

Troll Factories: The Internet Research Agency and State-Sponsored Agenda Building

Darren L. Linvill

Department of Communication, Clemson University

darrenl@clemson.edu

Patrick L. Warren

John E. Walker Department of Economics, Clemson University

pwarren@clemson.edu

Thanks to:

[The Social Media Listen Center](#), Clemson University

Brandon C. Boatwright, University of Tennessee Knoxville

Abstract

We document methods employed by Russia's Internet Research Agency to influence the political agenda of the United States from June 19, 2015 to December 31, 2017. Using the Social Studio social media listening platform, we downloaded tweets from handles identified by Twitter as being associated with the Agency. We identified five handle categories: *right troll*, *left troll*, *news feed*, *hashtag gamer*, and *fearmonger*. Within each type, accounts were used consistently, but the behavior across types was radically different, both in terms of "normal" daily behavior and in how they responded to external events. In this sense, the Internet Research Agency's agenda-building effort was "industrial"-- mass produced from a system of interchangeable parts, where each class of part fulfilled a specialized function.

In February 2018, the U.S. Justice Department indicted 13 Russian nationals for interference with the 2016 U.S. Presidential election (Barrett, Horwitz, & Helderman, 2018). The indictment named the Internet Research Agency (IRA), based in St. Petersburg, as central to a Russian effort, beginning in 2014, to sow discord in the U.S. political system, largely through social media. Evidence of IRA-linked accounts was found on Facebook, Twitter, Instagram, and Google. These accounts masqueraded as American citizens and were used to try to divide voters along a range of issues.

Cobb and Elder (1971) define *agenda building* as the process by which actors endeavor to move issues from their own agenda onto the agendas of policymakers. In many analyses of agenda-building, media content is used to measure the existing agenda, influenced by varying constituencies (for review, see Denham, 2010). But scholars have also examined the role that media, including social media, play in driving an agenda (Lariscy, Avery, Sweester, & Howes, 2009, Parmelee, 2014). National efforts to use media to influence foreign citizens is not new; Japan broadcasted to U.S. troops throughout World War II, and Voice of America has for decades been an global mouthpiece of the U.S. government. However, Russia's work on social media has taken agenda-building efforts by nations into a new context. The purpose of this study is not to look at the agendas behind IRA efforts, but rather to better understand the structure of the IRA's agenda building campaign. Given the covert nature of this campaign, such understanding is essential.

This study asks two questions about the IRA's behavior on the social media platform Twitter:

RQ1: Can IRA Twitter handles employed between June 19, 2015 and December 31, 2017 be categorized by their content into discrete types, and if so, what characterizes those types?

RQ2: If so, are differing types of Twitter handles employed by the IRA in ways that are different from one another?

Method

We employed an exploratory, sequential mixed methods design (Creswell, 2014), first applying qualitative analysis to the data and then using quantitative analysis to explore how behavior varied over time. We used Salesforce's Social Studio social listening platform to download all tweets in the study period from known IRA accounts (for API information see: Introducing Social Studio REST API, n.d.). Variants of this software have been used to examine social media content in political dialogue (Edgerly, Thorson, Bighash, & Hannah, 2016), crisis events (Black, Dietz, Stirratt, & Coster, 2015), and educational settings (Linville, Boatwright, & Grant, 2018).

Sample

On June 18, 2018, the U.S. House Intelligence Committee released an updated list of 3,841 Twitter handles associated with IRA activity (Permanent Select Committee on Intelligence, 2018). These handles identify what are commonly called *troll* accounts, meaning they have human operators, as opposed to *bot* accounts which are computer operated. We collected all tweets from these accounts between June 19, 2015, the earliest date searchable at the time of study, and December 31, 2017. Source-based searches are unreliable on suspended accounts. Keyword searches using the term "tweet from" followed by the handle, however, yielded all tweets from a given account. Keyword searches also returned clutter from non-IRA affiliated accounts, most notably variations on IRA handles; e.g. searches for the IRA handle @Broooooke returned tweets from @_Broooooke_.

@_Broooooke_, @_Broooooke_X_ and seven other variations. We removed 17,980 tweets from non-IRA handles.

We identified 24 handles with tweets not associated with IRA agenda building. Six handles engaged in commercial activity (four marketed exercise and diet related activities, another marketed payday loans, and a sixth marketed essay writing services). 18 accounts appeared to engage in normal human behavior.¹ We removed 154,832 tweets associated with these 24 handles.

828,219 tweets from 1,472 separate handles tweeted predominantly in a language other than English. The majority of these were Russian language handles, but handles also tweeted in German, Italian, Arabic, French, and Spanish. In an effort to keep the focus of the current study on the IRA's U.S. operations, these handles were removed.

Finally, 1,034 handles on the list did not tweet during our sample period. 1,875,029 tweets associated with 1,311 IRA handles remained for analysis.

