

Bots and Political Influence: A Sociotechnical Investigation of Social Network Capital

DHIRAJ MURTHY¹

University of Texas at Austin, USA

ALISON B. POWELL

London School of Economics, UK

RAMINE TINATI

University of Southampton, UK

NICK ANSTEAD

London School of Economics, UK

LESLIE CARR

SUSAN J. HALFORD

MARK WEAL

University of Southampton, UK

This study explains how bots interact with human users and influence conversational networks on Twitter. We analyze a high-stakes political environment, the UK general election of May 2015, asking human volunteers to tweet from purpose-made Twitter accounts—half of which had bots attached—during three events: the last Prime

Dhiraj Murthy: dhiraj.murthy@austin.utexas.edu

Alison B. Powell: a.powell@lse.ac.uk

Ramine Tinati: r.tinati@soton.ac.uk

Nick Anstead: n.m.anstead@lse.ac.uk

Leslie Carr: lac@ecs.soton.ac.uk

Susan J. Halford: Susan.Halford@soton.ac.uk

Mark Weal: mjw@ecs.soton.ac.uk

Date submitted: 2016-08-30

¹ The authors gratefully acknowledge the feedback from the day-long workshop Algorithms, Automation and Politics, organized by the European Research Council–funded Computational Propaganda project of the Oxford Internet Institute and held as a preconference to the International Communication Association Meeting in Fukuoka, Japan, in June 2016. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the European Research Council.

Copyright © 2016 (Dhiraj Murthy, Alison B. Powell, Ramine Tinati, Nick Anstead, Leslie Carr, Susan J. Halford, & Mark Weal). Licensed under the Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd). Available at <http://ijoc.org>.

Minister's Question Time before Parliament was dissolved (#PMQs), the first leadership interviews of the campaign (#BattleForNumber10), and the *BBC Question Time* broadcast of the same evening (#BBCQT). Based on previous work, our expectation was that our intervention would make a significant difference to the evolving network, but we found that the bots we used had very little effect on the conversation network at all. There are economic, social, and temporal factors that impact how a user of bots can influence political conversations. Future research needs to account for these forms of capital when assessing the impact of bots on political discussions.

Keywords: bots, political communication, capital, moral panics, experimental methods

In recent years, scholars have been fascinated with the role of algorithms and other automated processes in political life. Bot studies, and algorithm studies in general, have focused on the unknowable qualities of decision-making machines and the kinds of governance that these create (Barocas, Hood, & Ziewitz, 2013). In political communication, they may complicate the effectiveness of existing—and novel—means of understanding public opinion. Specifically, a proliferation of bots could challenge the robustness of using social media to understand large-scale population dynamics. Optimists see social media analysis as a way to understand the public in new and distinctive ways, which are more organic and require less researcher intervention than traditional methods (Anstead & O'Loughlin, 2014).

Other broader claims have been made about the possibility of social media enhancing democratic life more generally, providing new ways for citizens to engage and deliberate (Loader & Mercea, 2012). Additionally, social media monitoring is seen as having the potential to make government more interactive and better able to respond to the will of the public based on evidence. For example, the UK Cabinet Office (2015) argues, "More than 50% of people in the UK use social media every week, so it is an important source of information for policy makers. It can feed our policy understanding by building our knowledge and helping us to make informed decisions" (para. 3). Will public policy decisions be distorted by public opinion data corrupted by bots?

Pessimistic researchers question the value of such big data methods. For some, the data being gathered are not representative of the public overall (Gayo-Avello, 2013; Murthy, 2015). Others have highlighted that social media analysis reflects the structural inequalities of wider society. For example, political conversation on social media may be dominated by those with higher incomes and higher levels of education (Duggan & Brenner, 2013). Also, automated content analysis through machine learning faces challenges in parsing the nuances of human communication. Specifically, using natural language processing algorithms to effectively and accurately code many of the data sets being produced online may require "costly manual annotation" (Friedman, Rindfleisch, & Corn, 2013, p. 768). This line of thought, as Sells, Shieber, and Wasow (1991) explain, stretches back to Chomsky's critique "that finite-state grammars were inadequate models of natural language syntax" (p. 1). Thus, bots are potentially problematic, as they render these grammars even more fixed.

In this study, we evaluate the impact of bots as automated meaning makers that shift or twist political communications. The popular press regularly presents bots as distorting public conversation

(Finger, 2015). "Hackers also use bots for all sorts of nefarious reasons," writes Pullen (2015), "from lifting credit card numbers from an online store to scraping the text off an article and posting it on some random blog" (para. 4). Similar stories place bots at the center of geopolitical conflict. For instance: "Today, waves of trolls and bots regularly promote pro-Putin hashtags" (Chen, 2015, para. 26). These representations of bots are part of a larger moral panic about bots as malicious forces (Finger, 2015).

In this context, our project began with an assumption that an interdisciplinary team of political communication specialists, sociologists, science technology and society scholars, and computer scientists could experimentally model the impact of bots on a political conversation. A significant finding of our research is that the bots we built did not significantly influence the political conversation. Writing bot code in Python is relatively straightforward, and bots have become a significant population on diverse social media platforms (Fenton, 2016). Many social media companies promise to control bots, but we were under the impression that getting bots and deploying them would be relatively straightforward. However, we quickly realized that high amounts of capital—cultural, economic, temporal, and social—were needed to influence political discourse and that using bots in this way required some methodological sophistication that we had not anticipated. This finding is interesting in itself, as it indicates that those without significant financial resources and pre-existing bot knowledge are not able to readily influence political discourse. In other words, a combination of money, knowledge, time, and social networks is needed, raising questions about the uneven capacity of various actors to successfully leverage bot networks.

