Gender and Twitter Behavior in the Russian Disinformation Campaign in the United States, 2012-2018

Andrew W. Pierce, PhD

June 23, 2020

Abstract

One of the major goals of disinformation campaigns is to ignite negative emotions in citizens that might exacerbate tension in a target country. Evoking anger is a common goal of disinformation campaigns, and campaigns can exploit gender norms as a possible cleavage in society. In this paper, we theorize how gendered disinformation campaigns use gender signals to promote their objectives. We answer the question, is gendered differentiated in disinformation campaigns, by analyzing over 850,000 tweets released by Twitter as part of the investigation into the activities of the Russian Internet Research Agency (IRA). We find that gendered patterns of disinformation do exist both in the levels of Twitter activity but also in the types of content released. These results suggest that understanding how gender norms can be exploited is an important component of combating disinformation campaigns.

Introduction

With the advent of social media, governments, both foreign and domestic, have had increasing access to the public when it comes to the distribution of information. Since at least 2012, Russian agents have been posting content via Twitter in an effort to undermine governments abroad (Linvill and Warren 2020). Additionally, the US has seen increasing partisan polarization that frequently conflicts over issues of identity politics (Mason 2018). How then, might identity politics influence disinformation campaigns in the United States?

In this paper, we seek to understand specifically how gender differences manifest in disinformation campaigns. Using Twitter data released as part of an investigation into interference in the 2016 US Presidential election, we provide evidence that gender differences manifest in the behavior of disinformation campaigns. Drawing on a mixed methods approach, we use Natural Language Processing and common statistical tests to suggest that the Internet Research Agency (IRA) intentionally manipulated the behavior of their Twitter accounts to reflect gender differences. From this, we conclude that gender and other dimensions of identity are important for understanding the behaviors of disinformation campaigns as well as how to combat them.

We proceed as follows. First, we offer reasoning for why gender might be an important dimension to exploit when trying to achieve the objectives of a disinformation campaign. Next, we detail how we assigned gender to Twitter accounts attached to the IRA and identified a population of tweets to study. After detailing our methodology, we explore how male- and female-assigned accounts behave more or less differently or similarly in terms of activity, timing, and content. We conclude by acknowledging the limitations of this study and suggesting avenues for further research.

How Would Gender Play a Role in Disinformation Campaigns?

There are many different ways to conceptualize gender, especially as it relates to biological sex or sexual orientation. However, scholars have for many decades considered the social implications of gender and how gender shapes social behavior. In one influential article, West and Zimmerman (1987) explain how gender can be conceptualized as rules for behavior. As evidence, they cite studies of how children look to gender norms when deciding how to act as well as how transgender individuals must learn news ways of behaving. This conceptualization, where gender is a prescription for behavior, also implies that there are punishments for not acting according to this prescription.

Importantly, noncompliance with norms can potential lead to social sanctioning depending on the social environment and composition of social groups. Furthermore, noncompliance with norms can lead to emotional arousing among comforming individual or groups. Typically, one of the strongest reactions to noncompliance is anger (Carpenter and Matthews 2012). Thus, presenting tweets in a gendered way may be an effort to elicit emotional responses from the viewers of the tweets. Indeed, inciting anger has been a primary aim of disinformation campaigns in the United States (Tucker et al. 2018).

Anger and feelings of anger has recently be tied to the erosion of democratic norms and trust between groups. Weeks (2015), for example, has shown how anger can lead to increasingly partisan evaluations of information. Additionally, Webster (2020) has shown how anger leads to decreasing trust in the ability of the government to represent one's interests and discourages participation. Finally, Bail (2016) shows how emotional responses lead to higher levels of "viral" content on social media.

So why would disinformation agents take a gendered approach to their campaigns? By taking a gendered approach, disinformation agents may be able to elicit emotional responses or connect to group identities of consumers. Promoting tweets that are gender compliant may increase levels of trust in the material shared. Perhaps more importantly, gender norms can be challenged or reinforced in such a way as to provoke angry responses in the audience. This would lead to more shares and greater spread of disinformation, but would also undermine trust in institutions and fellow citizens. In other words, gender, by way of emotional evocation, can advance two objectives of disinformation campaigns: undermining trust and increasing virality of disinformation.

Is there Evidence of Gender Patterns in Disinformation Campaigns?

Given that taking a gendered approach to disinformation would advance the goals of disinformation campaigns, what evidence exists that supports such an approach is actually employed? To answer this question, we use a publicly available set of tweets attributed to the Russian Internet Research Agency(IRA) as part of a campaign to influence the 2016 US Presidential election.

How To Analyze Gender Differences in IRA Disinformation Campaign?