Data Analysis

We both worked to qualitatively analyze each handle, as recommended by Corbin and Strauss (2015), and placed handles into emergent categories. First, we engaged in a process of unrestricted open coding, examining, comparing, and conceptualizing the content. We considered elements of tweets including the hashtags employed by a handle, cultural references within tweets, as well as issues and candidates for which a handle advocated. Many tweets included external links, some of which were usable, and external pages were considered. Finally, the name of the

¹ The presence of handles operated by individuals with no relation to the IRA has been documented. A previous list published by the U.S. House Intelligence Committee in November, 2017 contained four individuals we worked with journalists to identify and speak to (Calderwood, Riglin, & Vaidyanathan, 2018). These individuals, and others, were removed from the updated June, 2018 Congressional list.

handle itself often contained information that helped us better understand its nature (e.g. @BLMSoldier). We conducted axial coding to identify patterns and interpret emergent themes. To verify the validity of results, near the end of axial coding, peer debriefing was conducted (Creswell & Miller, 2000). This involved bringing in an external individual familiar with the phenomenon to play devil's advocate.

178 handles with 10,838 total tweets could not be categorized due to insufficient activity. Some handles appeared briefly just after June 18, 2015, perhaps with tweets occurring prior to the beginning of our data that were unavailable to us. More common, however, were handles that simply tweeted very few times, often in single digits. We do not know if handles stopped tweeting voluntarily or if Twitter suspended the accounts.

The 1,133 remaining IRA-associated handles in our data were placed into one of five categories. We each independently coded a sub-sample of 50 handles and found a Krippendorff's alpha reliability of .92 (note: error occurred only in accounts with extremely low tweet counts).

The raw tweet data were imported into STATA 14.2 for quantitative analysis. The data were collapsed to account-by-day and account-by-hour units of observation, with total tweets tallied, and each account matched with the account-type codes derived in RQ1. We then analyzed the behavior by account type, both over the full period and in specific event windows.

Results

RQ1. We identified five categories of IRA-associated Twitter handles, each with unique patterns of behaviors: *Right Troll*, *Left Troll*, *News Feed*, *Hashtag Gamer*, and *Fearmonger*. With the exception of the *Fearmonger* category, handles were consistent and did not switch between categories.

Right Troll (617 handles, 663,740 tweets, M = 1075.75, SD = 2949.82). These handles broadcast nativist and right-leaning populist messages. They employ common hashtags used by similar real Twitter users, including #tcot, #ccot, and #RedNationRising. Following the nomination of Donald Trump, they uniformly supported his candidacy and his Presidency, e.g. @AmelieBaldwin retweeted on December 13, 2016, “No, Russia didn’t elect Donald Trump, the voters did <https://t.co/ce70G9gv4h> Repeat over and over disbelievers. PRESIDENT DONALD TRUMP!!” These handles regularly employed #MAGA, the acronym for “make America great again,” Donald Trump’s campaign slogan. They routinely denigrated the Democratic Party, e.g. @LeroyLovesUSA, January 20, 2017, “#ThanksObama We're FINALLY evicting Obama. Now Donald Trump will bring back jobs for the lazy ass Obamacare recipients.”

These handles’ themes were distinct from mainstream Republicanism. They rarely broadcast traditionally important Republican themes, such as taxes, abortion, and regulation, but often sent divisive messages about mainstream and moderate Republicans. During the Republican Party primaries, #GOPStop appears frequently in Right Troll tweets, e.g., @amalia_petty, December 16, 2015, “#VegasGOPDebate Asking who is gonna win #GOPDebate is like asking what sort of crap is your favourite?” Similarly, on October 6, 2016, @hyddrox retweeted “The House voted to impeach Koskinen but that JERK McConnell said he didn't have time to take it up on the senate Time to EXIT THE D.C.” in reference to Republican Senate Majority Leader Mitch McConnell.

This category also includes some themed accounts, including @itstimetoseced, which advocated for the secession of Texas, and @Jihadist2ndWife, a parody handle, which presented itself as the wife of an Islamic State fighter. The overwhelming majority of handles, however, had limited identifying information, with profile pictures typically of attractive, young women.

Left Troll (230 handles, 405,549 tweets, M = 1763.26, SD = 2468.32). These handles sent socially liberal messages, with an overwhelming focus on cultural identity. They discussed gender and sexual identity (e.g., #LGBTQ) and religious identity (e.g., #MuslimBan), but primarily focused on racial identity (e.g., #blacklivesmatter). Many handles, including @Blacktivists and @BlackToLive, tweeted in a way that mimicked the Black Lives Matter movement, with posts such as @Blacktivists, May 17, 2016, “Justice is a matter of skin color in America. #BlackTwitter”. Many such tweets seemed intentionally divisive, including @Blacktivists, May 10, 2016, “When you have been handcuffed for no good reason, all you can think about is how not to get shot. Never trust a cop”, or @BlackToLive, September 6, 2016, “they treat us today, not like fellow citizens, but as an insurgency which they must suppress...”.

