cascaded from that original content. Similarly, we rely on the timestamp attached to each tweet to estimate the variable time-to-retweet, calculated as the time elapsed between the original tweet and the ith retweet for cascade of size S.

Ethical Considerations

The research data examined in this study was collected via the publicly accessible Twitter Streaming and REST APIs. Although the information collected for this study is public, there are important ethical issues associated with harvesting public Twitter accounts (Zimmer, 2010).

Twitter profiles set to private were removed from our pool of users and no private information was examined in the analysis. While we have looked to preserve users' rights and interests, we ultimately decided to disclose the Twitter handles examined in this study whenever there was a reasonable level of certainty that we were dealing with Twitterbots, to which ethical considerations of privacy are immaterial. We also considered the potential sensitivity of some of the tweets examined in this study, but anonymizing the seeding accounts would impinge on our ability to understand the scope of the botnet and the strategies adopted by bot masters. Lastly, we have considered the ethical obligation not to display deleted Tweets, but we believe the content analyzed in this study is of public and scholarly interest, and that ethical considerations regarding users' rights to not have their deleted tweet made public are immaterial in the context of large botnets participating in politically contentious debates.

Results

From a total of 794,949 Twitter profiles that tweeted the Vote Leave and Vote Remain campaigns, we managed to identify the location of 60% of them (482,193) by triangulating information from geocoordinates embedded in tweets (i.e. reverse geocoding), geographic information tweeted by the users, and information that appeared in their profiles. From this cohort of users, only 30,122 users were identified as based in the U.K., a smaller population than the set of 40,031 accounts that have been deactivated, removed, blocked, set to private, or whose username was altered after the referendum. This latter group of accounts represents 5% of all users that tweeted the referendum and is divided as follows: 66% or 26,538 were users who have changed their username since the referendum but remained active on Twitter (designated

hereafter as repurposed or recycled accounts); 34% or 13,493 accounts were suddenly blocked or removed themselves from Twitter (deleted accounts).

Although repurposed/recycled accounts conspicuously interacted with deleted Twitterbot accounts, the focus of this study lies with the latter cohort. Notwithstanding, common to these two subgroups is the predominance of retweeted content that disappeared from the internet shortly after the referendum. Another commonality is the notable support for the Leave campaign, measured by the relative frequency of keywords and hashtags associated with each of the campaigns. While the ratio of messages using hashtags that supported the Leave and Remain campaigns was 31% and 11% for the entire network, recycled and removed accounts combined tweeted the Referendum hashtags to a ratio of 37% and 17% (or 2,434,077 and 840,726 versus 30,947 and 14,390 tweets for each of the campaigns, respectively). In what follows, we disentangle these groups to finally concentrate on a set of 13,493 accounts identified as bots.

Hyperpartisan and Perishable News

By annotating tweets using textual markers such as hashtags and keywords associated with the Leave and Remain campaigns, we found that the proportion of tweets supporting the Vote Leave campaign in the pool of removed accounts was yet higher, at 41% compared with 31% for active users, with the proportion of neutral tweets also being higher in the latter. Slogans associated with the Vote Leave campaign were also significantly more likely to have been tweeted by this pool of accounts in a ratio of 8:1. This subset of removed accounts was considerably more active in the period leading up to the referendum, with an average of 4.4 messages compared with 3.9 for the rest of the population ($\bar{x}=4.44$ $\sigma=33.3$, and $\bar{x}=3.99$ $\sigma=74.2$, respectively); and also less

active in the wake of the vote with an average of 2.4 tweets compared with 2.6 for the global population ($\bar{x}$=2.42 σ=9.0, and $\bar{x}$=2.61 σ=63.2, respectively).

Upon attempting to retrieve the webpages (RO1) tweeted by recycled and removed accounts, we found that most tweeted URLs (55%) no longer exist, cannot be resolved, or link to either a Twitter account or a webpage that no longer exists. Nearly one third (29%) of the URLs link to Twitter statuses, pictures, or other multimedia content that is no longer available and whose original posting account has also been deleted or blocked, a marker of the perishable nature of digital content at the center of political issues (Walker, 2015). From this total, 1% of all links was directed to user @brndstr, one of the few accounts appearing in the communication network of recycled accounts that remains active under the same username. This account is managed by Dubai-based "Bot Studio for Brands," a company specialized in providing bots for social media campaigns.

A closer inspection of the accounts sourcing content to the pool of recycled and removed accounts reveals the markedly short shelf life of user generated content. These are Twitter accounts invested in spreading dubious news stories sourced from a circuit of self-referencing *blews* (Gamon et al., 2008): a combination of far-right weblog such as WorldTribune.com and traditional tabloid media such as express.co.uk. However, the few webpages we managed to retrieve indicate that the content tweeted by this large pool of recycled and removed accounts does not conform with the notion of fake news designating news stories that are intentional, misleading half-truths and/or outright lies (Benkler, et al., 2017). Instead, the content is in line with a form of storytelling that blurs the line between traditional tabloid journalism and user-generated content, which is often anonymous, fact-free, and with a strong emphasis on simplification and spectacularization (Rowe, 2011). User-generated content takes the lion's share

of hyperlinks tweeted by recycled and removed accounts. The content is often presented as a

professionally-looking newspaper by resorting to content curation services such as paper.li and is

likely to include Twitter multimedia (e.g., Twitter's native multimedia sharing service

twimg.com).

