International Journal of Communication 17(2023), 6269–6290 1932–8036/20230005

Russian Troll Social Media Attacks on Presidential Candidates during the 2016 U.S. Election: The Role of Frontrunner Status, Political Party, and Candidate Gender

LARISSA TERÁN◆
The Geena Davis Institute on Gender in Media, United States

HEATHER GAHLER
DANIEL MONTEZ
KATE KENSKI
STEPHEN A. RAINS
University of Arizona, United States

This study examines how the Russian Internet Research Agency (IRA) troll attacks against the 2016 U.S. presidential candidates on social media varied based on three important characteristics: frontrunner status, political party affiliation, and the gender of the candidate. The frequency of attacks, types of attacks, and audience engagement via retweets were assessed. A content analysis of 4,518 IRA troll messages posted on Twitter (i.e., tweets) shows that frontrunners, Democrats, and the female candidate received the most attacks. In terms of attack types, attacks on character/integrity occurred most frequently and were more likely to be directed at frontrunners, Democrats, and the female candidate. Tweets attacking these three groups were also more likely to be retweeted than tweets without an attack.

Keywords: social media, candidate traits, political attacks, presidential campaign, frontrunners, gender

Fundamental changes to the political communication landscape have taken place in the past decade. Social media have created opportunities for nontraditional actors to influence political discourse and, with that, the attitudes, intentions, and behaviors of the public. Kenski and Jamieson (2017) observed, "During the past several decades, media systems and the loci of power both were national.

Larissa Terán: larissateran@arizona.edu
Heather Gahler: hgahler@arizona.edu
Daniel Montez: danielmontez@arizona.edu
Kate Kenski: kkenski@arizona.edu
Stephen A. Rains: srains@arizona.edu
Date submitted: 2022-11-06

By contrast, in the current environment media structures are increasingly transnational, and the nation-state model is challenged by nonstate actors" (p. 914). The current environment has created possibilities for states to use nontraditional actors to shape public opinion. In 2016, the Russian Internet Research Agency (IRA) deployed actors known as "trolls" on Twitter to disrupt the U.S. presidential election (Jamieson, 2018). Trolls infiltrated political discourse by presenting themselves as everyday citizens and crafting tweets of "division, discontent, and disconnection with reality among U.S. political discussions" (Linvill, Boatwright, Grant, & Warren, 2019, p. 292). Understanding how these actors engage the public is important for shifting agency and power back to the stakeholders within an electoral system and preventing future election interference.

Scholars exploring political troll behavior on Twitter have focused on elections, reporting that trolls tweeted general attacks against candidates, the media, and civil institutions (Linvill et al., 2019). Incivility was another tactic used by trolls to target candidates, with a peak in name-calling against Democratic nominee Hillary Clinton during the general election (Rains, Shmargad, Coe, Kenski, & Bethard, 2021). We build on this work by investigating how IRA trolls exploited the patterns of candidate trait attacks known to take place by candidates, campaigns, and news media in previous political campaigns. We examine troll attacks lodged toward the four most viable presidential candidates (i.e., Donald Trump, Hillary Clinton, Bernie Sanders, and Ted Cruz) to see whether the trolls followed certain patterns based on candidates' frontrunner status, political party affiliation, and gender. Our investigation is a case study of the role of new actors in the changing political landscape: trolls. The findings may inform other important research areas about digital interference.

Traditionally, news media, campaign advertisements, and groups within the United States have primed candidate information, privileging some traits over others. The IRA trolls were a novel source of attacks in 2016. They established themselves as social influencers for priming which candidate traits were considered important to the media and the general public. Although general election nominees Trump (Republican) and Clinton (Democrat) were often the targets of online incivility (Rains et al., 2021), we needed to know more about *the kinds of attacks* being made, whether they differed based on *candidate characteristics*, and the degree to which attacks *resonated* with Twitter users. This information illuminates the nuances of the IRA trolls' efforts and provides a benchmark for scholars studying political communication on social media.

Political Trolls

Online trolling is defined as an intentional effort to disrupt online discourse through manipulative behaviors, such as aggression, aggravation, and disruption (Uyheng, Moffitt, & Carley, 2022). The goal of trolls is to instigate discord among online communities (Rains et al., 2021). Unlike the subcultural trolls that seek to gain personal satisfaction by provoking anguish among individuals and communities, our study is primarily concerned with state-sponsored trolls who form part of strategic efforts to undermine democratic processes by spreading disinformation. Both types of trolls are motivated to subvert traditional systems and may engage in similar behaviors, but their ultimate goals are distinct. State-sponsored trolling involves broader goals that are larger than the content of individual posts or tweets (Starbird, 2019). We refer to these types of actors as *political trolls*. The Russian government-backed IRA used political trolls to amplify

U.S. social tensions and partisan polarization through a cross-platform strategy on social media, starting well before and throughout the 2016 U.S. presidential election.

An objective of state-sponsored trolls is to incorporate the lay public in their dissemination of disinformation (Uyheng et al., 2022). Ruck, Rice, Borycz, and Bentley (2019) found that 91% of Twitter users who retweeted IRA tweets first were *not* IRA bots. Not only did IRA trolls attack individuals and groups in 2016 but also assisted in boosting attitudes around the entities they supported. Badawy, Addawood, Lerman, and Ferrara (2019) analyzed how the lay public spread original content generated by trolls on social media, observing that conservative users were significantly more likely than liberal users to retweet Russian trolls. Those who engaged with IRA accounts significantly increased their Twitter activity after initial IRA contact (Dutta et al., 2021).

Examining the content produced by IRA trolls on social media helps us understand the environment in which state-sponsored actors attempt to influence American voters' evaluations of candidates and policies. Prior research has identified some general patterns in the behavior of trolls. Rains et al. (2021), for example, observed that Trump received the greatest proportion of name-callings during the preprimary and primary periods, but Clinton became the primary target during the general election. Although this work helps to understand trends in IRA troll behavior, finer-grained research is needed to examine the strategies exhibited by state-sponsored trolls to promote and attack candidates.

Trolls are a consequence of the politically mediated environment in which they act. It is possible that the negative influence of the IRA trolls in 2016 was complemented, not counterbalanced, by the style in which traditional media normally cover presidential elections. Although the media coverage for the two frontrunners—Trump and Clinton—was expressly negative, Trump was a press magnet for his sensationalism, which garnered three times the coverage Clinton received (Dunway & Graber, 2023). Rather than media stories focusing on leadership, experiences, or policy stances, most of the coverage focused on the horse race (42%), followed by controversies (17%). Clinton's negative media coverage doubled her positive media coverage, potentially shaping how American voters came to evaluate Clinton. In this light, we sought to understand how the IRA trolls attacked political candidates. We examined frontrunner status, party identification, and candidate gender to better understand the form and frequency of the attacks.