This project is tremendously important because work on political bots has been theoretical, and interdisciplinary empirical investigations remain lacking. Meanwhile, computational work on bots has focused more on automated methods for identifying bots and stripping them from the conversation than on tracing their impacts (Chu, Gianvecchio, Wang, & Jajodia, 2012; Davis, Varol, Ferrara, Flammini, & Menczer, 2016). Additionally, some experimental bot studies have ended up emulating humans and ultimately tricking users (Messias, Schmidt, Oliveira, & Benevenuto, 2013), methods that raise significant ethical questions. In contrast, we have taken an interdisciplinary sociotechnical approach to explore the processes involved in deploying bots and analyzing their effects.

Of course, lines of inquiry into bot detection, for example, are tremendously valuable. However, our experience offers insights into the challenges that researchers face when trying to organize bot-based experiments and raises some profound questions about how the power to mobilize bots to shape social networks is accumulated, if at all. First, we define bots and identify the assets required to create them. Then we describe how our research team set up the experiment and procured bots. Following this, we provide an analysis of the significance of our bots in the communication ecology. We draw on a Bourdieusian vocabulary to explain the performance of our bots in terms of capital. We conclude with some insights about automation within communication systems, bot studies, and social media methodology.

Building Political Bots

Defining a Bot

When we write about bots we refer specifically to social media bots—computer programs or algorithms controlling accounts on social media. Social bots are distinct from declared or spam bots, as they try to appear human (Boshmaf, Muslukhov, Beznosov, & Ripeanu, 2011). Some have claimed that social bots of this kind are created with the purpose of altering human behavior (Ferrara, 2015; Wald, Khoshgoftaar, Napolitano, & Sumner, 2013). These definitions are useful, and important empirical work on social bots has been conducted, but we still know little about bots' behavior, their relationship with human actors, and their ability to influence wider networks (Davis et al., 2016).

Bots, at their simplest, are social media accounts that are controlled either wholly or in part by software agents. When we talk about a bot, we often think of the account itself and the mechanisms controlling it as a single entity. However, to better understand the nature of bots, we must distinguish the account from the publishing algorithm. Social media systems operate in terms of accounts. These accounts can perform actions through interfaces such as application programming interfaces (APIs). Actions taken on APIs, whether instigated by people or algorithms, will often be largely opaque to the social media system itself.

We can see bots as a subcategory of algorithmic media elements because they are programmed to intervene in the way knowledge and information is communicated. The expanding "critical algorithms" literature (Gillespie & Seaver, 2016) focuses on the black-box nature of these automated processes and the significance they have for access to information and openness of speech. According to Gillespie, Boczkowski, and Foot (2014), our growing dependence on proprietary and invisible algorithms is "as momentous as having relied on credentialed experts, the scientific method, common sense, or the word of God" (p. 164). This quotation speaks to the significant power attributed to algorithms and other automated elements within the communication system and to the way that they become black boxes even by their critics and researchers.

On the Creation of Bots

Bots are added to social media systems in a variety of ways. A new account may be created with the explicit intention of having a bot control it. Depending on the social media system, this may even be possible automatically, but most social networking platforms try to ensure that only people can create new accounts. Existing user accounts can have bots attached to them when a human account holder—knowingly or not—passes over some of the control of the account to a software agent. An example of this might be a bot designed to repost tweets at timed intervals to increase the visibility of a person's posts. Alternatively, bots may capture a dormant account. In this case, existing accounts that the original creators are no longer actively using can be captured by hacking passwords and sold as supplementary followers. Bots may also be added as followers to accounts that the purchaser does not control.

On the Functionality of Bots

Once bots have partial or total control of a social media account, the design of the bot structures how it engages with the social network. If a user buys followers, the act of following another account might be the sole function of the bot. Once following an account, the bot might retweet or favorite tweets of the followed account. Alternatively, a bot might be designed to inject new content into the social media network, to circulate promotional material or spam, or to steer and disrupt discussion online (Woolley, 2016). Experimental bot studies on Brazilian news and politics have highlighted that bots can quickly gain more measurable influence than established, prominent users in specific domain areas (Messias et al., 2013). During the 2012 Mexican presidential elections, bots were used to disrupt and falsify political communication efforts (Bacallao-Pino, 2016). In the United States they have been used to spam via “astroturf” methods (Ratkiewicz et al., 2011). However, bots can also be used positively to assist in the mobilization of networks of supporters for a cause, emulating the techniques of well-known face-to-face call-to-action techniques (Savage, Monroy-Hernandez, & Hollerer, 2015).

Bot accounts may be individual computational actors but may also be the individual voices of a coordinated network of bot entities—a “botnet.” A study of 35 weeks of activity of a specific botnet controlling 130 individual bots tweeting messages relating to the Syrian civil war shows that it only achieved significant influence in the final third of its lifespan (Abokhodair, Yoo, & McDonald, 2015). This happened when the number of both bot accounts and bot-driven tweets were at a zenith. At this time the messages were produced at a rate far exceeding anything done by humans. By that point, the network had managed to influence real human accounts into retweeting its messages, particularly of opinion topics in which humans are more active.