Similarly, the few links that remained accessible six months after the referendum can

hardly be described as fake news. The hyperlinked material is rich in rumors, unconfirmed

events, and human-interest stories with an emotional and populist appeal that resembles tabloid

journalism, except for the added complexity that audiences play a pivotal role in curating and

distributing the content. The sources we managed to inspect, though not representative of the

much larger universe of content tweeted by this population of users, and that unfortunately has

mostly vanished from Twitter, is much akin to hyperpartisan tabloid journalism, with a topical

emphasis on highly-clickable, shareable, and human-interest driven stories (Bastos, 2016). Table

1 summarizes the URLs tweeted by this cohort of users.

TABLE 1 ABOUT HERE

Although 17% of weblinks point to Twitter accounts that are still active, a random sample shows

that the original message is frequently no longer available, thus preventing any determination of

the nature of the content originally tweeted. A good example is the tweet ID

740138870092750848 which generated a cascade of several hundred retweets and whose posting

user is still active. Although the user account seeding the cascade remains active, the original

tweet has been removed (together with the relevant retweet cascade). With Internet Archive

having no record of this specific tweet, it is no longer possible to know what the original image

conveyed. The scale of deleted content applies both to weblinks tweeted by this population as

well as to user accounts, a worrying development given the importance and contentious nature of the referendum (Walker, 2015).

Brexit Botnet

Turning to the removed user accounts, we relied on metrics discussed in the relevant literature (Bessi & Ferrara, 2016; Ratkiewicz, Conover, Meiss, Gonçalves, Patil, et al., 2011) to determine whether the pool of deleted accounts comprised a large network of Twitterbots. Upon inspecting the account creation date, we found that 83% of accounts in the botnet had been created in the previous 2 years compared with 43% for the subset of active accounts and 48% for accounts that ended up being recycled. We interpreted the result as an important indication of bot activity. The highest rate of tweet to retweet was found in the campaign accounts @iVoteStay and @iVoteLeave, with a retweet to tweet ratio of 90% and a total number of retweets of nearly 60K messages. These accounts did author original content though, and we do not feel confident they can be classified as bots despite the extraordinary high levels of activity and the high likelihood that some form of automation was used to relay content.

When analyzing retweet rate across groups, we found that the baseline for accounts that remained active after the referendum was of one retweet to each 3 tweets ($\tilde{x}=0.33$ and $\bar{x}=0.45$), while the ratio for accounts that changed their usernames is twice as high ($\tilde{x}=0.61$ and $\bar{x}=0.54$). The group that significantly deviates from this baseline is the set of accounts removed after the referendum (i.e., the botnet). For such accounts, the retweet rate is of 1 retweet for every tweet, with 54% of accounts never having authored any tweet related to the referendum (i.e., only retweets were registered for these accounts, another marker of bot activity). For this group of accounts, the median tweet to retweet rate is 1 ($\bar{x}=0.6302$). Table 2 shows the metrics used to

classify this subset of accounts as bots in a network of 13,493 Twitterbots that tweeted a total of 63,797 messages. The variables in the table indicate tweets per user, retweet to tweet ratio, incidence of known words in usernames, propensity to post using web clients, @-mention indegree to outdegree, retweet indegree to outdegree, account creation date, outdegree and indegree transitivity, retweet reciprocity, modularity score, mean and maximum cascade size, number of cascades triggered, share of triggered cascades, and cascade mean time.

TABLE 2 ABOUT HERE

Botnet Lifecycle

After establishing that this subnet consists primarily of automatic posting protocols, we approached RO2 by contrasting activity levels in active accounts and the botnet. While the activity of active users presents the usual seasonal patterns associated with Twitter activity, including the ebb and flow of messages resulting from daily patterns associated with work and leisure (Puschmann & Bastos, 2015), tweets posted by bots follow no such variation. The absence of seasonal patterns provides an indication that the time signature observed in the botnet deviates considerably from the remainder of the user base tweeting the referendum. Moreover, botnet activity is marked by a higher level of activity compared with active users in the period leading up to the referendum, and a sharp decrease immediately afterwards, an indication that this pool of accounts was either retired or expelled from the platform. Figure 1 shows the tweet activity for the period of June 10 to July 10 (two weeks before and after the referendum of June 23), with botnet activity markedly higher in the period before the referendum followed by a sharp decline thereafter.

FIGURE 1 ABOUT HERE

The histogram in Figure 1 presents a sharp decline in bot activity immediately after the referendum, much in line with the assumptions surrounding RO2. In addition to the short timespan the botnet was kept alive, the structural network of retweets disseminated by Twitterbots shows greater integration when compared with the network of active users, with significantly lower measures of transitivity (i.e., measure of clustering) and higher modularity (i.e., tendency for a network to be organized into subnetworks, as shown in Table 2) due to an abundance of hub-to-spoke formations found across the bot subnet. The botnet appears organized in specialized tiers dedicated to replicating tweets originating either from active users or bots. Such formations are abundant but restricted to bot-retweet-bot and bot-retweet-user formations, with little crossover between such formations.