Political Attacks

Political attacks are targeted statements that create a negative image of a candidate by highlighting negative aspects of the candidate's record, character, or position (Pfau, Parrott, & Lindquist, 1992). Research has examined attacks by political campaigns using television ads and public stages (Jamieson, 1995). Given that candidate traits influence voting decisions (Kenski, Hardy, & Jamieson, 2010) and that political social media attacks are more effective at generating engagement (e.g., retweets) than nonattack posts (Hemsley, 2019), trait attacks on social media may be important because of their potential to shape voting behavior. Traits also play a significant role in how candidates are attacked (Fridkin & Kenney, 2011), and characteristics such as party and gender may make certain types of attacks more likely. In this study, we considered differences in IRA attacks on candidates by frontrunner status, party affiliation, and gender.

Frontrunners

The competitive status of candidates may affect the media coverage and attacks they receive. Research analyzing traditional media suggests that although political frontrunners often receive more news coverage than less prominent candidates (Steger, 1999), they are also more likely to receive negative coverage and be the targets of attacks by trailing candidates (Haynes & Rhine, 1998).

Research has analyzed how frontrunner status affects online attacks against candidates (Rossini et al., 2018). Gross and Johnson (2016) suggest that past patterns are applicable in online spaces. In 2016, Stein and Benoit (2021) found that incumbents attacked fewer than challengers. Rossini and colleagues (2018) observed that leading gubernatorial candidates were less likely to attack their competitors on Facebook and Twitter than candidates who were tied or trailing. Much of this work has focused only on the attacks initiated by opposing candidates against frontrunners. To our knowledge, no empirical analyses have examined the types of state-sponsored troll attacks on political frontrunners in comparison to trailing candidates. We argue that in their attempts to sow discord, malicious online state-sponsored actors attacked frontrunners more than trailing candidates, following patterns already engaged by media and political opponents.

Party Affiliation

Candidate traits are not only representative of themselves as individuals, but of their respective parties. The American public sees Republicans as stronger leaders and more moral, whereas Democrats are perceived as more empathetic and compassionate (Hayes, 2005). Party stereotypes may be placed on individual candidates, who are evaluated based on how well they represent party values (Bartels, 2002).

Goren (2002) observed that partisan bias leads party identifiers to focus on the perceived character weaknesses of opposition candidates. Traits are evaluated not only against a candidate's opponent but also against party stereotypes. This appeared to take place in 2016, particularly on social media. By posing as members of the voting public, IRA trolls utilized partisan attacks on Twitter (Linvill et al., 2019), which brought forth salient party traits and candidate traits. Rains et al. (2023) showed that messages from partisan IRA trolls related to the 2016 presidential election on Twitter were particularly likely to be shared by the lay public via retweeting. Their study, however, did not capture nuances in the attacks made by trolls on candidates.

Candidate Gender

The 2016 U.S. presidential election was the first to have a major party female nominee. Female presidential candidates have often received biased press treatment, such that their presence is often treated as a novelty rather than as a serious option (Falk, 2018). Female candidates are more likely to be seen as possessing traits associated with warmth, such as compassion and empathy, whereas males are more likely to be viewed as possessing traits associated with competence, such as leadership (Hayes, 2011). Gender stereotypes may place females in a double bind (Jamieson, 1995). That is, women may be perceived as gullible or spineless but also disliked when presenting traits that appear too masculine.

Trait perceptions differ when gender comes into play. When male candidates exhibit power-seeking behaviors, they are viewed as more competent, assertive, and strong, eliciting a male candidate preference (Okimoto & Breskoll, 2010). When similar power-seeking traits are shown in female candidates, they are viewed as lower in competency and agency, with people expressing more anger, disgust, and contempt toward the candidate. Banwart and Kearney (2018) found that in 2016, benevolent sexism (i.e., comments that seem positive but imply inferiority to men) was not a predictor of perceptions of Clinton's traits, but it did significantly and positively predict perceptions of the same traits for Trump. Moreover, hostile sexism (i.e., blatant and disrespectful) was a significant predictor of both negative evaluations of Clinton's traits and positive evaluations of Trump's traits. These findings suggest that gender was an important candidate difference and may have influenced the IRA trolls in 2016 as it did citizens.

The Present Study: The Case of the IRA Trolls and the 2016 U.S. Election

Research on Russian interference in 2016 has generally focused on messages shared by trolls on Twitter (Linvill et al., 2019; Rains et al., 2021). Questions remain about the types of attacks trolls made against candidates in their efforts to sow discord. In this study, we focus on tweets made by IRA trolls about leading candidates Trump, Clinton, Sanders, and Cruz. We chose these candidates because news coverage and polling, showing them as the frontrunners of their parties. Understanding the nuances of the trolls' efforts in targeting these candidates is important because such information offers a valuable benchmark from which to understand political communication on social media and may help to develop defenses against future attacks.

We first examine how important the characteristics of presidential candidates are related to the specific candidates whom the trolls attacked most frequently and the nature of their attacks. Drawing from the scholarship on political communication, we expect that trolls may have targeted the same characteristics that are traditionally exploited by rival campaigns and the media. To enhance authenticity, trolls were likely to target candidates in ways that elicited reactions from average citizens. Research on political campaigns indicates that leading candidates are more likely to be the focus of attacks than others (Haynes & Rhine, 1998), and challengers often go on the attack when they are behind (Rossini et al., 2018). We expect this pattern to continue among the trolls, leading to our hypothesis.

H1: *Frontrunner presidential candidates will be more likely to be attacked by IRA trolls than will challengers.*

Research on gender and party affiliation offers less of a guide about the volume and nature of attacks made by IRA trolls. Although both Republicans and Democrats have been affected by perceptions of moral rectitude or competence based on the individual actions or circumstances that have taken place while in office, one would be hard-pressed to declare that one party's candidates should receive more competence or character attacks than another as a general expectation. Similarly, existing research demonstrates that female candidates are evaluated differently than male candidates (e.g., Hayes, 2011), but it is unclear whether they ought to receive more specific types of attacks from IRA trolls on Twitter. Accordingly, we pose research questions about differences in the quantity and nature of attacks based on party affiliation and gender.

RQ1: Is there a difference in the likelihood of an attack on (a) Republican candidates compared with Democratic candidates or (b) male candidates compared with female candidates made by IRA trolls?

RQ2: Is there a difference in the types of attacks on (a) frontrunner presidential candidates compared with challengers, (b) Republican candidates compared with Democratic candidates, or (c) male candidates compared with female candidates launched by IRA trolls?