On the Detection of Bots

Significant research from the computer science community has been directed at the automatic detection of bots as algorithms become more complex and adept at simulating the behavior of a human controller. A DARPA challenge has even been mounted in this area (Emerging Technology from the arXiv, 2016). A variety of techniques have been used to identify bots, with measurable success in detecting the strict timing of bot postings (Chu et al., 2010; Costa, Yamaguchi, Traina, Caetano Traina, & Faloutsos, 2015). However, timing is under computational control and can be disguised if bot designers seek to avoid detection. The BotOrNot service uses machine learning techniques to judge whether an account is controlled by a bot based on six feature dimensions: network, user, friends, time, content, and sentiment. This method of combining multiple approaches can yield comparatively low false positive rates (Ferrara, Varol, Davis, Menczer, & Flammini, 2014).

Method and Research Questions

In our experimental study, student volunteer participants were asked to set up new Twitter accounts and comment on high-profile broadcast events using the hashtags linked to those events. Bots were linked to the accounts of some of the participants while other accounts remain untouched. We

then monitored the activity of the participants, the bots, and the overall pattern of the conversational network. This method allowed us to address several research questions:

- RQ1: Do accounts with bots attached behave differently from those that are untouched?*
- RQ2: Do accounts with bots attached have greater social influence within conversations than those without?*
- RQ3: Do different types of bots exert different types of influence? Are particular types of bots more effective?*
- RQ4: How do bots interact differently with different types of participants?*

Experiment Setup: Research Design

Our experiment attempts to better understand the role of bots within a series of related events on the microblogging platform Twitter. In this experiment, we defined bots as a set of automated processes—attached to half of the participants' Twitter accounts. These automated processes simulated a number of system-level functions within the Twitter platform, including the features retweeting, following, and favoriting. There are several ways to deploy bots, all of which we evaluated in light of ethical considerations. The first method would be to fully create our own bots, the method used by Messiah et al. (2013). The clear advantages to this are full design control of each bot. However, as with the controversial Tay—the renegade Microsoft chatbot—it can be very difficult to envision what a bot will do in the wild. Second, in our study, we wanted to experiment with the supposed ease of access to bots portrayed in the popular press, and we sought to create bots using straightforward and ethical methods.

Another alternative for deploying bots is to approach advanced hacker communities. Though the bots deployed by these groups leverage some of the most influential botnets over hijacked accounts, this option contravenes both acceptable research ethics protocols and Twitter's terms of service. The other method is to purchase bots legally from a digital marketing firm that has a clear framework for how they deploy their bots and can provide examples of past projects. This relative transparency ensures certain constraints on bots when releases in the wild. We opted for this last method as the means to deploy our bots both to emulate ease of use and to meet institutional ethics guidelines. We used the MonsterSocial software to achieve this. It provides a user interface to manage and configure a number of Twitter accounts, with options such as automatic retweeting, following, and favoriting.

Our experiment was focused around three events occurring within the United Kingdom, defined by three hashtags. These events included the political discussions within the BBC television show *BBC Question Time* (#BBCQT), the television airing of Prime Minister's Questions (#PMQ), and the television debate "Battle for Number 10" (#BattleForNumber10).

The experiment recruited 12 volunteer participants to write tweets and participate in the discussions relevant to the selected set of hashtags. Participants were each asked to create a new Twitter account with a name of their choosing; we asked participants not to use their real names. Participants

contributed to all three events, and the bot configuration remained the same during all phases of the experiment. Using the MonsterSocial software, we attached bots to half of the participants' Twitter accounts, which were then configured to automatically follow and retweet accounts within a given set of hashtags. The follow process ran every five minutes and tried to follow up to 100 random Twitter accounts that tweeted with a given hashtag. Similarly, the retweet process ran every five minutes and retweeted tweets that contained the chosen hashtag.

To capture the Twitter data during the three events, we used the EPrints tweet harvesting service Tweepository.² The service is designed to collect Twitter data based on a given keyword, including hashtags. It uses the Twitter search API in combination with its own harvesting strategy to collect publically posted Tweets and also captures a snapshot of the follower network of a set of specified accounts at similar intervals. This is important in revealing how the network changes during the course of the experiment. Using Tweepository, we created entries corresponding to each of the three event hashtags. Table 1 provides an overview of three data sets collected. Finally, we exported the data in a JSON structure and transformed the data into a series of temporal network graphs (Tinati, Carr, Hall, & Bentwood, 2012) constructed from the mention and retweets extracted from the collected tweets.

Ethics

To ensure that our study was ethically sound, we employed the London School of Economics' (LSE, 2016) research ethics self-certification form. This process requires reflexivity on the part of researchers to identify and mitigate any risks in a proposed study. During this process, we identified three ethical issues. The first related to the types of bots we were going to use for the experiment. A variety of bots are commercially available. However, at least some of these bots use accounts that have been hacked and hijacked from human users. To use bots of this kind would certainly have been unethical, illegal, and against Twitter's terms and conditions. Instead, we used bots from MonsterSocial that provide automatic following mechanisms but require humans to perform tweets, unlike the fully automated bots used in other studies (Messias et al., 2013).