In order to understand any potential manipulation of gender cues, we must first explore what differences in Twitter account behavior exist between genders if at all. Doing so requires first classifying accounts by gender and then examining differences between the behaviors of those accounts. Accordingly, analysis will proceed by first classifying the gender of accounts, performing exploratory data analyses of qualitative differences between account types, and then testing for meaningful quantitative differences between account types.

Data Source

To better understand the use of gender signaling in a disinformation campaign, we will be examining Twitter posts from Twitter identified as being maintained by the Russian Intelligence Agency (IRA)[1]. These tweets were released by Twitter as a result

[1]This data was made available through the GitHub account for 538.com. The original data was provided to 538 by authors Linvill and Warren (2020).

of counterintelligence operations in the US who identified accounts used by Russian intelligence to sow discord in the United States and elsewhere. While the entirety of the data totals over 9 million tweets, we will focus our analysis on English tweets in the US. Limiting the scope of this analysis to English tweets in the US holds constant cultural and political factors that vary by country and might confound any gender differences between countries[2]. Of the 9.04 million tweets released by Twitter, roughly 1.84 million represent English tweets in the US. Of the 3,841 accounts in the original data set, 1,264 remain for analysis after subsetting.

Assigning Gender to Twitter Handles

Coding gender from social media is a rich and diverse field of study with many different approaches, both simple and complex. For the task at hand, we make several simplifying assumptions about the coding of gender. We have erred on the side of accuracy in identifying gender over inclusion whenever possible, while acknowledging these decisions lead to limitations in analysis.

In order to code gender, we first employ Social Security Administration (SSA) data from 1932 to 2012, which contains name frequency broken down by gender for births from 1932 to 2012. Over 90,000 names exist within this dataset, but not all are necessarily common enough for an audience to identify. Thus, we are limiting names to those that occur more than 5000 times within the data set, meaning at least 5000 people born between 1932 and 2012 had that name. This limits the pool to 3498 names, but covers over 77.5% of the total population.

Next, for each unique name in the SSA subset, we identify all Twitter handles in which that name is wholly present. This flavor of exact matching catches fewer matches than a fuzzy match because many real-world names are translated when included in handles, perhaps being abbreviated ("Michael" to "Mchl") or substituting numbers for letters ("Michael" to "M1cha3l"). There are also cases where names may be subsets of other names, such as "Drew" being a subset of "Andrew." Because matches were made

[2]Future analyses might very well leverage cross-cultural differences in order to compare and contrast how gendered language and cues shape disinformation campaigns

if a name was contained anywhere in the handle, multiple names may be matched to the same handle. When this occurred, we chose the match with the longest[3] number of characters. So, for example, @1d_nicole_ contains the names "Cole," "Nicole," and "Nico." While all three names were matched to the handle, for gender classification purposes, "Nicole" would be used because it is name with the longest number of characters. Since some short names, such as "Al" tend to be very common in other names, we limit our matching to names with character length 4 or longer.[4]

After the names are assigned to the data, gender is assigned according the total proportion of a name that belongs to each gender in the data set. If a name is a certain gender more than 50% of the time, then that name is assigned that gender. This works very well in cases where a name is much more frequently assigned to one gender than another. For example, 99.6% of children named "James" were male and only 0.4% were female. However, this method is much less successful for more gender-neutral names, such as "Jessie" (49.9% male, 50.1% female). While this method lacks the nuance of gender that exists in the real world, it does allow for classifying the most likely gender impression that a person would have when reading a tweet.[5]

Of the 1,264 accounts that tweeted in English in the United States, our matching process was able to assign gender to 862 accounts. This represents 68% coverage of the handles included in the original subset. Of the 1.84 million tweets included in the original subset, 859,029 can be attached to a gendered name, which represents 46.5% coverage of the original tweets.

[3]Ties were broken alphabetically.

[4]We also filter out common news names, such as "online" and "today" which appear in handles such as @oaklandonline or @houstontoday. A subset of 100 handles were checked for accuracy with a 71% accuracy rate.

[5]This categorization process has the benefits of being simple and easily implemented, but admittedly lacks sophistication with respect to the way gender works in the real world. Gender, for most behaviors, acts more like a spectrum of participation than a dichotomy. However, absent sophisticated tools allowing for categorization of an accounts behaviors as more or less feminine/masculine, this cruder measure will have to suffice.

Behavioral Differences Between Accounts Assigned to Different Genders

If gender does indeed play an important role in disinformation campaigns, the lowest level of evidence required to support this claim would be simple difference in tweet activity. Considering the null hypothesis, that disinformation agents ignore gender, we would expect that tweet activity is assigned randomly to accounts. Thus, the first step in supporting the hypothesis that gender is dimension exploited by disinformation agents would be to establish differences between tweet activity between gender-assigned accounts.