We also considered the degree to which their efforts resonated with Twitter users in the form of retweets. As a measure of audience engagement, retweets capture the number of users who shared a particular message authored by an IRA troll with their social network on Twitter. Although the literature does not directly tell us whether greater engagement in more retweets of frontrunner attacks relative to challenger attacks should be expected, it is reasonable that the attacks on the leading candidates would result in greater engagement. This logic follows from research showing that frontrunners receive more news coverage (Steger, 1999) and are more likely to be attacked on social media (Gross & Johnson, 2016) than trailing candidates. We assume that the same mechanism that yields news coverage or attacks is operative when people decide to retweet attacks. However, the existing research on party affiliation and gender has been insufficient to predict the implications of these characteristics for retweets, leading us to pose research questions.

H2: IRA troll attacks on frontrunner presidential candidates will receive greater engagement relative to attacks on challengers.

RQ3: Is there a difference in audience engagement with attacks by IRA trolls on (a) Republican relative to Democrat candidates or (b) male relative to female candidates?

Method

Data and Sample

The data for this project were acquired directly from Twitter. The procedure used to identify IRA trolls was not disclosed by Twitter, but the data were from the same accounts that Twitter supplied to the U.S. Congress as part of the inquiry into election interference. To identify tweets about Trump, Clinton, Sanders, and Cruz, we first searched the text of approximately 3 million troll tweets written in English for the first and last names of each candidate. For Clinton, we also included the common misspelling "Hilary" as well as the initial "HRC." We excluded retweets to focus on messages originating from IRA trolls and limited our sample to tweets made between January 1 and election day (November 8) in 2016. There were 35,731 original tweets about one or more of the four candidates posted during this time period: Trump (*n* = 19,716), Clinton (*n* = 14,710), Sanders (*n* = 5,025), and Cruz (*n* = 3,126).[1]

Next, we constructed a stratified random sample of tweets about the four candidates from January 1 to election day. First, we identified three time periods corresponding to key phases of the election: primaries (January 1–June 14), conventions (June 15–July 28), and general election (July 29–November 7). Within each timeframe, we randomly sampled 10% of the tweets addressing each of the four candidates. Because some candidates were rarely discussed during certain periods (e.g., Cruz during the general election), we retained a minimum of 100 tweets for each candidate and time period. This process resulted in 4,518 total tweets: Trump (*n* = 2,587), Clinton (*n* = 2,083), Sanders (*n* = 836), and Cruz (*n* = 544).[2] These tweets served as the sample for the present study.

Candidate Traits

We operationally defined an attack as a negative message targeting a candidate and calling attention to a candidate's weakness (Pfau & Kenski, 1990). Given the dearth of research on IRA troll attacks on candidates and the unique nature of the data, we used a grounded theory approach to develop a codebook (Strauss & Corbin, 1990). Over a five-week span, we completed an interactive process in which we reviewed sample IRA tweets, discussed the attacks that were most frequent, and generated categories and operational definitions. This process led to the identification of seven types of attacks: competence, character/integrity, age, appearance, sexual reference, gendered name-calling, and general (nonspecific) attacks.[3] The operational definition and sample tweets for each category can be found in Table 1.

[1] The total for each candidate sums to greater than the total sample size because each tweet can address more than one candidate. The names of Clinton and Trump, for example, may appear in the same tweet.
[2] Again, the total for each candidate sums to greater than the stratified random sample size because each tweet can address more than one candidate.
[3] We included an eighth category involving e-mails but have excluded it from this project because it was only relevant to Clinton and appeared infrequently.

Table 1. Attacks Codebook.

Category	Definition	Examples
Appearance	Refers to the candidate's appearance with negative or positive[4] descriptors that may objectify the candidate. Appearance-based attacks are not a relevant trait to consider for the role of POTUS and are insulting.	Example 1: "If I'll vote for hillary she would reveal the secret why she is so ugly ? #HillaryPickUpLines" (personal communication, April 18, 2016). Example 2: "#HighSchoolTaughtMe not to go near snakes, especially if they're ugly. So why do you think I stay away from #HillaryClinton ….!" (personal communication, August 12, 2016).
Sexual reference	Contains a sexual undertone, often describing the candidate as sexually active. It also includes acts of sexual aggression or describes the candidate as a sexual predator. Sexual reference attacks take away from the qualities needed to fulfill the job and instead attack or mock a candidate's personal life. Sexual references signal disrespect to the candidate.	Example 1: "#HillaryPickUpLines I think you know, that there are a lot of blow jobs in the Oval Office" (personal communication, April 18, 2016). Example 2: "#TrumpsFavoriteHeadline BREAKING: Leaked Ivanka Sex Tape" (personal communication, August 17, 2016).
Competence	Questioning the qualifications, expertise, or experience of the candidate. Competence attacks make candidates seem dumb, unprepared, and/or amateurish.	Example 1: "#donttellanyonebut I still believe in unicorns. I'm more likely to meet one than Hillary is to meet people's expectations" (personal communication, August 10, 2016). Example 2: "#politics Obama: trump "woefully unprepared for presidency (personal communication, August 2, 2016).
Character/integrity	Questioning the trust, honesty, morality, ethics, carelessness, or integrity of the candidate. It can also include questioning the candidate's goodwill or concern about others. Character/integrity attacks make candidates appear immoral and/or dishonest.	Example 1: "trump used charity's money to settle his legal disputes #politics." (personal communication, September 20, 2016). Example 2: "#politics Donald trump: 'ted cruz is a total liar'" (personal communication, January 31, 2016).

[4] We consider a comment about appearance to be an attack when it focuses on superficial traits about the candidate rather than traits regarding competence or ability to fulfill the job.

Age	Discussion of candidate's deficiencies in relation to age. Includes questions about mental capacities, physical conditions, or health conditions. Age attacks make the candidate appear frail and unfit for office.	Example 1: "poll: 71% of 250 physicians seriously concerned about #hillaryshealth" (personal communication, September 8, 2016). Example 2: "old socialist buys $600k summer house after the sellout to Hillary Clinton #feelthebern #nomorerefunds" (personal communication, August 10, 2016).
Gendered name-calling	Name-callings that are sexist in nature and meant to be targeted toward males or females. Often includes vulgar and crass language. Gendered name-calling attacks discriminate based on gender and use specific language reserved strictly for men or women.[5]	Example 1: "trump that bitch #trump2k16 #grabherbythepussy, everybody a closet trump fan he gon win, fuck hillary , fuck hillary #trumpforpresident" (personal communication, November 8, 2016). Example 2: "Trump calls Cruz a hypocrite: 'a nasty guy' #NewYork" (personal communication, January 17, 2016).
General (nonspecific) attacks	Attacks a candidate without substance, context, or elaboration. Insufficient information is placed into other categories.	Example 1: "I'm sick of libtards but Hillary is even worse! She is pure evil!" (personal communication, June 20, 2016). Example 2: "The Benghazi butcher should go to hell. #NeverHillary" (personal communication, September 3, 2016).