Just as the Right Troll handles attacked mainstream Republican politicians, Left Troll handles attacked mainstream Democratic politicians, particularly Hillary Clinton. Tweets such as @Blacktivists, October 31, 2016, “NO LIVES MATTER TO HILLARY CLINTON. ONLY VOTES MATTER TO HILLARY CLINTON” and a retweet from @JerStoner, October 7, 2016, “#ClintonBodyCount if anyone else had her rap sheet - they'd be on death row”. Such tweets undermined Clinton’s credibility and spread questionable information about her campaign prior to the 2016 election. In contrast, these handles were supportive of Bernie Sanders prior to the election, with posts such as @blacneighbor, June 13, 2016, “I think many folks took @BernieSanders for granted. I've never seen a politician so passionate about the people!”

News Feed (54 handles, 567,846 tweets, M = 10,515.67, SD = 9275.05). These handles overwhelmingly presented themselves as U.S. local news aggregators and had descriptive names such as @OnlineMemphis and @TodayPittsburgh. They linked to legitimate regional news sources and tweeted about issues of local interest, such as @KansasDailyNews, December 9, 2015,

“#news Barton County finds new revenue with oil well” and on the same day, “#news SW Kansas sheriff says he’s getting calls about welfare of some horses”.

A small number of these handles, including @SpecialAffair and @WarfareWW, tweeted about global issues, often with a pro-Russia perspective. The handle @todayinsyria tweeted on October 11, 2015, “2 civilians killed by terrorists’ gunfire in Sweida countryside <http://t.co/lHbleruLq3>” and on the next day “Russian Air Force destroys 53 targets for ISIS in several areas in Syria <http://t.co/aSBbcfwQkT>”. These link directly to the Syrian Arab News Agency, a Syrian state agency allied with the Russian government.

Hashtag Gamer (110 handles, 216,895 tweets, M = 1955.31, SD = 3176.10). These handles are dedicated almost entirely to playing hashtag games, a popular word game played on Twitter. Users add a hashtag to a tweet (e.g., #ThingsILearnedFromCartoons) and then answer the implied question (Haskell, 2015). These handles also posted tweets that seemed organizational regarding these games, e.g. @AmandaVGreen’s quote tweet, August 31, 2016, “15 minutes till we play @TheHashtagGame with @HashtagRoundup & @HashtagZoo! Who's ready to #hashtag!”. Many of these tweets were mundane, including @DonnieLMiller, April 12, 2017, “#OffendEveryoneIn4Words fart in your face.” Others, however, often using the same hashtag, were socially divisive, including @DonnieLMiller, April 12, 2017: “#OffendEveryoneIn4Words undocumented immigrants are ILLEGALS.” Many tweets from Hashtag Gamers were overtly political, e.g. @LoraGreen, July 11, 2015, “#WasteAMillionIn3Words Donate to #Hillary”. While many tweets shared themes seen in the Right Troll category, Left Troll themes also appeared, e.g., @LoraGreen, January 25, 2016, “#ItsSoWhiteOutsideThat Donald Trump thought it was a meeting of his followers.”

Fearmonger (122 handles, 10,161 tweets, M = 82.79, SD = 60.06). These accounts spread news of a fabricated crisis event— that salmonella-contaminated turkeys were produced by Koch Foods, a U.S. poultry producer, near the 2015 Thanksgiving holiday. The tweets described the poisoning of individuals who purchased these turkeys from Walmart. These included @RitterTra, November 26, 2015, “OMG Obama and Koch bros. are trying to steal our holidays! nice. #USDA” and also @Peter_Downs_Up, November 27, 2015, “woow Whut? Poisoned #turkey on Thanksgiving?! #KochFarms #foodpoisoning #USDA”. Koch Foods has no connection to the Koch brothers, and the story was an IRA fabrication (Washington, 2018).

Handles were coded as Fearmongers if they tweeted exclusively about Koch Foods or if their account ceased tweeting after a period of extensive tweets about this issue. Several handles tweeted briefly in a manner consistent with the Right Troll category but switched completely to tweeting as a Fearmonger and were coded as the latter.

RQ2: Figure 1 displays the daily number of tweets from Left Troll and Right Troll accounts. Figure 2 displays the daily number of tweets from News Feed and Hashtag Gamer accounts. These figures illustrate differences in how the IRA employed account types. The Left Troll and Right Troll accounts were very active, with mean aggregate daily output of 437 and 716 tweets, respectively. But their day-to-day was also highly variable, with between-day standard deviations of 655 and 1347. The News Feed handles, by contrast, tweeted at a quite consistent rate throughout the period, averaging at total of 613 tweets per day and a standard deviation of only 380. The Hashtag Gamer handles’ output rises in early 2016 and are very active during and after the election season, but by summer of 2017 they are nearly silent. They tweet, on average, 234 times per day, with a standard deviation of 610. A Kruskal–Wallis test rejects the null that the daily tweet totals for these four accounts types were pulled from the same distribution ($p < .001$).