Second, because we used human volunteers as part of our experiment, we needed to ensure their safety and well-being. As other scholars have noted (Burnap & Williams, 2015; Murthy & Sharma, 2015), online political conversations can sometimes create aggressive spaces, with conversations often degenerating into hate speech of various kinds. To protect our participants from such eventualities, we provided them with our contact details and made clear they could cease participating in the experiment at any time. In addition, we provided them with contact details of appropriate support services at the LSE.

Our final concern related to the content that participants were posting, particularly the possibility that they might publish anything that had the potential to create risks for themselves or our institutions. To prevent this from occurring, the Twitter accounts of participants were monitored at all times during the

² See Hitchcock (2013) for some applications of Tweepository and other harvesting services.

experiment by a member of the research team, who also—if it would have been necessary—had the ability to shut down any participant’s Twitter account.

Results

Data Set Overview

Table 1. Crossover of Unique Twitter Users Between #BBCQT, #PMQ, and #BattleForNumber10.

	BBCQT	PMQ	BattleForNumber10
BBCQT	...	6,371 (46,408)	8,248 (67,444)
PMQ	8,734 (71,773)
BattleForNumber10

Note. A crossover user is identified by the subset of users that exists between two datasets. Parentheses indicate the total numbers of unique users between both data sets.

Table 2. Descriptive Statistics for Twitter Data Sets #BBCQT, #PMQ, and #BattleForNumber10.

Metric	BBCQT	PMQ	BattleForNumber10
Collection start	03/13/2015 02:54	03/11/2015 12:16	03/21/2015 15:03
Collection end	03/30/2015 20:00	03/30/2015 19:58	03/2015 20:02
Tweets	63,309	69,171	99,200
Unique Users	23,980	28,798	51,710
Retweets	42,004	45,553	70,749
Mentions	4,860	6,186	5,852
Cascades	1,615	1,612	1,131
Longest cascade	231	185	290

Note. All values represent data prior to any preprocessing or data cleansing.

As Table 2 illustrates, the networks show similar proportions of retweets and of unique users to total tweets and are similar in the numbers and sizes of the cascades that emerged (the number of times that a single tweet was retweeted). With regards to the three different datasets, interestingly, we do see crossover between the different users who participated, showing that there was some consistency in terms of users within the network.

Bot Participation

Table 3. Twitter Account Statistics for Bot Accounts Used During the Three Experiments.

Twitter User	Following	Followers	BBCQT (T/RT/M)	PMQ (T/RT/M)	BattleForNumber10 (T/RT/M)
SoniaPritchett	0	1	13/0/0	7/0/0	11/0/0
n_hart45	0	2	6/0/0	9/0/0	16/0/0
MissCat7327	0	1	17/0/0	0/0/0	6/0/0
tomm_sanders	0	0	23/0/0	4/0/0	16/0/0
JohnsonZhang520	71	12	10/0/0	0/0/0	10/0/0
HamishMiller95	97	3	5/0/0	2/0/0	12/0/0
Politizy	42	8	21/0/0	19/0/0	70/0/0
Koalaparkpark	28	1	8/0/0	15/0/0	26/0/0
Toriga	71	2	13/0/0	0/0/0	17/0/0

Note. Purple rows are accounts without bots, and green rows are accounts with bots. T = tweet; RT = retweet; M = mention. Total tweet count may be higher, as all tweets were captured using search API in real time. This is the most accurate way to collect topic-specific tweets compared to the streaming API, but some tweets might not be captured.

Table 3 illustrates—independent of whether an account had a bot (bot accounts are those in purple) attached or tweeting in the conversations—that these accounts were not able to gain any retweets or mentions during the hashtag conversations. The users of these accounts were not told if they were going to have bots or not. As the results show, accounts that had bots attached gained slightly more followers than those without bots because the bots were following Twitter users during the tweeting sessions, making it more likely that the account was also followed.

Network Analysis

Next, we conducted a network analysis of mentions. Figure 1 reveals that only two of our accounts were in the edge list and not connected to anyone with 15 mentions or accounts mentioning. This indicates that our accounts were quite distant from anyone mentioned or doing mentioning: They were not proximate to any influencers. This suggests that bots are not necessarily generic and that the location of a bot in the network likely matters significantly. These findings are consistent with other studies examining the network position of bots (Ferrara et al., 2014). Additionally, bots with lots of followers using these hashtags would likely have a much higher impact on influencing the conversations we studied. In other words, bots could have domain expertise and impact based on their network positions.




Figure 1. Network visualization of our edge list.
Source: Authors' visualization using data analyzed.

These findings fit within existing literature studying bots that employs social network analysis. For example, Varvello and Voelker (2010) in their work on bot behavior and bot detection argue that network centrality measures are important to understanding the role of bots in online social networks.

Information Diffusion

We arranged our data to produce temporal visualizations to specifically study information diffusion. Figure 2 provides a temporal visualization of the speed (in minutes) at which a tweet was shared via the retweet functionality. The graph contains all tweets that were retweeted two or more times; each line represents a unique tweet. The steeper the line, the faster the information was spread, and the height of the line (y axis) indicates the number of times it was retweeted. As highlighted in Figures 2, 3, and 4, several repetitive diffusion characteristics can be found within the three information diffusion networks: (a) fast and abrupt diffusion of information shared by few actors; (b) slow diffusion occurring over a long period of time, shared by few actors; and (c) fast diffusion of information, shared by many actors.