Note. More than one type of attack may appear in a tweet.

Candidate *competence* was defined in terms of tweets that questioned the qualifications, expertise, or experience (e.g., past performance) of the candidate. *Character/integrity* included any tweets that questioned the trust, honesty, morality, ethics, carelessness, or integrity of a candidate. Attacks related to *age* discussed the candidate's supposed age-related deficiencies. The *appearance* category included any reference to a candidate's physical appearance. Appearance attacks can use negative (e.g., ugly) or positive[6] (e.g., attractive, beautiful) descriptors. The *sexual reference* category had a sexual undertone, including a description of the candidate being sexually active or otherwise involved in acts of sexual aggression or predation. The *gendered name-calling* category included gender-based name-callings targeting male (e.g., dick, pussy, nasty guy) or female (e.g., nasty woman, bitch, slut, prude) candidates. Finally, the *general attacks* category included attacks that targeted a candidate without substance, context, or elaboration.

[5] Gendered name-calling has specific words used to discriminate against men and women; for instance, pussy, bitch, and "nasty woman" are offensive terms mostly used towards women whereas dick and "nasty guy" are offensive terms mostly used towards men.

[6] Bringing attention to the candidate based on their appearance was deemed an attack. Seemingly positive descriptors can be attacked when focused on appearance because they are commonly used as a form of benevolent sexism. However, we did not identify any of these attacks in the sample.

These tweets did not include enough information to be placed into any other category (e.g., calling a candidate "crazy" or a "loser").

An intercoder agreement was established across the three authors who served as coders. Human coders, rather than computer-based forms of coding, were valuable in our research because coders needed context to inform their decision-making processes about whether a tweet constituted an attack. For example, various tweets included hints of sarcasm that a dictionary-based or machine-learning approach could not capture. Coders were instructed to evaluate tweets for the presence or absence of each measure. Multiple attacks could appear in a single tweet. Once the group achieved an acceptable level of agreement across all categories, the 4,518 tweets in the stratified random sample were evaluated. Each coder examined approximately 500 tweets, with 150 tweets overlapping to assess agreement at multiple time points. This process was repeated three times to evaluate the entire stratified random sample.

Because some variables appeared infrequently, we used Gwet's AC1 to determine agreement. Gwet's is a preferable tool for intercoder reliability when there is a skewed distribution in which one category is over- or under-represented, yet there is also high agreement between coders (Gwet, 2018). This was the case in our sample. Gwet's AC1 ranges from 0 to 1, with large values indicating higher levels of agreement. Coefficients between 0.80 and 1.0 are described as "very good." Intercoder agreement for all variables across the first ($AC1_{mean} = .97$, $AC1_{min} = .88$, $AC1_{max} = 1$), second ($AC1_{mean} = .97$, $AC1_{min} = .88$, $AC1_{max} = 1$), and third ($AC1_{mean} = .98$, $AC1_{min} = .93$, $AC1_{max} = 1$) iterations of coding was acceptable. The agreement estimates for each individual category and round appear in the Open Science Framework (OSF) dedicated to this project.[7]

Frontrunner Status, Party Affiliation, and Candidate Gender

Frontrunner status was determined based on polling during the primaries in which Trump led Cruz in the Republican primary (Real Clear Politics, 2016) and Clinton led Sanders in the Democratic primary (FiveThirtyEight, 2016). Tweets that addressed Trump or Clinton were consolidated into a frontrunner group (n = 3,850), and tweets about Sanders or Cruz were aggregated into a challenger group (n = 1,362).[8] Party affiliation was determined based on the party membership of each candidate. Tweets about Trump or Cruz formed the Republican group (n = 2,878), and tweets about Clinton or Sanders formed the Democratic group (n = 2,545). Finally, tweets about Clinton were included in the female candidate group (n = 2,083), and tweets about any of the three remaining candidates were aggregated in the male candidate group (n = 3,581). Tweets that addressed both groups in a set (e.g., Republican and Democrats; male and female candidates) were omitted from the analyses.

---

[7] https://osf.io/zj3x6/?view_only=7ff98bd77cf746deaecfeca4cd1516a8

[8] We recognize that Clinton might be considered the frontrunner during the general election, based on polling at the time. To address this issue, we reanalyzed the data involving the frontrunner variable limiting the data to only tweets shared during the primary and convention time periods. The results of those analyses for the frontrunner variable followed the same trends as the results reported in this article with the tweets from all time three periods included.

Audience Engagement

Audience engagement in the form of retweets was evaluated using meta-data supplied by Twitter. Retweet statistics were included, showing the number of times each troll tweet had been retweeted. Because Twitter removed all troll accounts from the service at the same time, the retweet values reflected the final total number of retweets that a given tweet accrued.

Procedure for Data Analysis

H1, RQ1, and RQ2 were evaluated using pairwise proportion tests. For each specific type of attack (e.g., competence, character/integrity, etc.), tweets were assigned a value of 1 when the attack was present and 0 when the attack type was absent; for the overall attack variable, tweets were assigned a value of 1 when any of the specific types of attacks were included and 0 when all attack types were absent. The total number of tweets about an attack target group was identified along with the proportion that included each type of attack. For frontrunners and challengers, for example, the total number of tweets was identified for each group, and then the proportion of tweets that included each type of attack (e.g., competence) was computed. The difference in proportions between groups was then evaluated for each attack type.

H2 and RQ3 were evaluated using multilevel modeling. The unit of analysis was individual tweets. The outcome variable was the number of retweets a tweet received. Random intercepts were included in the multilevel models for the troll account from which a tweet was made and the time period from which tweets were sampled (e.g., primary, convention, and general elections). This approach made it possible to account for the individual troll generating the tweet and its proximity to the election. Negative binomial models were used because retweets are count data with the potential for many zeros (indicating no retweets). The R script used to conduct the analyses can be found on the OSF page for this project.

Results

Preliminary Analyses

We conducted a series of preliminary analyses to examine the prevalence of attacks across the four candidates. As illustrated in Table 2, the four presidential candidates were attacked in 4.7% and 42.3% of the tweets that addressed them, respectively. In terms of individual candidates, 42.3% of the tweets about Clinton included an attack. The other three candidates were attacked in less than 6% of the tweets.

Table 2. Proportion and Raw Number of Tweets About Each Candidate Containing an Attack.