Beyond hashtags, the content tweeted by bots presents a clear slant towards the Leave campaign, with the three most frequent words being "Brexit," "referendum," and "VoteLeave," which together account for 9% of all terms tweeted by the botnet. In comparative terms, 31% of tweets posted by bots included the term "leave" as opposed to 16% for the group of active users. Remarkably, 17% of tweets posted by the botnet also included the term "remain" compared with only 11% for the group of active users, an indication that although the botnet tweeted predominantly pro-Leave messages, it also tweeted more pro-Remain campaign messages. The result alludes to the patent partisanship espoused by such accounts in comparison with the larger community of users tweeting the referendum.

We addressed RO3 by analyzing the tweet-to-retweet ratio and the cascade mean time across groups, which provides an indicator of whether bots were successful at generating large retweet cascades. We found that despite the imbalance in the ratio of retweeted messages for these accounts (Table 2), bot activity was largely successful to the extent that it attained a

retweet ratio of 2:1 between active users and bots. In other words, the much larger pool of active users retweeted on average every second message posted by bots. This estimate appears unusually high and we detail the causes for this surprisingly high ratio when addressing RO5 and RO6.

Retweet Cascades

The largest retweet cascade (*S*=13,417) was authored by a user making a direct reference to @brndstr, the Dubai-based startup specialized in social media bots. The message read "I #VoteIn for the #Brexit #EURef vote with @Brndstr & unlocked my own Flag Profile pic! What will you vote? #ivoted https://t.co/iFlZyhrzLd" (the link directs to @brndstr Twitter account). Another tweet with the same content but starting with "I #VoteOut for the #Brexit #EURef" is the third largest cascade in the data. In short, @brndstr messages were directed at both Leave and Remain campaigns with the purpose of placing the company within the larger conversation, automated or otherwise. Despite these large cascades, @brndstr is connected to a relatively small number of cascades, particularly in view of the magnitude of the botnet which includes over 10K accounts devoted to supporting one or the other side of the referendum campaign.

FIGURE 2 ABOUT HERE

We could not reject the assumption underlying RO3 that bots would be less effective at triggering retweet cascades. In fact, the botnet is considerably more effective at joining successful cascades, with a mean cascade participation of $\bar{x}=18$ (compared with $\bar{x}=6$ for active users). The botnet is also just as effective at generating large retweet cascades compared with active, regular Twitter users and/or accounts that have been repurposed during the period of this study. While such regular users started 36% of cascades, the botnet claimed a total of 30% of

20

cascades (34% were generated by accounts repurposed at the end of the referendum period). Interestingly, the same long-tailed distribution of hyperactive accounts among regular users is observed in the botnet, where a small share of bots was found to have triggered most retweets, with the remainder of the bots being strategically although peripherally positioned to retweet the initial cascade (Figure 2). While the group of active users tweeted 97% of the messages and initiated 7.5% of all cascades, the botnet tweeted under 1% of the total messages in the dataset, but accounted for a comparable share of 6% of all cascades.

We approached RO4 by comparing the distribution of user activity and cascade mean time between active users and bots. As is often the case with Twitter data, user activity, hashtag use, and cascade time (calculated as the average time T for a cascade of size S) follow a long-tailed distribution. User activity and cascade time are particularly skewed in the botnet subgraph, whose power-law distribution comprises a single user responsible for 4% of all tweets (@trendingpls) followed by a set of four accounts which together account for an additional 5% of the activity (namely, @EuFear, @steveemmensUKIP, @uk5am, and @no_eusssr_thx). These few accounts have been removed or deleted after the referendum, although the accounts @trendingpls and @uk5am have been recreated in October 2016 and now operate under a new user ID. Beyond these highly-active accounts, the botnet presents a mean of 5 tweets per bot (compared with 1.2 per human user).

Despite the different distributions, the average cascade time is comparable between the two groups, with botnets starting and completing cascades of size 5, 10, and 20 retweets just one minute faster than active users. The mean cascade time provides an indication that bots mimic the average timespan of retweet cascades, or more likely, that they retweet real-world accounts to maximize exposure to the message or to the user posting the original content (RO4). In fact,

taken together, botnet and the active user groups have not just similar cascade size and time, but also similar averages for cascade mean time at 69h for active users and 65h for the botnet. For medium-sized cascades (S=40, S=80, and S=160), the botnet completes the cascade 20 minutes faster, but it is with large cascades—S>320 and S<640—that larger temporal differences are observed, with the botnet completing such large cascades 1.5-2 hours faster than the active user base. Figure 3 unpacks this relationship: while cascade time linearly grows for the active user base, the fitted linear regression for cascades seeded by bots, and the bulk of the observations, falls close to the baseline of just a few minutes or hours.

FIGURE 3 ABOUT HERE

We approached RO5 by examining the difference between the average bot cascade and the outliers in the botnet. These differences shed light on the existence of at least two tiers or clusters of fundamentally different bots. The first group is dedicated to replicating automated content, hence achieving a much faster cascade turnaround compared with active user-generated cascades. The second group is deeply embedded in human-driven activity. These two types of bot activity are depicted in Figure 4 which foregrounds the substantially different groups of bots. While the formation on the left is geared towards replicating content exclusively from active users, the formation on the right is dedicated to replicating content seeded by other automatic posting protocols. Both accounts succeeded at generating medium (S>50) and large cascades (S>100), but their typical retweeting patterns indicate they were created and deployed to meet fundamentally different objectives.