Having assigned accounts to different genders, there are at first review meaningful differences in the activity levels between groups. Following the classification procedure defined above, there is relative balance in the number of accounts assigned to each gender; 401 accounts are classified as female names (48.2%) and 431 accounts are classified as male names (51.8%). However, there exists a larger imbalance in the overall volume of tweets. Of the roughly 859,000 tweets, 500,326 (58.2%) come from accounts classified as male and only 358,703 (42.8%) come from accounts classified as female. This suggests that attempts at spreading disinformation were more likely to come from male-assigned accounts, but also that there is gender dimension with respect to disinformation activity.

While the evidence suggests that male-assigned accounts are more active than the female-assigned accounts in aggregate, it may be the case that the distribution of activity between groups is similar save for a few hyperactive accounts among the male-assigned accounts that do not exist among the female assigned accounts. However, considering the distributions in Table 1, the difference in activity between male- and female-assigned accounts does not seem to be driven by outliers. Indeed, the account with the largest volume of tweets, @AMELIEBALDWIN, is actually assigned female. Furthermore, when looking at the distributions, the 25th percentiles for both gender groups is relatively similar. However, the 75th percentile is much larger for me. This suggests that much of the difference in tweet activities between gender is driven by a number of more active male-assigned accounts and not a few outliers.

Summary Statistics of Tweet Activity by Gender						
Account Type	Min	25th Percentile	Median	Mean	75th Percentile	Max
Female-Assigned	1	19	69	948	211	34307
Male-Assigned	1	24	79	1112	897	14833

Table 1: **Male-Assigned Accounts Tweet More** In looking at the summary statistics for tweet behavior, male-assigned accounts do tweet more on average, and this is driven by extra activity by high frequency accounts. The account with the most tweets was actually a female-assigned account, but the overall average for male-assigned accounts is still higher.

One way to quantify the level of similarity between the distribution of tweet activity between genders is using the Overlapping Coefficient. This measure is calculated as the area of overlap between the two distributions, $\int min(f(x), g(x))dx$, where $f(x)$and $g(x)$ represent the two distributions. Theoretically, this ranges from 0 to 1, with 1 representing perfect overlap. The overlapping coefficient for the distributions presented in Figure 1 is 0.518, which suggests a balance of similarity and difference between the two distributions. Reviewing the visualized distributions, both groups contain clusters of high-activity and low-activity accounts, but the high-activity account is more dense for male-assigned accounts. This supports the conclusion that male-assigned accounts were, in aggregate, more active that female-assigned accounts.

Gender Differences in Timing of Tweets

Broadly speaking, accounts of both genders tweeted similar amounts on similar days. Considering Figure 2, the number of tweets from both types of accounts seem to follow a similar pattern. There is relatively little tweet activity until 2015, which sees a increase during the 2016 US Presidential campaign. The largest spike occurs right before the November election, and there are higher levels of activity until Summer 2017, when there is a large drop. There is brief spike in July 2017, but roughly no activity afterwards.

Additionally, the volume of tweets is highly correlated between male- and female-