	Clinton		Trump		Cruz		Sanders	
	Prop. attack	Total attacks	Prop. attack	Total attacks	Prop. attack	Total attacks	Prop. attack	Total attacks
Overall	0.423	881	0.056	145	0.118	64	0.047	39
Age	0.015	31	0.000	0	0.000	0	0.005	4
Appearance	0.001	3	0.002	5	0.000	0	0.000	0
Character/integrity	0.285	593	0.030	77	0.064	35	0.038	32
Competence	0.035	72	0.003	9	0.013	7	0.001	1
Gendered name-calling	0.004	8	0.000	0	0.002	1	0.000	0
General attacks	0.103	215	0.020	53	0.042	23	0.012	10
Sexual reference	0.010	20	0.005	12	0.006	3	0.000	0

Notes. Prop. = proportion. Proportion values reflect the number of tweets about a candidate containing an attack relative to all tweets that address a candidate.

Differences in Attacks and Attack Types

H1 predicted that frontrunners would be more likely to be attacked than challengers. RQ1 asked about differences in the likelihood of an attack on (a) Republican candidates compared with Democratic candidates and (b) male candidates compared with female candidates. The results of the pairwise comparison tests reported in Table 3 indicate differences between all three target groups. A significantly greater proportion of tweets about frontrunners included attacks compared with tweets about challengers. H1 was supported. A significantly greater proportion of tweets about Democrats included attacks relative to tweets about Republicans (RQ1a). A significantly greater proportion of tweets about the female candidate included attacks compared with tweets about male candidates (RQ1b).

RQ2 inquired about differences in the specific types of attacks directed at the target groups. The results reported in Table 3 indicate that tweets about frontrunners were more likely to include attacks on character/integrity, competence, general attacks, and sexual references than tweets about challengers (RQ2a). Democrats were more likely to be targeted with attacks on character/integrity, competence, and general attacks relative to tweets about Republicans (RQ2b). Although the number of such attacks was smaller, Democrats were also more likely to be the target of age attacks, gendered name-calling, and sexual references. The female candidate was more likely to be attacked based on character/integrity, competence, and general attacks than the male candidates (RQ2c). She was also significantly more likely to be attacked based on age and the recipient of gendered name-calling, but the absolute numbers of attacks for these variables were small.

Table 3. Difference in Attacks Between Frontrunners and Challengers, Republicans and Democrats, and Male and Female Candidates.

	Front runner			Challenger				
	Prop. attacks	Total attacks	Total tweets	Prop. attacks	Total attacks	Total tweets	*z*	*p*
Overall	0.261	1,003	3,850	0.075	102	1,362	208.19	**<.001**
Age	0.008	31	3,850	0.003	4	1,362	3.22	.073
Appearance	0.002	8	3,850	0.000	0	1,362	1.64	.200
Character/integrity	0.170	655	3,850	0.048	66	1,362	123.93	**<.001**
Competence	0.021	81	3,850	0.006	8	1,362	12.90	**<.001**
Gendered name-calling	0.002	8	3,850	0.001	1	1,362	0.42	.518
General attacks	0.067	257	3,850	0.023	32	1,362	35.13	**<.001**
Sexual reference	0.008	30	3,850	0.002	3	1,362	4.15	**.042**
	Republicans			Democrats				
	Prop. attacks	Total attacks	Total tweets	Prop. attacks	Total attacks	Total tweets	*z*	*p*
Overall	0.072	206	2,878	0.358	912	2,545	662.18	**<.001**
Age	0.000	0	2,878	0.013	34	2,545	38.42	**<.001**
Appearance	0.002	5	2,878	0.001	3	2,545	0.00	1.000
Character/integrity	0.038	110	2,878	0.244	621	2,545	702.84	**<.001**
Competence	0.006	16	2,878	0.029	73	2,545	73.79	**<.001**
Gendered name-calling	0.000	1	2,878	0.003	8	2,545	8.40	**.004**
General attacks	0.026	74	2,878	0.086	218	2,545	171.10	**<.001**
Sexual reference	0.005	15	2,878	0.008	20	2,545	5.43	**.020**

(table continues)

	Male candidates			Female candidate				
	Prop. attacks	Total attacks	Total tweets	Prop. Attacks	Total attacks	Total tweets	*z*	*p*
Overall	0.068	242	3,581	0.423	881	2,083	1,014.43	**<.001**
Age	0.001	4	3,581	0.015	31	2,083	36.58	**<.001**
Appearance	0.001	5	3,581	0.001	3	2,083	0.03	.857
Character/integrity	0.039	140	3,581	0.285	593	2,083	488.67	**<.001**
Competence	0.005	17	3,581	0.035	72	2,083	43.32	**<.001**
Gendered name-calling	0.000	1	3,581	0.004	8	2,083	4.80	**.029**
General attacks	0.023	81	3,581	0.103	215	2,083	94.10	**<.001**
Sexual reference	0.004	15	3,581	0.010	20	2,083	1.09	.296

Note. Prop = proportion.

Audience Engagement With Candidate Attacks

H2 predicted that attacks on frontrunners would receive greater engagement compared with attacks on challengers. RQ3 questioned differences in audience engagement about attacks on (a) Republican candidates relative to Democrats and (b) male candidates compared with female candidates. Because the different attack target groups were correlated, separate multilevel models were conducted to evaluate each pair of groups. To evaluate the unique implications of attacks, dummy variables were constructed for each of the predictors to compare tweets relating to groups that were attacked with tweets and those that were not (e.g., tweets when frontrunners were attacked and when challengers were attacked relative to tweets where neither group was attacked). This approach made it possible to examine the outcomes of attacks on retweets relative to the baseline of tweets made about the two groups that did not contain an attack. The results of the model can be found in Table 4.

Table 4. Predicting Differences Between Groups in Retweets of Attack Tweets.

	Retweets			Retweets			Retweets		
	IRR	*CI*	*p*	IRR	*CI*	*P*	IRR	*CI*	*p*
(Intercept)	0.94	0.59–1.48	.776	0.094	0.59–1.48	.780	0.94	0.59–1.48	.778
Challenger attacked	0.92	0.68–1.24	.577						
Frontrunner attacked	1.19	1.06–1.33	**.002**						
No attack (reference group)									
Democrat attacked				1.22	1.09–1.38	**.001**			
Republican attacked				0.87	0.68–1.10	.251			
No attack (reference group)									
Female attacked							1.22	1.09–1.38	**.001**
Male attacked							0.93	0.75–1.15	.498
No attack (reference group)									
Random Effects									
ICC		0.80			0.80			0.80	
Marginal R^2/conditional R^2		0.001/0.796			0.001/0.801			0.001/0.801	

Notes. Outcome variable = retweet count. IRR = incident rate ratio. CI = 95% confidence interval. ICC = intraclass correlation. The outcome variable is the number of times a tweet has been retweeted. Random intercepts were included for individual troll accounts and time periods (i.e., primary, convention, general election). Marginal R^2 only accounts for the fixed effects; Conditional R^2 accounts for the fixed and random effects.