The account types also differentially react to political circumstance. Figure 3 zooms in on a four-day period centered on midnight UTC, heading into Oct. 7, 2016. The News Feeds have their normal low and consistent output, while the Hashtag Gamers spike the evening of the 5th and then stop tweeting. Finally, the Trolls (more Left than Right) begin massively producing tweets for at least 14 hours in a row, beginning at noon UTC on the 6th, and continuing with a second spike around 8:00 a.m. UTC on the 7th. WikiLeaks released the first batch of the hacked Podesta emails around 8:30 p.m. UTC on the 7th. A very similar pattern of behavior can be seen in the enormous spike in Right Troll accounts in late July and early August, 2017, when the other account types quit operations as the IRA focused resources in Right Trolls.

Figure 4 zooms in on a three-day period centered on Oct. 19, 2016, the day of the final Presidential debate between Donald Trump and Hillary Clinton (which began at 1:00 a.m. UTC on the 20th). This figure shows spikes in IRA use of Hashtag Gamer handles and includes many tweets addressing the debate, such as @Gab1Aldana, “#RejectedDebateTopics how to turn your Foundation into money laundering machine.” There was very limited activity by other account types over this period, and, as with the pre-WikiLeaks efforts, major output occurred in advance of the event.

Finally, Figure 5 zooms in on a ten-day period beginning on Sept. 11, 2016. Over this entire period, Left Trolls activity is low. Even within work periods, the troll operators are specializing on one account type at a time, beginning with Hashtag Gamers and following up with Right Trolls on Sept. 12, but reversing the order on Sept. 14. The enormous plateau of Right Troll activity beginning around noon on the 16th and rising through the end of the 17th is regular Right Troll activity, but significantly spikes on the 18th in response to the Chelsea bombing at 1:30pm UTC on the 17th.

These three periods make clear that the IRA allocated their efforts amongst the account types very differently when faced with different political circumstances or shifting goals. In all three periods, a Kruskal–Wallis test rejects the null that the daily tweet totals for these four accounts types were pulled from the same distribution ($p < .001$). Without better information about their goals, it is difficult to speculate about exactly what underlying strategy is driving these shifts, but accounts of different types are not good substitutes for each other--- each plays a different role and is used differently.

Conclusion

The IRA efforts in our sample period were conducted systematically. Handles were built into one of five groups and then used as interchangeable parts depending on allocation needs or Twitter bans. Effort was reallocated amongst account types in response to shocks, depending on the segment of the U.S. electorate the IRA wished to engage, changing IRA strategic goals, or both.

Understanding how governments work to influence other nations' domestic agendas is important, and the IRA Twitter operation is an important example from the digital age. At a February 13, 2018 U.S. Senate Intelligence Committee hearing, Senator Mark Warner stated that social media companies have been “slow to recognize the threat” that Russian influence poses (Nakashima & Harris, 2018, para. 7). At that same hearing, U.S. Director of National Intelligence Daniel Coats said of Russian efforts to disrupt the 2016 election, “There should be no doubt that Russia perceives its past efforts as successful” (para. 9). The Director then warned of the certainty of future Russian interference.

For this reason, future research will need to examine IRA efforts further. The data collected for this study can be used to analyze the qualitative nature of individual tweets and to give a more detailed understanding of the effectiveness of this campaign over time. These data can also be used to better understand how the IRA's tactics adapt over time and, by analyzing the non-English tweets, in various contexts.

Russia's attempts to distract, divide, and demoralize has been called a form of political war (Galeotti, 2018). This analysis has given insight into the methods the IRA used to engage in this war. One former employee of the IRA described the feeling of working there as though "you were in some kind of factory that turned lying, telling untruths, into an industrial assembly line" (Troianovski, Helderman, Nakashima, & Timberg, 2018). The systematic and organized nature of the messaging we have analyzed here suggests this employees feeling was correct. The IRA is engaging in what is not simply political warfare, but industrialized political warfare.

References

- Barrett, D., Horwitz, S. & Helderman, R. S. (2018, February 17). Russians Indicted in 2016 election interference. *The Washington Post*. p. 1A.
- Black, D.R., Dietz, J.E., Stirratt, A.A., & Coster, D.C. (2015). Do social media have a place in public health emergency response? *Journal of Emergency Management*, 13, 217–226. doi:10.5055/jem.2015.0235
- Calderwood, A., Riglin, E., & Vaidyanathan, S. (2018, July 20). How Americans ended up on Twitter's list of Russian bots. *Wired*. Retrieved from <https://www.wired.com/story/how-americans-wound-up-on-twitters-list-of-russian-bots/>
- Cobb, R. W., & Elder, C. D. (1971). The politics of agenda-building: An alternative perspective for modern democratic theory. *Journal of Politics*, 33, 892-915. doi:10.2307/2128415
- Corbin, J., & Strauss, A. (2015). *Basics of qualitative research*. Thousand Oaks, CA: Sage.
- Creswell, J. W. (2014). *Research design: Qualitative, quantitative, and mixed methods approaches*. Thousand Oaks, CA: Sage.
- Denham, B. E. (2010). Toward a conceptual consistency in studies of agenda-building processes: A scholarly review. *The Review of Communication*, 10, 306-323. doi:10.1080/15358593.2010.502593
- Edgerly, S., Thorson, K., Bighash, L., & Hannah, M. (2016). Posting about politics: Media as resources for political expression on Facebook. *Journal of Information Technology & Politics*, 13, 108–125. doi:10.1080/19331681.2016.1160267