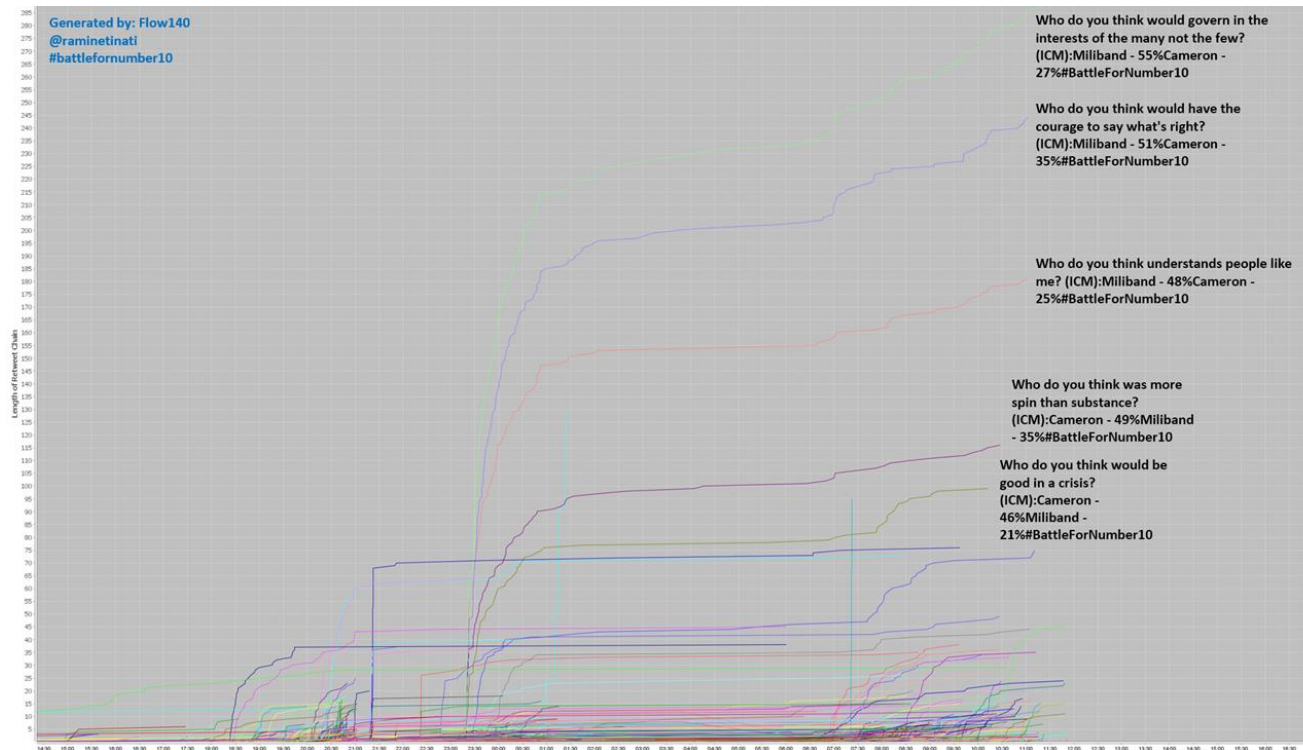


Figure 2. Information diffusion of #BattleForNumber10.

We pay particular attention to the information cascade profiles where information diffuses quickly and is shared by a large number of actors. As Figure 2 illustrates, during the #battlefornumber10 Twitter conversation, the most prevalent tweets shared across the network contained information about the polling status between two political candidates. We see that the five most retweeted tweets share the

same cascade structure, with virtually identical content, bar the change in polling results. If we compare this to the #BBCQT and #PMQ hashtags, as seen in Figures 3 and 4, we find similar profiles of fast diffusing information across many actors, yet the content is far broader, ranging from personal views on political issues to the reporting of news and comments made during the television broadcast.

This is in keeping with previous research in two strands of political science. First, work on public reactions to broadcast events on social media indicates that viewer-commenters' reactions will be conditioned by and reflective of what they are viewing (Anstead & O'Loughlin, 2011; Elmer, 2012; Freelon & Karpf, 2015). As such, it is not surprising that different broadcast types will see different types of comments dominating. Second, #thebattlefornumber10 was not a regular broadcast, as it was at least something of an election debate. In contrast, #BBCQT and #PMQS are both regularly scheduled parts of the political week. This means not only that the latter two broadcast events had well-established networks of commentators using social media to talk about them but also that #battlefornumber10 was viewed as an election debate, with reactions inevitably being much about winners, losers, and electoral impact, especially in a multimedia environment (Chadwick, 2010). It also raises interesting issues related to the temporal nature of social media conversations: Because our volunteer tweeters only participated in conversations during the specific time periods covered by the programs, they potentially lacked the network influence that they might have accrued as regular tweeters during PMQs or *BBC Question Time*.




Figure 3. Information diffusion of #BBCQT.




Figure 4. Information diffusion of #PMQ.