FIGURE 4 ABOUT HERE

This finding sheds light on RO6, which seeks to distinguish patterns of human and bot activity along with interactions that might cut across the two subgraphs. While the first subset of bots

was associated with accounts that leveraged retweet behavior to amplify the reach of a small set of users and rarely if ever started any cascade themselves, the other subset of bots had a narrower scope of operation only retweeting other bots in the botnet and thereby producing many medium-sized cascades that spread significantly faster than the remainder of the cascades. As shown in Figure 4, both bots deployed to retweet active users and bots developed to retweet other bots exhibit different retweeting patterns. Although both @trendingpls and @nero are bots, the first only retweets active users whereas the retweet activity of the latter is restricted to other bots, likely deployed in conjunction with the head node. Each of the bot subnets play a specialized role in the network, and both feed into the larger pool of regular accounts brokering information to @vote_leave, the official Twitter account of the Vote Leave Campaign, and arguably the most prominent point of information diffusion associated with the Leave Campaign.

As to its time distribution, the retweet activity was mostly concentrated in the period leading up to the referendum vote (Figure 1). Most of it consisted of organic retweets from and to accounts in the active user base. Bots operated in the same period both by retweeting active users and retweeting other bots, chiefly in the week preceding the vote (June 16-23) and in the eve of the referendum (June 21), when we observed a peak in retweet activity between bots. There was a sharp decline in retweet activity after the referendum, principally among active users who ceased to trigger or join retweet cascades. On the other hand, bots remained operational and activity peaks are observed on Jul 12-15: first retweeting active users, then replicating bot content, only to tail off in the following weeks when the botnet is retired, deactivated, or removed entirely from the Twitter platform. In fact, head nodes of the bot-to-bot subnet such as @NoThanksEU, @wnwmy, @Foresight1st, @nero, @horrorscreens00, and @Dugher101 disappear after the end of the referendum (only @NoThanksEU was reactivated in

November 2016). This is the critical period (June 2016) when content tweeted by such bots and the webpages linked to their tweets disappeared from the internet and Twitter public and enterprise APIs (search API and GNIP, respectively).

FIGURE 5 ABOUT HERE

The timewise evolution of large cascades in Figure 5 sheds light on RO6 by showing the peak of retweet activity between active users, between bots, and the hybrid variant of bots retweeting active users which account for a considerable portion of retweet cascades in the botnet. Notwithstanding the apparent success achieved by the botnet, we caution against extrapolating the analysis of cascade triggering to the effects these bots might have had on the referendum debate. Although the botnet can trigger small and medium-sized cascades effortlessly, and even participate in large ones, bots are particularly ill-suited to starting large cascades due to the unique constraints associated with growing and maintaining a botnet. In fact, the botnet failed to generate any large cascade of 1K retweets, while the active user base successfully generated nearly one hundred such cascades. Similarly, while the botnet generated only 5 cascades of $S > 500$, the active user base generated as many as 237 of them. Thus, we advise caution in diagnosing the potential impact of a botnet on the larger conversation on Twitter. Although these accounts can impinge on the broader conversation and boost the reach of a subset of users, we have yet to find evidence that they can manipulate the Twittersphere.

Conclusion

In this paper we uncovered a network of Twitterbots comprising 13,493 accounts that tweeted the U.K. E.U. membership referendum which were deactivated or removed by Twitter shortly after polling stations closed. We have shown that the botnet tweeted mainly messages supporting

24

the Leave campaign and argued that such bots may likely be repurposed from one campaign to the next by the social media analytics outfits. The botnet can thus be operated as an army of sockpuppet accounts deployed to amplify a defined group of user by aggregating and retweeting content tweeted by seed users, which may conceivably be bots themselves, a process that corporate literature refers to as "false amplification" chiefly orchestrated by "fake accounts" (Weedon, Nuland, & Stamos, 2017).

Nonetheless, our analysis has not found evidence of widespread fake news diffusion with political bots. Instead, we found a combination of what appears to be a Twitter botnet feeding and echoing user-curated and hyperpartisan information. The material that remained available after the referendum points to a significant milestone in tabloid journalism, which is incorporating audience feedback while undergoing a transition from the strong editorial identity of tabloid newsprint to content curation that is both user-generated and created by paid staff members (Bastos, 2016). The hyperpartisan content pushed by the botnet epitomizes an ongoing trend to push viral content that is mostly short, shareable, accessible with mobile devices, and that accentuates polarized identities and balkanizes readerships into like-minded groups (Sunstein, 2001).

The botnet exhibited clear patterns of specialization with sets of accounts dedicated to retweeting active users and another set of bots positioned to echo campaign slogans and follow communication tactics directed by other bots. The likely overhead involved in setting up such a specialized botnet pays off during the automation of retweets, which allow the botnet to trigger small to medium-sized cascades in a fraction of the time required by active users to start cascades of comparable size. Despite the botnet's capacity to rapidly trigger such cascades, we have not found evidence supporting the notion that bots can substantively alter campaign

communication, as the activity of the botnet—at least of this defunct botnet in particular—was relatively minor with respect to the larger conversation about the referendum that took place on Twitter.