Galeotti, M. (2018, March 5). I'm sorry for creating the 'Gerasimov Doctrine'. *Foreign Policy*. Retrieved from <http://foreignpolicy.com/2018/03/05/im-sorry-for-creating-the-gerasimov-doctrine/>

Nakashima, E. & Harris, S. (2018, February 13). The nation's top spies said Russia is continuing to target the U.S. political system. *The Washington Post*. Retrieved from https://www.washingtonpost.com/world/national-security/fbi-director-to-face-questions-on-security-clearances-and-agents-independence/2018/02/13/f3e4c706-105f-11e8-9570-29c9830535e5_story.html?utm_term=.9d39f53cf636

Haskell, W. (2015). People explaining their 'personal paradise' is the latest hashtag to explode on Twitter. *Business Insider*. Retrieved from <http://www.businessinsider.com/hashtag-games-on-twitter-2015-6>

Introducing Social Studio REST API (n.d.). Retrieved from https://developer.salesforce.com/docs/atlas.en-us.api_social.meta/api_social/1-introduction-to-the-social-studio-marketing-cloud.htm

Lariscy, R. W., Avery, E. J., Sweetser, K. D., Howes, P. (2009). An examination of the role of online social media in journalists' source mix. *Public Relations Review*, 35, 314-316. doi:10.1016/j.pubrev.2009.05.008

Linville, D. L., Boatwright, B. C., & Grant, W. J. (2018). "Back-stage" dissent: Student Twitter use addressing instructor ideology. *Communication Education*, 67, 125-143. doi:10.1080/03634523.2018.1428998

Matsakis, L. (2017, November 3). Twitter told congress this random American is a Russian propaganda troll. *Vice*. Retrieved from

https://motherboard.vice.com/en_us/article/8x5mma/twitter-told-congress-this-random-american-is-a-russian-propaganda-troll

Parmelee, J. H. (2014). The agenda-building function of political tweets. *New Media & Society*, 16, 434-450. doi:10.1177/1461444813487955

Permanent Select Committee on Intelligence (2018, June 18). Schiff statement on release of Twitter ads, accounts and data. Retrieved from: <https://democrats-intelligence.house.gov/news/documentsingle.aspx?DocumentID=396>

Troianovski, A., Helderman, R. S., Nakashima, E., & Timberg, C. (2018, February 17). The 21st-century sleeper agent is a troll with an American accent. *The Washington Post*. Retrieved from https://www.washingtonpost.com/business/technology/the-21st-century-russian-sleeper-agent-is-a-troll-with-an-american-accent/2018/02/17/d024ead2-1404-11e8-8ea1-c1d91fcec3fe_story.html?noredirect=on&utm_term=.d5906ace8983

Washington, B. (2018, February 22). Inside Russia's fake news HQ. *The Australian*. p. INQUIRER 11.