Hashtag Analysis

Unfortunately, the use of hashtags is not consistent on Twitter. The hashtags #bbcqt, ##bbcqt, and #bcqt are all used for *BBC Question Time*. Similarly, #PMQ, #pqs, and #pmqs are used for Prime Minister's Question Time. This is problematic for hashtag analysis but is consistent with other hashtag-based Twitter work that highlights potential data loss due to misspelled or varying hashtags (Weller, Dröge, & Puschmann, 2011). This raises questions about methods of collecting data. Specifically, if hashtags are used in complex ways, understanding the actors involved in a conversation is not as simple as using one hashtag to track the conversation. This is an important point given how much social media-based research of politics uses hashtags as selection criteria (Jungherr, 2014).

For instance, #battlefornumber10 was quite clearly a dominant hashtag, but #bbcqt had several variants in use, including ##bbcqt and #bcqt. If we had access to the full Twitter Firehose, the data stream with all tweet data, we hypothesize that we might have detected more variants. Several data collection issues can already affect the way we interpret conversations and interactions.

Discussion: Capital Required to Influence Social Networks with Bots

Overall, our bots were far less successful than we expected. In retrospect, we interpret this in terms of the capital required to deploy bots effectively and suggest that social influence, even over technologies that allow bots, is a product of capital.

First, it became apparent that the new accounts we set up lacked the social capital to be effective. An account may have social capital in its network of followers: The more followers someone has, the more social capital they have. Indeed, people buy bot followers to appear as though they have extended networks, more influence, or more social capital. However, it became clear in our experiment that whether bots impact the spread of information or the shape of social networks depends on the type of bot deployed and the type of account it is attached to. The bots that we used were attached to our participants' new accounts and operated by retweeting our participants' posts. They repeated the messages that our students posted, ensuring that they appeared more than once on their followers' timelines, making it more likely that followers would notice a post. If a Twitter user is following 500 users, then the chances that he or she will see any one post from one of those followed accounts may be quite slim. Repetition of the post increases the likelihood that it will be seen. In this case, the effectiveness of the bot rests on the existing social capital of the account to which it is attached. It is possible that by repeating messages from that account it will attract more attention for the user, maybe even further retweets or new followers, but the original driver is the number of followers that the account had in the first place.

To be more effective, we could have bought a different kind of bot. For example, other bots work by hacking into dormant accounts: genuine human user accounts that have been set up but abandoned. These accounts could then have been set to follow our students' accounts and to retweet all the students' posts to the existing network of the dormant account. In this case, then, the social capital of the original dormant account is transferred to the bot and mobilized on behalf of the bot purchaser. As novices, we had not realized how significant these differences were. The bots that hijack accounts are often illegal, certainly violate Twitter's terms and conditions, and would never have been approved by our university ethics committee. These types of bots have been used in election contexts where the bots clearly had success in influencing the political landscape. They are similar to those used in the 2012 Mexican general election (Orcutt, 2012) and those that silenced anti-Kremlin dissent in Russia (BBC Technology News, 2012).

In these contexts, we might reflect on the capital already possessed by the entities effectively using bots. In our first experiment in this area, we lacked the cultural capital and the technical capital to invest in bots that would have been more effective (Halford & Savage, 2010). These more effective bots are also more expensive, raising the question of economic capital. Even if we had the technical knowledge to choose the right bots, we were not in a position to spend a lot of money on them. It is also possible that Twitter as a service would intervene in the process. We do not know a great deal about how Twitter organizes content, but it is feasible that the company moderates follower timelines, biasing these toward tweets that are already popular. Previous studies have highlighted that the Twitter API has various biases, some of which are challenging to discern (Joseph, Landwehr, & Carley, 2014).

Overall, we might say that our low cultural and technical capital—understanding which bots to buy—combined with the low social capital of the account holders caused our bots to fail to have influence. These were fresh Twitter accounts created for the experiment, so they had few friends, followers, and social interactions initially and therefore not could not convert the cash we paid for the bots into social capital (enhanced networks) on Twitter. Clearly, the effect of bots is shaped by a complex interplay of social and technical factors mobilized in interaction with each other.

Conclusion

This was a very small experiment, and it may be that with more participants or more sophisticated bots, we would have generated more influence. However, our findings also point to a different definition of political bots. Traditional definitions assume bots are social because they use interaction to convince people they are human. However, they could equally be seen as social because they are a product of the social, political, and economic context that created them.

Clearly, bots do have influence in the political context. The cases of Mexico (Orcutt, 2012), Venezuela (Forelle, Howard, Monroy-Hernández, & Savage, 2015), Turkey (Saka, 2014), and Russia (BBC Technology News, 2012) highlight this capacity. In many of these cases, hijacked accounts and hacker botnets have been involved. If the bots most successful in shifting conversations are those that are constructed from hijacked accounts, how can researchers effectively study the influence of bots? We were very mindful of the fact that our group of interested social media scholars found it impossible to assemble the appropriate financial, (un)ethical, social, and platform-based resources that might make it possible to influence a political conversation.

Does this mean that moral panic about bots is unfounded? Or—more worryingly—does it mean that these concerns are well-founded, but that the kinds of entities shifting online conversations have resources and understanding at their disposal to make them even more difficult for critical researchers to investigate?