Author Information

Marco Bastos is Lecturer in Media and Communication at City, University of London and an affiliate of Duke University's Network Analysis Center and the Data Science Initiative of the University of California at Davis. Email: Marco.Bastos@city.ac.uk

Dan Mercea is Senior Lecturer in the Department of Sociology at City, University of London. He holds a PhD in Communication Studies from the Department of Sociology, University of York. Email: Dan.Mercea.1@city.ac.uk

References

Abokhodair, N., Yoo, D., & McDonald, D. W. (2015). *Dissecting a Social Botnet: Growth, Content and Influence in Twitter.* Paper presented at the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing, Vancouver, BC, Canada.

Asenbaum, H. (2016). Cyborg activism: Exploring the reconfigurations of democratic subjectivity in Anonymous. *New Media & Society, 0*(0), 1461444817699994. doi: doi:10.1177/1461444817699994

Asthana, A., Quinn, B., & Mason, R. (2016). UK votes to leave EU after dramatic night divides nation *The Guardian*. Retrieved from

https://www.theguardian.com/politics/2016/jun/24/britain-votes-for-brexit-eu-referendum-david-cameron

Bakshy, E., Hofman, J. M., Mason, W. A., & Watts, D. J. (2011). *Identifying influencers on twitter.* Paper presented at the Fourth ACM International Conference on Web Seach and Data Mining (WSDM).

Bastos, M. T. (2016). Digital Journalism and Tabloid Journalism. In B. Franklin & S. Eldridge (Eds.), *Routledge Companion to Digital Journalism Studies*. London.

Bastos, M. T., & Mercea, D. (2016). Serial Activists: Political Twitter beyond Influentials and the Twittertariat. *New Media & Society*. doi: 10.1177/1461444815584764

Becker, S., O., Fetzer, T., & Novy, D. (2016). Who voted for Brexit? A Comprehensive District-Level Analysis (pp. 1-69). Coventry: Centre for Competitive Advantage in the Global Economy, University of Warwick.

Benkler, Y., Faris, R., Roberts, H., & Zuckerman, E. (2017). Breitbart-led right-wing media ecosystem altered broader media agenda: Columbia Journalism Review.

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

Boykoff, M. T. (2008). The cultural politics of climate change discourse in UK tabloids. *Political Geography, 27*(5), 549-569. doi: http://dx.doi.org/10.1016/j.polgeo.2008.05.002

Celli, F., Stepanov, E. A., Poesio, M., & Riccardi, G. (2016). *Predicting Brexit: Classifying Agreement is Better than Sentiment and Pollsters.* Paper presented at the Workshop on Computational Modeling of People's Opinions, Personality, and Emotions in Social Media (PEOPLES), Osaka, Japan.

Cheng, J., Adamic, L. A., Dow, P. A., Kleinberg, J., & Leskovec, J. (2014). *Can cascades be predicted?* Paper presented at the 23rd International Conference on World Wide Web (WWW'14), Seoul, Korea.

Cummings, D. (2016). On the Referendum #20: The Campaign, Physics, and Data Science - Vote Leave's 'Voter Intention Collection System' (VICS) now available for all. Retrieved from https://dominiccummings.wordpress.com/2016/10/29/on-the-referendum-20-the-campaign-physics-and-data-science-vote-leaves-voter-intention-collection-system-vics-now-available-for-all/

Doherty, M. (2016, July 15, 2016). Should Making False Statements in a Referendum Campaign Be an Electoral Offence? *Lancashire Law School*.

Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2016). The Rise of Social Bots. *Communications of the ACM, 59*(7), 96-104. doi: 10.1145/2818717

Gamon, M., Basu, S., Belenko, D., Fisher, D., Hurst, M., & König, A. C. (2008). *BLEWS: Using blogs to provide context for news articles.* Paper presented at the 2nd International AAAI Conference on Weblogs and Social Media (ICWSM'08), Seattle, WA.

González-Bailón, S., Borge-Holthoefer, J., Rivero, A., & Moreno, Y. (2011). The Dynamics of Protest Recruitment through an Online Network. *Scientific Reports, 1*. doi: 10.1038/srep00197

Howard, P. N., & Kollanyi, B. (2016). *Bots, #StrongerIn, and #Brexit: computational propaganda during the UK-EU Referendum*. SSRN.

Hwang, T., Pearce, I., & Nanis, M. (2012). Socialbots: voices from the fronts. *interactions, 19*(2), 38-45. doi: 10.1145/2090150.2090161

Inglehart, R. F., & Norris, P. (2016). *Trump, Brexit, and the Rise of Populism: Economic Have-nots and Cultural Backlash*. Paper presented at the American Political Science Association Annual Meeting, Philadelphia, USA. https://research.hks.harvard.edu/publications/getFile.aspx?Id=1401

Inglehart, R. F., & Norris, P. (2017). Trump and the Xenophobic Populist Parties. *Perspectives on Politics*(June).

Kriesi, H., & Frey, T. (2008). The United Kingdom: moving parties in a stable configuration. In E. Grande, H. Kriesi, M. Dolezal, R. Lachat, S. Bornschier & T. Frey (Eds.), *West European Politics in the Age of Globalization* (pp. 183-207). Cambridge: Cambridge University Press.

Kumar, S., Cheng, J., Leskovec, J., & Subrahmanian, V. S. (2017). *An Army of Me: Sockpuppets in Online Discussion Communities*. Paper presented at the Proceedings of the 26th International Conference on World Wide Web, Perth, Australia.