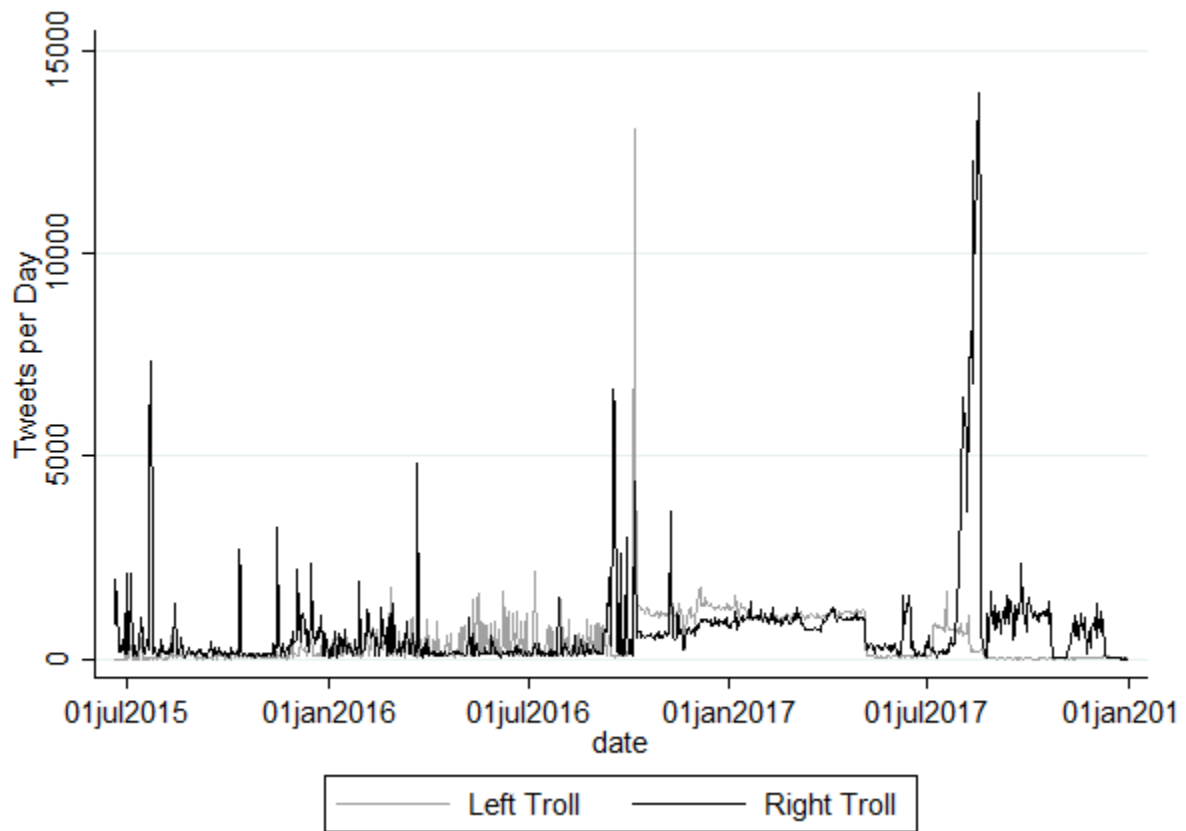


Figure 1. Daily tweets by Left Troll and Right Troll accounts, Jun 19, 2016 – Dec 31, 2017.

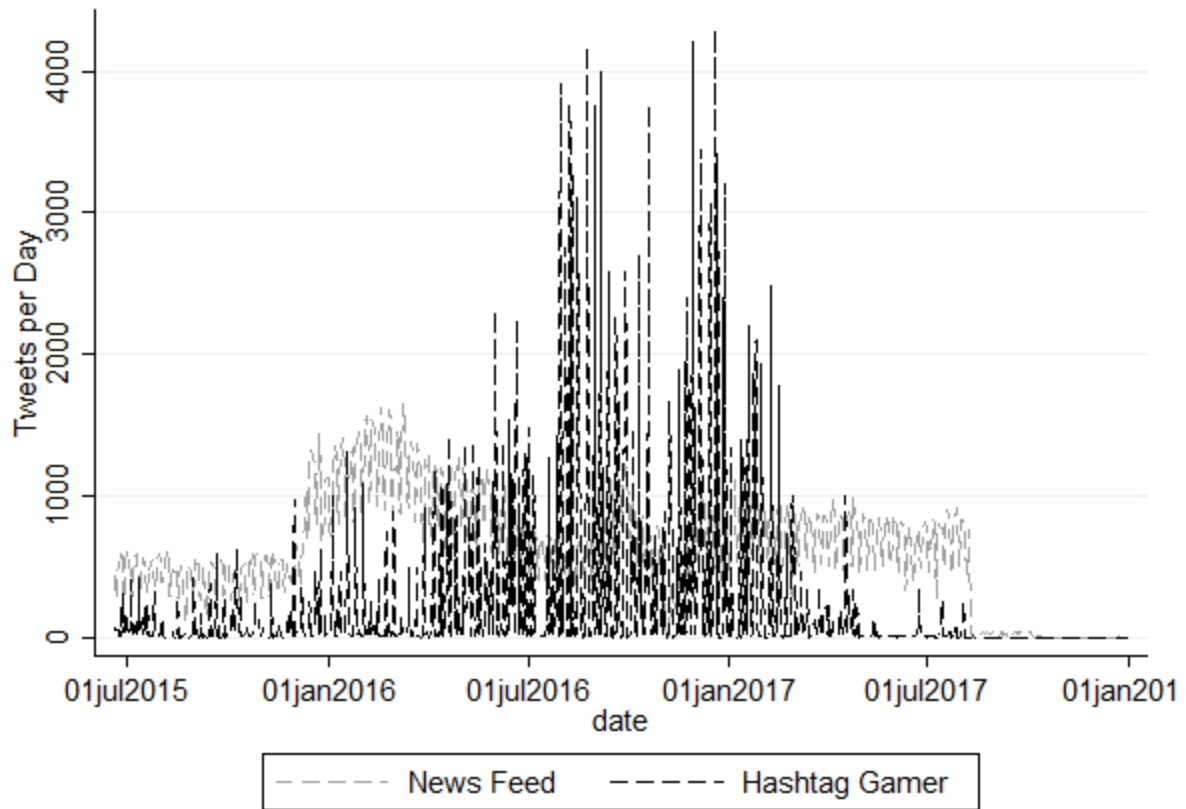


Figure2. Daily tweets by News Feed and Hashtag Gamer accounts, Jun 19, 2016 – Dec 31, 2017.