Our experiment also demonstrates that our student volunteers lacked social capital within the conversational networks at hand. They also lacked the temporal capital they could have accumulated had they regularly tweeted using the same hashtags. These may be reasons why they—and the bots attached to their accounts—were unable to exert much influence. We can extend this logic to real-world examples. For instance, pro-Russian bot activity online can be seen as a distortion of online conversations and networks. At the same time, this is all also a manifestation of real geopolitical circumstances. How do bots really work in social media? Our experiment demonstrates that they may have negligible impact, and that the most effective bots may be the ones we cannot study.

References

- Abokhodair, N., Yoo, D., & McDonald, D. W. (2015). Dissecting a social botnet: Growth, content, and influence in Twitter. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work* (pp. 839–851). Vancouver, Canada: ACM.
- Anstead, N., & O'Loughlin, B. (2011). The emerging viewertariat and *BBC Question Time*: Television debate and real-time commenting online. *The International Journal of Press and Politics*. doi:1940161211415519.
- Anstead, N., & O'Loughlin, B. (2014). New directions in Web analysis: Can semantic polling dissolve the myth of two traditions of public opinion research? In D. Hansen, B. Schneiderman, & M. A. Smith (Eds.), *Analyzing social media data and Web networks* (pp. 256–275). New York, NY: Springer.
- Bacallao-Pino, L. M. (2016). Radical political communication and social media: The case of the Mexican #yosoy132. In T. Dezelan & I. Vobic (Eds.), *(R)evolutionizing political communications through social media* (pp. 56–74). Hershey, PA: IGI Global.
- Barocas, S., Hood, S., & Ziewitz, M. (2013). *Governing algorithms: A provocation piece*. Social Science Research Network. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2245322
- BBC Technology News. (2012, March 8). Russian Twitter political protests swamped by spam. *BBC News*. Retrieved from <http://www.bbc.co.uk/news/technology-16108876>
- Boshmaf, Y., Muslukhov, I., Beznosov, K., & Ripeanu, M. (2011). The socialbot network: When bots socialize for fame and money. In *ACSAC '11: Proceedings of the 27th Annual Computer Security Applications Conference* (pp. 93–102). Orlando, FL: ACM. doi:10.1145/2076732.2076746
- Burnap, P., & Williams, M. L. (2015). Cyber hate speech on Twitter: An application of machine classification and statistical modeling for policy and decision making. *Policy & Internet*, 7(2), 223–242.
- Chadwick, A. (2010). Britain's first live televised party leaders' debate: From the news cycle to the political information cycle. *Parliamentary Affairs*, 64, 24–44.
- Chen, A. (2015, June 2). The agency. *The New York Times Magazine*. Retrieved from www.nytimes.com/2015/06/07/magazine/the-agency.html
- Chu, Z., Gianvecchio, S., Wang, S., & Jajodia, S. (2010). Who is tweeting on Twitter: human, bot, or cyborg? In *Proceedings of the 26th Annual Computer Security Applications Conference* (pp. 21–30). Austin, TX: ACM. doi:10.1145/1920261.1920265

- 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.
- Costa, A. F., Yamaguchi, Y., Traina, A. J. M., Caetano Traina, J., & Faloutsos, C. (2015). RSC: Mining and modeling temporal activity in social media. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 269–278). Sydney, Australia: ACM.
- Davis, C., Varol, E., Ferrara, E., Flammini, A., & Menczer, F. (2016). BotOrNot: A system to evaluate social bots. In *Proceedings of the 25th International Conference Companion on World Wide Web* (pp. 273–274). Geneva, Switzerland: International World Wide Web Conferences Steering Committee. doi:10.1145/2872518.2889302
- Duggan, M., & Brenner, J. (2013, February 14). The demographics of social media users. *Pewinternet.org*. Retrieved from <http://www.pewinternet.org/2013/02/14/the-demographics-of-social-media-users-2012/>
- Elmer, G. (2012). Live research: Twittering an election debate. *New Media & Society*. doi:1461444812457328.
- Emerging Technology from the arXiv. (2016). How DARPA took on the Twitter bot menace with one hand behind its back. *MIT Technology Review*. Retrieved from <https://www.technologyreview.com/s/546256/how-darpa-took-on-the-twitter-bot-menace-with-one-hand-behind-its-back/>
- Fenton, N. (2016). The Internet of me (and my “friends”). In J. Curran, N. Fenton, & D. Freedman (Eds.), *Misunderstanding the Internet* (pp. 145–172). New York, NY: Routledge.
- Ferrara, E. (2015). Manipulation and abuse on social media. *arXiv*. Retrieved from <https://arxiv.org/abs/1503.03752>
- Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. (2014). The rise of social bots. *arXiv*. Retrieved from <http://arxiv.org/abs/1407.5225>
- Finger, L. (2015, February 17). Do evil: The business of social media bots. *Forbes*. Retrieved from <http://www.forbes.com/forbes/welcome/?toURL=http://www.forbes.com/sites/lutfinger/2015/02/17/do-evil-the-business-of-social-media-bots/&refURL=https://www.google.es/&referrer=https://www.google.es/>
- Forelle, M., Howard, P., Monroy-Hernández, A., & Savage, S. (2015). Political bots and the manipulation of public opinion in Venezuela. *arXiv*. Retrieved from <http://arxiv.org/abs/1507.07109>
- Freelon, D., & Karpf, D. (2015). Of big birds and bayonets: Hybrid Twitter interactivity in the 2012 presidential debates. *Information, Communication & Society*, 18(4), 390–406.