Mercea, D., & Bastos, M. T. (2016). Being a Serial Transnational Activist. *Journal of Computer-Mediated Communication, 21*(2), 140–155. doi: 10.1111/jcc4.12150

Metaxas, P., & Mustafaraj, E. (2010). *From Obscurity to Prominence in Minutes: Political Speech and Real-Time Search.* Paper presented at the Proceedings of the WebSci10: Extending the Frontiers of Society On-Line, Raleigh, NC: US.

Murthy, D., & Dawsonera. (2013). *Twitter: social communication in the Twitter age*. Cambridge: Polity.

Murthy, D., Powell, A. B., Tinati, R., Anstead, N., Carr, L., Halford, S. J., et al. (2016). *Automation, Algorithms, and Politics| Bots and Political Influence: A Sociotechnical Investigation of Social Network Capital* (Vol. 10).

Mustafaraj, E., & Metaxas, P. T. (2017). *The Fake News Spreading Plagues: Was it Preventable?* . arXiv.org.

Newman, N., Richard Fletcher, Levy, D. A. L., & Nielsen, R. K. (2016). Reuters Institute Digital News Report 2016: Reuters Institute for the Study of Journalism.

Nied, A. C., Stewart, L., Spiro, E., & Starbird, K. (2017). *Alternative Narratives of Crisis Events: Communities and Social Botnets Engaged on Social Media*. Paper presented at the Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing, Portland, Oregon, USA.

Ofcom. (2017). News consumption in the UK: research report (pp. 94). New York: The Office of Communications.

Paulson, L. D. (2006). Hackers Strengthen Malicious Botnets by Shrinking Them. *Computer, 39,* 17-19.

Puschmann, C., & Bastos, M. (2015). How Digital Are the Digital Humanities? An Analysis of Two Scholarly Blogging Platforms. *PLoS ONE, 10*(2), e0115035. doi: 10.1371/journal.pone.0115035

R. (2014). R: A Language and Environment for Statistical Computing (Version 3.0.3). Vienna, Austria: CRAN. Retrieved from http://www.R-project.org

Ratkiewicz, J., Conover, M., Meiss, M., Gonçalves, B., Patil, S., Flammini, A., et al. (2011). *Truthy: mapping the spread of astroturf in microblog streams.* Paper presented at the 20th International Conference Companion on World Wide Web, Hyderabad, India.

Ratkiewicz, J., Conover, M. D., Meiss, M., Gonçalves, B., Flammini, A., & Menczer, F. (2011). *Detecting and Tracking Political Abuse in Social Media.* Paper presented at the 5th International AAAI Conference on Weblogs and Social Media (ICWSM11), Barcelona.

Rowe, D. (2011). Obituary for the newspaper? Tracking the tabloid. *Journalism, 12*(4), 449-466. doi: 10.1177/1464884910388232

Shorey, S., & Howard, P. N. (2016). Automation, Algorithms, and Politics| Automation, Big Data and Politics: A Research Review. [literature review, critical, big data, algorithms, politics, automation]. *International Journal of Communication, 10*, 5032–5055.

Silva, S. (2016). Trump's Twitter Debate Lead Was 'Swelled by Bots' Retrieved 15 December 2016, 2016, from http://www.bbc.co.uk/news/technology-37684418

Starbird, K. (2017). *Examining the Alternative Media Ecosystem through the Production of Alternative Narratives of Mass Shooting Events on Twitter.* Paper presented at the 19th International AAAI Conference on Weblogs and Social Media, Venice, Italy.

Subrahmanian, V. S., Azaria, A., Durst, S., Kagan, V., Galstyan, A., Lerman, K., et al. (2016). The DARPA Twitter Bot Challenge. *Computer, 49*(6), 38-46. doi: 10.1109/MC.2016.183

Sunstein, C. R. (2001). *Republic.com*: Princeton.

Swinford, S. (2016, August 10, 2016). Britain could be up to £70 billion worse off if it leaves the Single Market after Brexit, IFS warns. *The Telegraph*.

Temple Lang, D. (2016). RCurl: General Network (HTTP/FTP/...) Client Interface for R (Version 1.95-4.8). Retrieved from https://CRAN.R-project.org/package=RCurl

Varol, O., Ferrara, E., Davis, C. A., Menczer, F., & Flammini, A. (2017). *Online Human-Bot Interactions: Detection, Estimation, and Characterization.* Paper presented at the 11th International AAAI Conference on Weblogs and Social Media, Motreal, Canada.

Walker, S. (2015). *The Complexity of Collecting Social Media Data in Ephemeral Contexts*. Paper presented at the Internet Research 16, Phoenix, AZ.

Weedon, J., Nuland, W., & Stamos, A. (2017). Information Operations and Facebook (F. Security, Trans.): Facebook.

Woolley, S. C. (2016). Automating power: Social bot interference in global politics. *First Monday, 21*(4).

Woolley, S. C., & Howard, P. N. (2016). *Automation, Algorithms, and Politics| Political Communication, Computational Propaganda, and Autonomous Agents — Introduction* (Vol. 10).

Zheng, X., Lai, Y. M., Chow, K. P., Hui, L. C. K., & Yiu, S. M. (2011, 14-16 Oct. 2011). *Sockpuppet Detection in Online Discussion Forums.* Paper presented at the 7th International Conference on Intelligent Information Hiding and Multimedia Signal Processing.