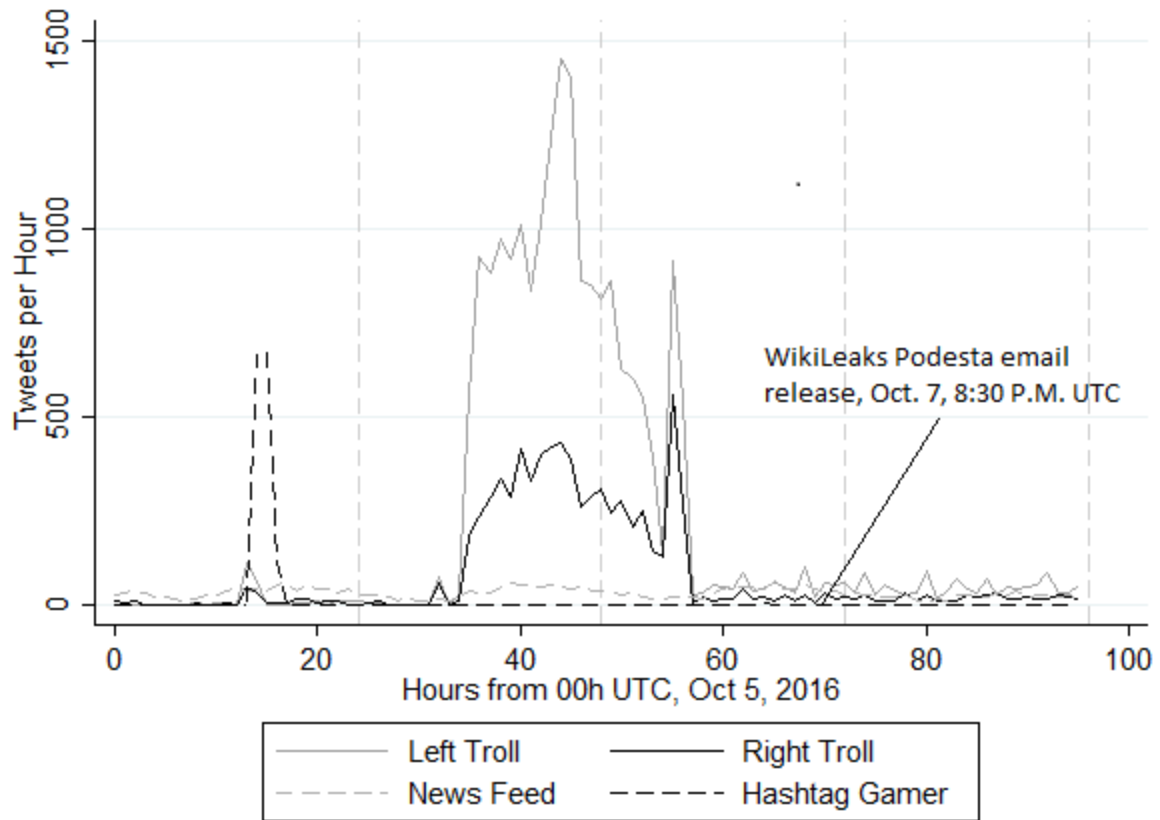


Figure 3. Hourly tweets by four largest account types, Oct 5-Oct 8, 2016.