Tweets containing attacks on frontrunners generated significantly more retweets than tweets without an attack. The incident rate ratio (IRR) indicated that tweets attacking frontrunners yielded 1.19 times the number of retweets as tweets without an attack. There was no difference in retweets containing attacks on challengers compared with tweets with no attacks. These results are consistent with H2.

About RQ3, tweets attacking Democrats yielded significantly more retweets than tweets about Democrats and Republicans that did not contain an attack, but tweets attacking Republicans did not differ in retweets from tweets that lacked an attack. Tweets attacking Democrats generated 1.22 times the number of retweets compared with those in which neither group was attacked (RQ3a). Finally, tweets attacking the female candidate received significantly more retweets than tweets that did not attack the male or female candidates, but tweets attacking the three male candidates did not differ from tweets without an attack. Tweets attacking Clinton generated 1.22 times the number of retweets as tweets that did not attack either group of candidates (RQ3b).

Discussion

The purpose of our research was to further contribute to the literature about the role that IRA trolls played on Twitter in 2016. We investigated how IRA trolls' attacks varied based on candidate characteristics and which attacks received the most audience engagement. Readers should note that our goal was not to generalize to all candidates across all elections but to better understand how IRA trolls attacked presidential candidates on social media during the 2016 election. Our results provide three general findings.

First, we found that troll tweets about frontrunners (i.e., Trump and Clinton) contained more attacks than tweets about challengers (i.e., Cruz and Sanders). The most prevalent types of attacks against frontrunners included attacks based on character/integrity, competence, general attacks, and sexual references. Tweets attacking frontrunners also garnered a greater number of retweets relative to tweets that did not attack either frontrunners or challengers. There was no difference in retweets between tweets attacking challengers and tweets without an attack. Second, tweets about Democrats (i.e., Clinton and Sanders) contained more attacks than tweets about Republicans (i.e., Trump and Cruz). Specific attacks that were more prevalent against Democrats included attacks based on character/integrity, competence, general attacks, age, gendered name-calling, and sexual references. Tweets attacking Democratic candidates garnered more retweets than tweets about Democrats and Republicans without an attack. There was no difference in retweets between tweets attacking Republicans and tweets without an attack. Finally, tweets about Clinton, the only female candidate, contained more attacks than tweets about the male candidates. She received the most age-based and gendered name-calling attacks. Additionally, tweets attacking Clinton received more retweets than tweets about male or female candidates without an attack. There was no difference in retweets between tweets attacking male candidates and tweets in which neither group was attacked.

Because of the accessibility of social media, changes in the political landscape have allowed new actors to shape political discourse and voting behaviors (Kenski & Jamieson, 2017). The IRA trolls acted as a group without national standing to set forth specific narratives about Clinton, Trump, Sanders, and Cruz. They were primarily concerned with attacking frontrunners, Democrats, and the female candidate—all of which ultimately encompass Clinton as the main target. The patterns in our data indicate that trolls wanted to create a narrative that framed Clinton as an incompetent candidate who lacked character compared with her opponents. Our results suggest that the audience responded favorably to the trolls' attacks against Clinton because they received more audience engagement than tweets lacking such attacks. As traditional sources (e.g., news media, ads, and candidates) have established how the public thinks about candidates through specific attacks, our study suggests that online trolls offer a new source for attacks that can perhaps be predictive of political outcomes. Below, we unpack each of these main findings in greater detail.

Supporting past research analyzing traditional (Haynes & Rhine, 1998) and social media (Stein & Benoit, 2021), frontrunners Trump and Clinton received more attacks than trailing candidates Sanders and Cruz from IRA Twitter accounts. Of all tweets mentioning frontrunners, 26% were attacks. In contrast, of all tweets mentioning challenging candidates, just 7.5% were attacks. Our results show that the longstanding strategy of attacking frontrunners in traditional media extends to social media trolls. When challenging candidates attack frontrunners, they face a potential backlash from voters who disapprove of their behavior, sometimes indirectly resulting in candidacy withdrawal (Hinck, Hinck, Dailey, & Hinck, 2013). Trolls, however, can evade such penalties and continue attacks for the duration of the election.

Frontrunners received more attacks than challenging candidates in the areas of character/integrity, competence, general attacks, and sexual references. Character/integrity attacks (e.g., "The greatest political accomplishment of Hillary is to avoid prosecution. #NeverHillary"; personal communication, June 23, 2016) by far made up the greatest proportion of attacks (17%) on frontrunners; competence attacks made up the second greatest proportion of attacks (2.1%; e.g., "how to win an election in #USA #ThingsHillaryGoogles"; personal communication, September 21, 2016). General attacks usually occurred in the form of hashtags (e.g., #NeverHillary, #NeverTrump) or profanity against frontrunners. In future elections, not only do frontrunners need to be concerned about the level of attacks launched by all other trailing candidates, but they must also fight the potential fire of nefarious nonstate actors online. Trolling efforts could cause future campaigns to spend exponentially increased resources addressing slights on candidate image rather than explicating policy. These adjustments in frontrunners' campaign strategies, driven by trolls' excessive attacks and the audience engagement they generate, may also reinforce media attention and constituent interest in candidate traits, leaving even less room to discuss policy issues that directly affect citizens.

In terms of audience engagement, troll tweets attacking frontrunners retweeted significantly more than tweets about frontrunners and challengers that did not contain an attack. Hemsley (2019) confirmed that among tweets emanating from political candidates, attack messages generally receive more retweets than advocating tweets. To our knowledge, retweets of tweets attacking frontrunners in comparison with challenging candidates have not been considered in past research, making our findings a unique contribution to not only how trolls attack frontrunners but also how troll followers, who may largely be regular Twitter users, engage with attack tweets.

Additionally, the IRA trolls attacked Democratic candidates more often than Republican candidates on all fronts, but particularly character/integrity (e.g., "Dems are racists, they divide people so they can easily control them. And Bernie Sanders wants to revive Marx also"; personal communication, March 7, 2016), competence (e.g., "Hillary can't hire professionals even for her own campaign. Think she'll somehow be OK for presidency? #Unqualified"; personal communication, May 25, 2016), and general attacks (e.g., "I'm sick of libtards but Hillary is even worse! She is pure evil!"; personal communication, June 20, 2016). Some of these tweets contained simultaneous attacks against both competence and character in the same phrases (e.g., "Raging alcoholic, not fit for president #MAGA #HillaryForPrison2016 #Election2016"; personal communication, November 8, 2016).