Zimmer, M. (2010). "But the data is already public": on the ethics of research in Facebook. *Ethics and information technology, 12*(4), 313-325.

Table 1: Weblinks tweeted by deleted accounts

1	Dead link (external)	54.30%	**11**	facebook.com	0.50%
2	Valid link (Twitter)	17.20%	**12**	dailymail.co.uk	0.50%
3	Dead link (Twitter)	8.60%	**13**	twimg.com	0.50%
4	express.co.uk	1.70%	**14**	Suspended account	0.30%
5	theguardian.com	1.60%	**15**	cnn.com	0.20%
6	youtube.com	1.40%	**16**	petition.parliament.uk	0.20%
7	bbc.co.uk	1.10%	**17**	virgin.com	0.20%
8	@brndstr	1.00%	**18**	paper.li	0.20%
9	telegraph.co.uk	0.60%	**19**	thesun.co.uk	0.20%
10	bloomberg.com	0.50%	**20**	reuters.com	0.20%

Table 2: Metrics of automated activity for active users, recycled accounts, and bots. Predictors of bot activity shown in bold

	Active	Recycled	Botnet
users	1,641,472	26,538	13,493
tweets	6,546,998	103,606	63,797
tw2user	3.988492	3.904062	**4.72815**
tw2rtMean	0.451069	0.539010	**0.63021**
commonWords	0.138760	0.137341	**0.12963**
webClient	0.233227	0.241327	**0.34726**
mentionOut2In	1.204246	1.984704	**2.92922**
retweetIn2Out	1.035133	1.669028	**1.82125**
newAccount	0.428071	0.476256	**0.82619**
rtTransOut	0.005570	0.000876	0.00097