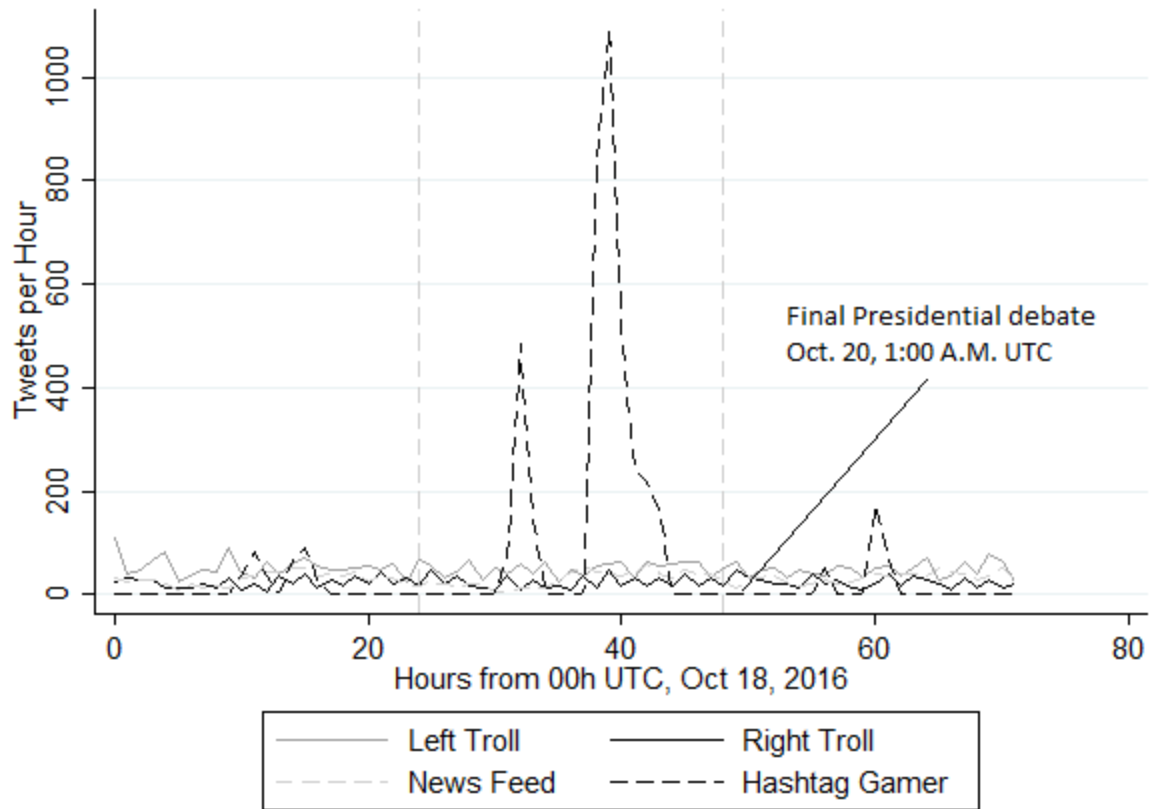


Figure 4. Hourly tweets by four largest account types, Oct 18-Oct 20, 2016.

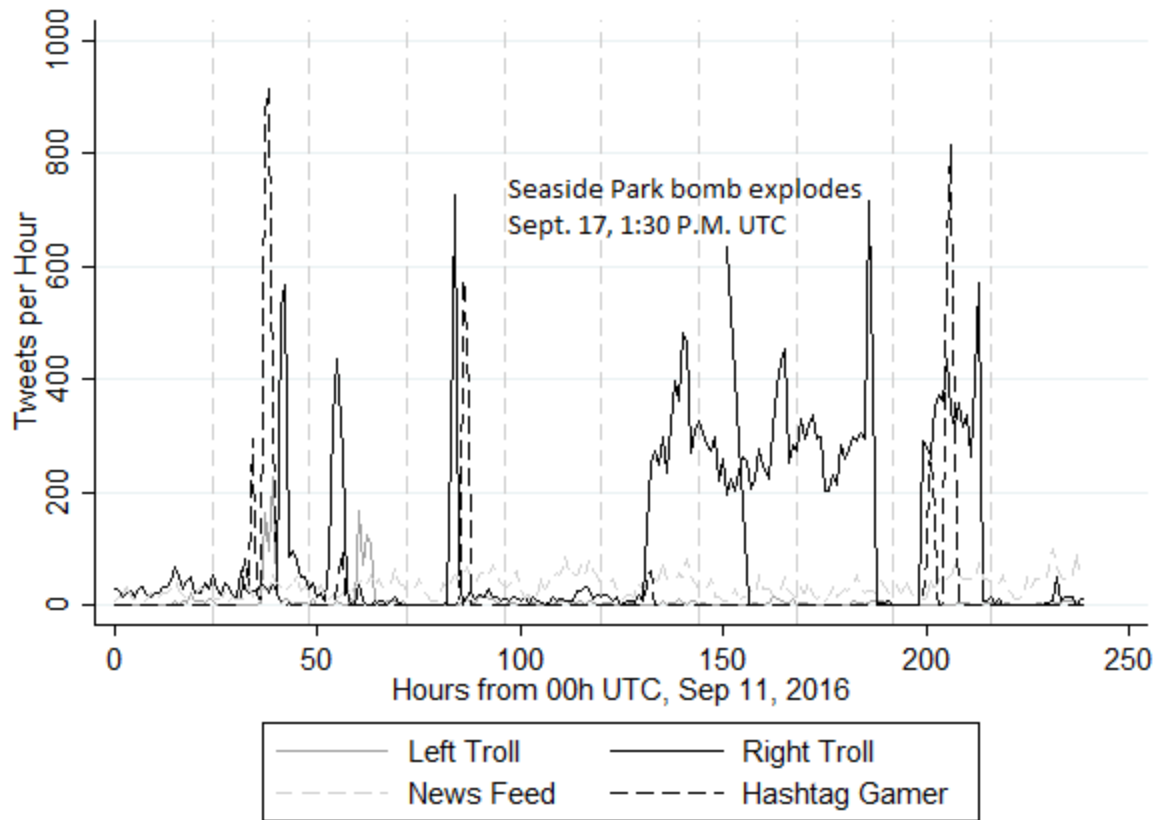


Figure 5. Hourly tweets by four largest account types, Sep 11-20, 2016.