Our findings align with party perceptions. Partisan bias against the other party's perceived strengths may allow for particular attention to these characteristics, such as character weaknesses, in Democratic candidates (Goren, 2002). Further, because Republicans are seen as strong leaders, the IRA trolls might have attacked what they saw as a lack of strong leadership in the form of competence among Democratic candidates.

Not only were these attacks against Democratic candidates more frequent, but they were also more retweeted than tweets without an attack. The consequences of this assault against Democratic candidates, compared with Republican candidates, may have increased visibility and primacy for thoughts on Democratic weaknesses compared with Republican ones. The IRA trolls attacked candidates on traits that have been stereotypically used to critique parties as a whole. This presented the trolls with the opportunity to understand and exploit those traits that were most likely to amplify partisan bias and infighting. Although the trolls appeared to prefer the Republican candidates, who were attacked less frequently, it is unclear whether this preference ought to be attributed to their preference for Republican policies generally or their desire to sow discord by undermining the candidate perceived at the outset of the election to have the best chances of winning and happened to be the Democrat. These attacks may have implications for future elections, with other bad actors attempting to sow discord or elevate preferred candidates by disproportionally attacking candidates based on party traits and stereotypes. Doing so may place these perceived weaknesses, particularly in competence and character, at the forefront of the public's mind, which may then impact voting behaviors when the time comes.

In addition to increased attacks against Democratic candidates in general, Clinton was the only woman and frontrunner in the party. She received more attacks than the male candidates, and tweets attacking Clinton were retweeted more than tweets without an attack. Intentional or not, the IRA trolls capitalized on hostile sexism as they attacked Clinton. To illustrate, although the appearance category was relatively rare, most of these tweets were aimed at Clinton (e.g., "SHOCKING: Leaked photo of Hillary Clinton without makeup #Mstrumprally"; personal communication, August 24, 2016), thus conforming to the sexist ideology that women should be evaluated by what they look like rather than what they can do. Although Clinton had extensive political experience, she received the most competence and character attacks indicating that she was "unfit" to be president (e.g., "Keep #CrookedHillary out of the wh! #HillarysUnfit to be elected dog catcher. #MAGA #TrumpPence16 #ImNotWithHer"; personal communication, August 12, 2016). Some tweets even described Clinton as unfit because of the emotional trope associated with being a woman (e.g., "Just a little reminder—Trump is telling the truth. Hillary speaks emotionall stuff which people wnant to hear"; personal communication, February 22, 2016). Although Sanders is six years older than Clinton, she received more attacks for her age (e.g., "@HillaryClinton How the hell a woman who can't climb a few steps & struggles to open a loosened jar can be our LEADER?!"; personal communication, August 23, 2016) than her male counterparts.

The candidate gender findings suggest that IRA trolls may not have felt that the cultural climate of 2016 welcomed explicit sexist attacks. Instead, trolls used more subtle attacks based on character, competence, and age. These indirect sexist attacks would have been particularly potent, as Trump himself used sexist language during the 2016 campaign to provoke voters (Banwart & Kearney, 2018). The attacks against Clinton did not include significant levels of gendered name-calling, an overt form of sexism. This permitted trolls, and indeed general Twitter users, to hide behind denials of sexism. Nevertheless, covert sexism is still sexism. Further, the data collected from the IRA trolls have implications for not only Clinton's electoral loss but also the future of women in politics. For Clinton, the IRA trolls played on the power of covert sexism to sow doubt about her candidacy. By doing so, they may have placed future female presidential candidates at a disadvantage, and Clinton will become the primary comparison. This has broader implications for the IRA trolls' actions beyond the 2016 election.

Limitations & Future Directions

We acknowledge the three main limitations of our study. First, we analyzed only tweets during a single election. This, of course, is attributed to the novel role that the IRA trolls played in the 2016 U.S. election. Because the trolls only dedicated attention to leading candidates, our analyses necessarily had to be limited in scope to the four elected officials we examined. Second, we coded and analyzed tweets with an overlap in categories. It was possible for multiple categories of attacks to appear in a single tweet. Third, Trump and Clinton had distinct features that could have impacted the outcomes of our results. Gender is especially salient to expand beyond the 2016 election because Clinton has been the only female general election presidential candidate. This case study, however, is important for understanding the environment that future female candidates may face.

Despite these caveats, our study offers possibilities for future research. Although our study examined engagement via retweets, we did not examine replies to tweets. Future research should consider how Twitter users reply to tweets containing attacks and how replies vary per candidate. Qualitative analysis of such tweets could offer further insight into how the public processes and responds to attacks. It would also be worthwhile to expand the analyses to other state-sponsored campaigns on social media. Because our results are limited to the IRA and the 2016 U.S. presidential election, it is unclear whether the trends we observed are unique to this group and contest or represent broader patterns in the way state-sponsored trolls engage in election interference. Future research should investigate female candidates as targets of foreign interference, in particular, and how sexism takes place in an online environment.

Conclusion

The IRA trolls made a concerted effort to interfere in the 2016 U.S. presidential election. This study provides a nuanced understanding of who they were attacking, the nature of their attacks, and the degree to which their attacks resonated with Twitter users. The IRA trolls were the most likely to attack frontrunners, Democrats, and the female candidate. Such results may extend beyond the 2016 election and—particularly about the gender-specific findings—highlight the progress society has yet to make. We hope that a better understanding of troll behavior in 2016 will leave us more informed about state-sponsored attacks on our electoral process.

References

Badawy, A., Addawood, A., Lerman, K., & Ferrara, E. (2019). Characterizing the 2016 Russian IRA influence campaign. *Social Networking Analysis and Mining, 9*(1), 1–12. doi:10.1007/s13278-019-0578-6

Banwart, M. C., & Kearney, M. W. (2018). Social dominance, sexism, and the lasting effects on political communication from the 2016 election. In B. R. Warner, D. G. Bystrom, M. S. McKinney, & M. C. Banwart (Eds.), *Democracy disrupted: Communication in the volatile 2020 presidential election* (pp. 419–440). Santa Barbara, CA: Praeger.

Bartels, L. M. (2002). The impact of candidate traits in American presidential elections. In A. King (Ed.), *Leaders' personalities and the outcomes of democratic elections* (pp. 44–69). Oxford, NY: Oxford University Press.

Dunway, J., & Graber, D. A. (2023). *Mass media and American politics* (11th ed.). Washington, DC: CQ Press.