rtTransIn	0.005553	0.001292	0.00140
rtReciprocity	0.064109	0.001830	**0.00076**
modularScore	0.574058	0.666160	0.602234
cascadeMean	7.370980	8.005936	8.087834
cascadeMax	13417	1344	860
patZero	7.4%	7.2%	6.2%
patZeroShare	36%	34%	30%
ccdMeanTime	69h	102h	**65h**
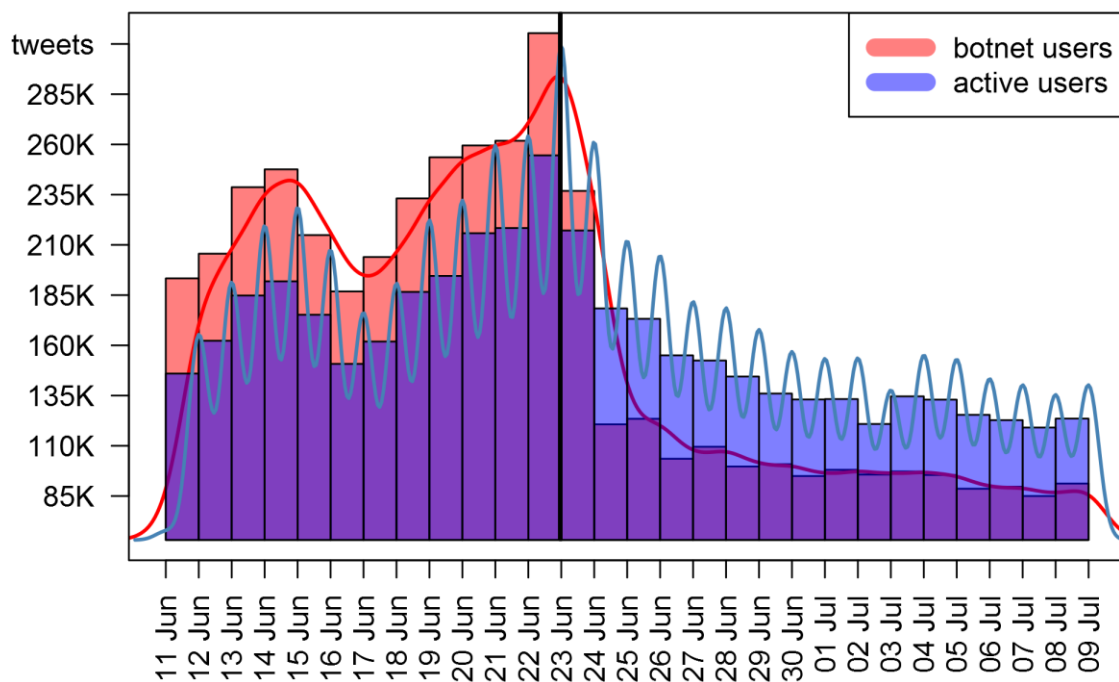
Figure 1: Twitter activity before and after the referendum. The vertical black line marks the date of the referendum (23 June 2016), with blue and red horizontal lines showing the density curves for active users and bots, respectively. The dark purple area in the histogram shows the relative proportion of user activity that overlaps with bot activity
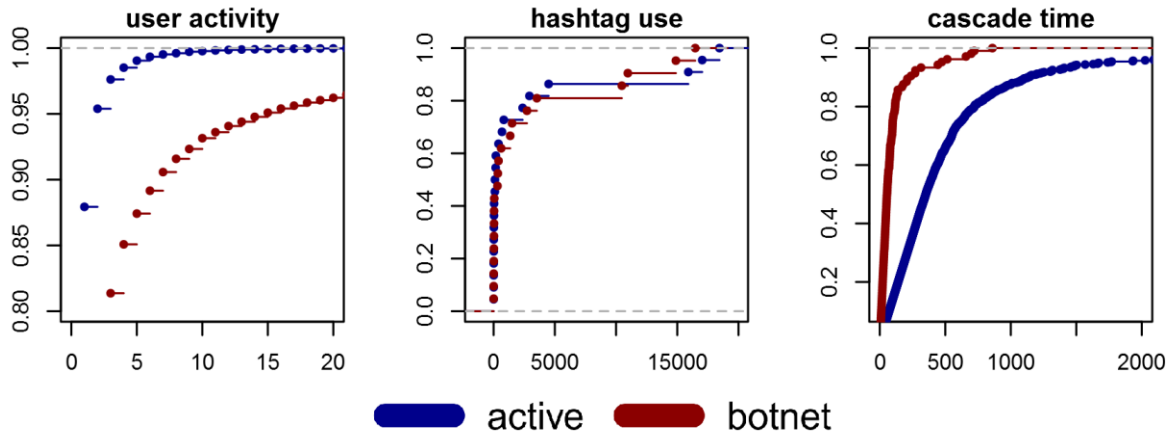
Figure 2: CDF of tweets, hashtags, and retweet cascades for active users and Twitterbots
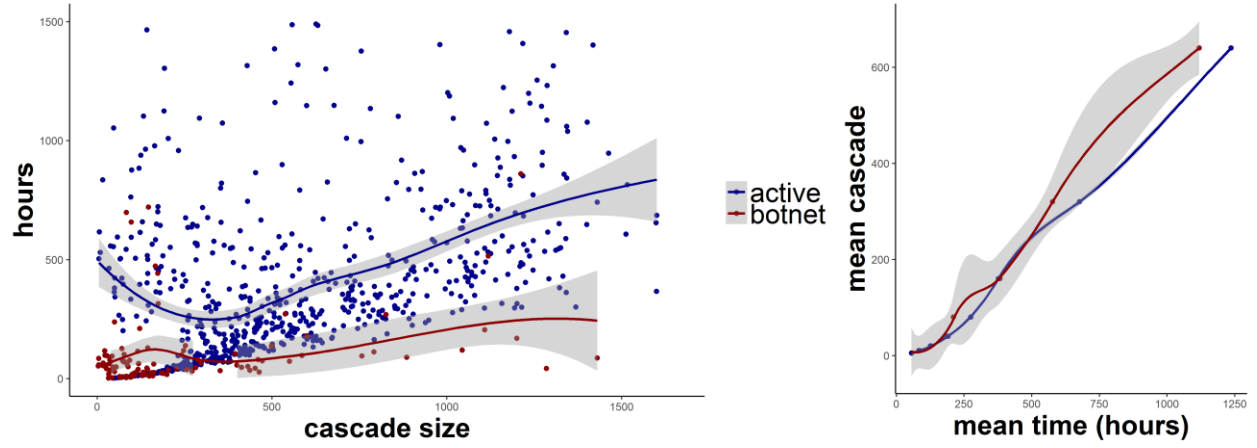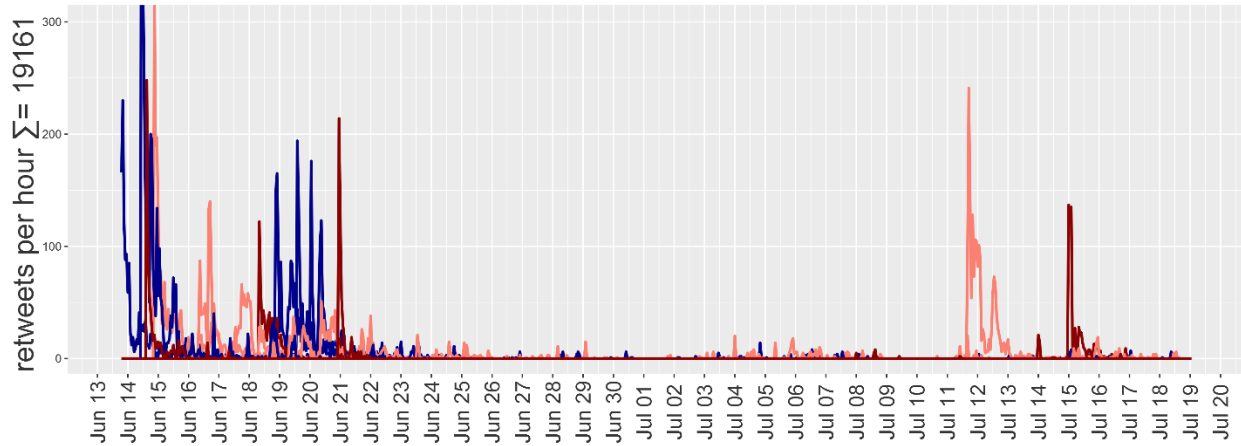
Figure 3: Time-to-cascade and mean cascade time for active users and Twitterbots

Figure 4: Two-tiered botnet, with bots specialized in retweeting active users and bots dedicated to retweeting other bots. Vertice and edge color identify source of information

Figure 5: Large cascades (*S*>506) from user to user, bot to bot, and user to bot