- Friedman, C., Rindflesch, T. C., & Corn, M. (2013). Natural language processing: State of the art and prospects for significant progress. *Journal of biomedical informatics*, 46(5), 765–773.
- Gayo-Avello, D. (2013). A meta-analysis of state-of-the-art electoral prediction from Twitter data. *Social Science Computer Review*, 1(6), 647–679. doi:0894439313493979.
- Gillespie, T., Boczkowski, P. J., & Foot, K.A. (Eds.). (2014). *Media technologies: Essays on communication, materiality, and society*. Cambridge, MA: MIT Press.
- Gillespie, T., & Seaver, N. (2016, August 30). Critical Algorithm Studies: A reading list. *Socialmediacollective.org*. Retrieved from <https://socialmediacollective.org/reading-lists/critical-algorithm-studies/>
- Halford, S., & Savage, M. (2010). Reconceptualizing digital social inequality. *Information, Communication & Society*, 13(7), 937–955.
- Hitchcock, S. (2013). *Collecting and archiving tweets: A datapool case study*. Southampton, UK: University of Southampton.
- Joseph, K., Landwehr, P. M., & Carley, K. M. (2014). Two one per cents don't make a whole: Comparing simultaneous samples from Twitter's streaming API. In *Social Computing, Behavioral-Cultural Modeling, and Prediction: Proceedings of SBP 2014, the 7th International Conference*, (pp. 75–83). Washington, DC: Springer International.
- Jungherr, A. (2014). *Twitter in politics: A comprehensive literature review*. Social Science Research Network. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2402443
- Loader, B. D., & Mercea, D. (2012). *Social media and democracy: Innovations in participatory politics*. New York, NY: Routledge.
- London School of Economics. (2016). Research ethics policy, procedures and guidance. London School of Economics. Retrieved from <http://www.lse.ac.uk/intranet/researchAndDevelopment/researchDivision/policyAndEthics/ethicsGuidanceAndForms.aspx>
- Messias, J., Schmidt, L., Oliveira, R., & Benevenuto, F. (2013). You followed my bot! Transforming robots into influential users on Twitter. *First Monday*, 18(7). Retrieved from <http://firstmonday.org/ojs/index.php/fm/article/view/4217/3700#p4>
- Murthy, D. (2015). Twitter and elections: Are tweets predictive, reactive, or a form of buzz? *Information, Communication & Society*, 18(7), 816–831.

- Murthy, D., & Sharma, S. (2015). Mapping race and YouTube comments. In *the XXXV Sunbelt Conference of the International Network for Social Network Analysis* (n.p.). Brighton, UK: Sunlight Media.
- Orcutt, M. (2012). Twitter mischief plagues Mexico's election. *MIT Technology Review*. Retrieved from <https://www.technologyreview.com/s/428286/twitter-mischief-plagues-mexicos-election/>
- Pullen, J. P. (2015, May 28). How bad bots are destroying the Internet. *TIME*. Retrieved from <http://time.com/3899612/internet-bots/>
- Ratkiewicz, J., Conover, M. D., Meiss, M., Goncalves, B., Flammini, A., & Menczer, F. M. (2011). Detecting and tracking political abuse in social media. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media* (pp. 297–304). Menlo Park, CA: AAAI Press.
- Saka, E. (2014). The AK Party's social media strategy: Controlling the uncontrollable. *Turkish Review*, 4(4), 418.
- Savage, S., Monroy-Hernandez, A., & Hollerer, T. (2015). Botivist: Calling volunteers to action using online bots. *arXiv*. Retrieved from <http://arxiv.org/abs/1509.06026>
- Sells, P., Shieber, S. M., & Wasow, T. (1991). *Foundational issues in natural language processing: Introduction*. Cambridge, MA: MIT Press.
- Tinati, R., Carr, L., Hall, W., & Bentwood, J. (2012). Identifying communicator roles in Twitter. In *the Proceedings of the 21st international conference companion on World Wide Web* (pp. 1161–1168). New York, NY: ACM.
- UK Cabinet Office. (2015, August 4). Open policy making toolkit: Social media listening and analysis. *Gov.uk*. Retrieved from <https://www.gov.uk/guidance/open-policy-making-manual>
- Varvello, M., & Voelker, G. M. (2010). Second life: A social network of humans and bots. In *The Proceedings of the 20th International Workshop on Network and Operating Systems Support for Digital Audio and Video* (pp. 9–14). Amsterdam, Netherlands: ACM.
- Wald, R., Khoshgoftaar, T. M., Napolitano, A., & Sumner, C. (2013). Predicting susceptibility to social bots on Twitter. In *the IEEE 14th International Conference on Information Reuse and Integration* (n.p.). San Francisco, CA: IEEE.
- Weller, K., Dröge, E., & Puschmann, C. (2011). Citation analysis in Twitter: Approaches for defining and measuring information flows within tweets during scientific conferences. In *#MSM2011: The First Conference on Making Sense of Microposts* (pp. 1–12). Heraklion, Crete: WWW.
- Woolley, S. C. (2016). Automating power: Social bot interference in global politics. *First Monday*, 21(4). doi:10.5210/fm.v21i4.6161