Dutta, U., Hanscom, R., Zhang, J. S., Han, R., Lehman, T., Lv, Q., & Mishra, S. (2021). Analyzing Twitter users' behavior before and after contact by the Russia's Internet Research Agency. *Proceedings of the ACM on Human-Computer Interaction, 5*(CSCW1), 1–24. doi:10.1145/3449164

Falk, E. (2018). *Women for president: Media bias in eight campaigns*. Champaign: University of Illinois Press.

FiveThirtyEight. (2016). *2016 national primary polls*. Retrieved from https://projects.fivethirtyeight.com/election-2016/national-primary-polls/democratic/

Fridkin, K. L., & Kenney, P. J. (2011). The role of candidate traits in campaigns. *The Journal of Politics, 73*(1), 61–73. doi:10.1017/S0022381610000861

Goren, P. (2002). Character weakness, partisan bias, and presidential evaluation. *American Journal of Political Science, 46*(3), 627–641. doi:10.2307/3088404

Gross, J. H., & Johnson, K. T. (2016). Twitter taunts and tirades: Negative campaigning in the age of Trump. *PS: Political Science & Politics, 49*(4), 748–754. doi:10.1017/S1049096516001700

Gwet, K. (2018). *An evaluation of the impact of design on the analysis of nominal-scale inter-rater reliability studies*. AgreeStat Analytics. Retrieved from https://www.agreestat.com/papers/inter-rater%20reliability%20study%20design1.pdf

Hayes, D. (2005). Candidate qualities through a partisan lens: A theory of trait ownership. *American Journal of Political Science, 49*(4), 908–923. doi:10.1111/j.1540-5907.2005.00163.x

Hayes, D. (2011). When gender and party collide: Stereotyping in candidate trait attribution. *Politics and Gender, 7*(2), 133–165. doi:10.1017/S1743923X11000055

Haynes, A. A., & Rhine, S. L. (1998). Attack politics in presidential nomination campaigns: An examination of the frequency and determinants of intermediated negative messages against opponents. *Political Research Quarterly, 51*(3), 691–721. doi:10.1177/106591299805100307

Hemsley, J. (2019). Followers retweet! The influence of middle-level gatekeepers on the spread of political information on Twitter. *Policy & Internet, 11*(3), 280–304. doi:10.1002/poi3.202

Hinck, S. S., Hinck, R. S., Dailey, W. O., & Hinck, E. A. (2013). Thou shalt not speak ill of any fellow Republicans? Politeness theory in the 2012 Republican primary debates. *Argumentation and Advocacy, 49*(4), 259–275. doi:10.1080/00028533.2013.11821801

Jamieson, K. H. (1995). *Beyond the double bind: Women and leadership*. Oxford, NY: Oxford University Press.

Jamieson, K. H. (2018). *Cyberwar: How Russian hackers and trolls helped elect a president: What we don't, can't, and do know*. Oxford, NY: Oxford University Press.

Kenski, K., Hardy, B. W., & Jamieson, K. H. (2010). *The Obama victory: How media, money, and message shaped the 2008 election*. Oxford, NY: Oxford University Press.

Kenski, K., & Jamieson, K. H. (2017). Political communication: Looking ahead. In K. Kenski & K. H. Jamieson (Eds.), *The Oxford handbook of political communication* (pp. 913–918). Oxford, NY: Oxford University Press.

Linvill, D. L., Boatwright, B. C., Grant, W. J., & Warren, P. L. (2019). "The Russians are hacking my brain!" Investigating Russia's Internet Research Agency Twitter tactics during the 2016 United States presidential campaign. *Computers in Human Behavior, 99*, 292–300. doi:10.1016/j.chb.2019.05.027

Okimoto, T. G., & Brescoll, V. L. (2010). The price of power: Power seeking and backlash against female politicians. *Personality and Social Psychology Bulletin, 36*(7), 923–936. doi:10.1177/0146167210371949

Pfau, M., & Kenski, H. C. (1990). *Attack politics: Strategy and defense*. Westport, CT: Praeger.

Pfau, M., Parrott, R., & Lindquist, B. (1992). An expectancy theory explanation of the effectiveness of political attack television spots: A case study. *Journal of Applied Communication Research, 20*(3), 235–253. doi:10.1080/00909889209365334

Rains, S. A., Harwood, J., Shmargad, Y., Kenski, K., Coe, K., & Bethard, S. (2023). Engagement with partisan Russian troll tweets during the 2016 U.S. presidential election: A social identity perspective. *Journal of Communication, 73*(1), 38–48. doi:10.1093/joc/jqac037

Rains, S. A., Shmargad, Y., Coe, K., Kenski, K., & Bethard, S. (2021). Assessing the Russian troll efforts to sow discord on Twitter during the 2016 US election. *Human Communication Research, 47*(4), 477–486. doi:10.1093/hcr/hqab009

Real Clear Politics. (2016). *2016 republican presidential nomination*. Retrieved from https://www.realclearpolitics.com/epolls/2016/president/us/2016_republican_presidential_nomination-3823.html

Rossini, P., Stromer-Galley, J., Kenski, K., Hemsley, J., Zhang, F., & Dobreski, B. (2018). The relationship between race competitiveness, standing in the polls, and social media communication strategies during the 2014 U.S. gubernatorial campaigns. *Journal of Information Technology & Politics, 15*(3), 245–261. doi:10.1080/19331681.2018.1485606

Ruck, D. J., Rice, N. M., Borycz, J., & Bentley, R. A. (2019). Internet Research Agency Twitter activity predicted 2016 US election polls. *First Monday, 24*(7). doi:10.5210/fm.v24i7.10107

Starbird, K. (2019). Disinformation's spread: Bots, trolls and all of us. *Nature, 571*(7766), 449–450. Retrieved from https://media.nature.com/original/magazine-assets/d41586-019-02235-x/d41586-019-02235-x.pdf

Steger, W. P. (1999). Comparing news and editorial coverage of the 1996 presidential nominating campaign. *Presidential Studies Quarterly, 29*(1), 40–64. Retrieved from https://www.jstor.org/stable/27551958

Stein, K. A., & Benoit, W. L. (2021). A functional analysis of 2016 nonpresidential campaign tweets. *American Behavioral Scientist, 65*(3), 432–447. doi:10.1177/0002764221996770

Strauss, A., & Corbin, J. (1990). *Basics of qualitative research: Grounded theory procedure and techniques*. Thousand Oaks, CA: Sage Publications.

Uyheng, J., Moffitt, J. D., & Carley, K. M. (2022). The language and targets of online trolling: A psycholinguistic approach for social cybersecurity. *Information Processing & Management, 59*(5), 103011–103012. doi:10.1016/j.ipm.2022.103012