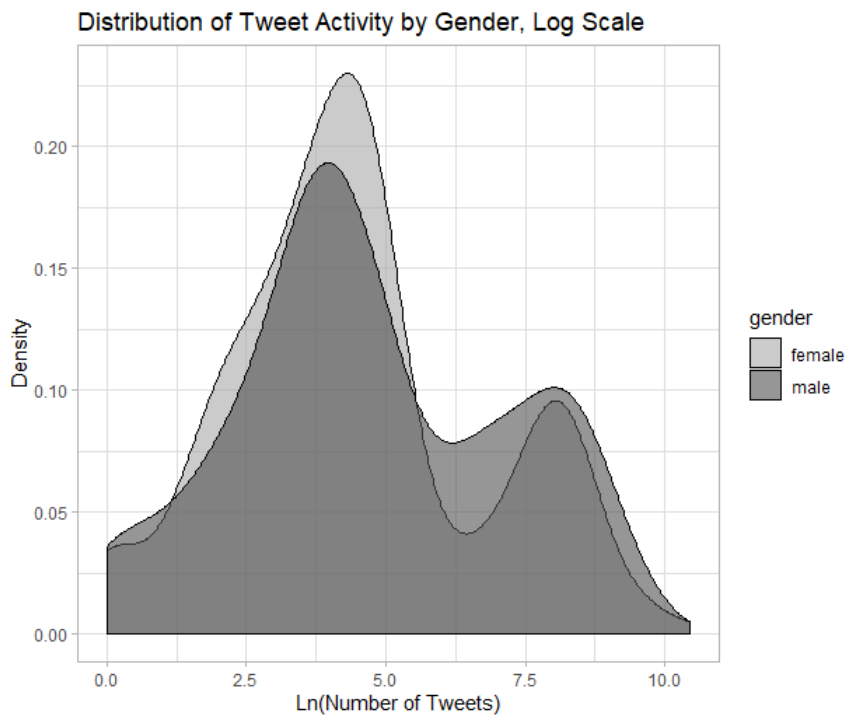

Figure 1: **Male-Assigned Accounts Tweet More** Comparing the distribution of tweet activity, both distributions appear bimodal. However, the higher-activity mode is more dense for male-assigned accounts than for female-assigned accounts

assigned accounts. The volume of tweet correlates a level of $\rho = 0.845$[6], which suggests that the timing of Twitter activity corresponds highly between male- and female-assigned accounts.

Given that both types of accounts tend to be active at the same time, might it be case that activity by one group of accounts precedes or predicts activity by the other group? In order to test this theory, we apply a series of Granger Causality Tests with lags of 3 days and one week. As can be seen in Table 2, in both cases activity in both male- and female- assigned accounts Granger causes activity in the other group, which suggests that activity in one group do not unilaterally predict activity in the other group. Rather, activity between each group seems to be high clustered together.

[6]Statistically significant at $\alpha \leq 0.001$ level.

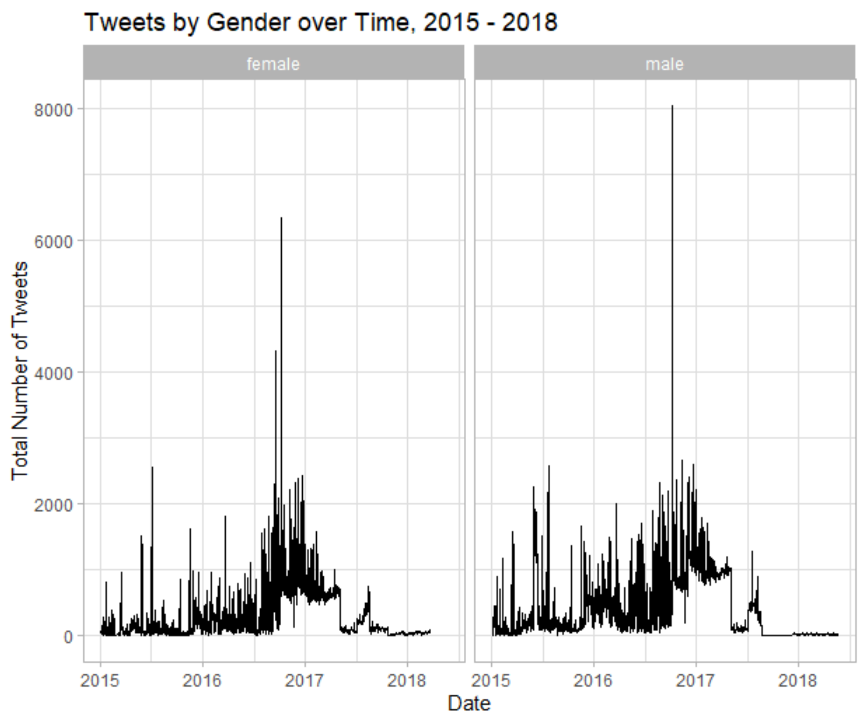

Figure 2: **Tweets Over Time by Gender** The Twitter activities of both types of account tend to be active at the same time. Both groups saw increases in the lead up the 2016 US Presidential election and decreases after the election.

	Granger Causality Tests for Predictive Activity by Gender		
Lag	Hypothesis	F-Statistic	p-Value
3	Male Activity Predicts Female Activity	12.57***	< 0.001
3	Female Activity Predicts Male Activity	36.85***	< 0.001
7	Male Activity Predicts Female Activity	20.05***	< 0.001
7	Female Activity Predicts Male Activity	16.05***	< 0.001

*** signifies statistical significance at $\alpha \leq 0.001$ level.

Table 2: **Accounts Assigned to Each Gender Granger Cause Each Other** Conducting Granger causality tests hypothesizing that activity by accounts to one gender predict the activity in accounts assigned to the other gender suggests that activity is clustered among genders. Activity by accounts assigned to either gender does not unilaterally predict activity in the other gender.

Gender Differences in Tweet Content

While there is evidence that there are gender differences in tweet activity, there may also be significant difference in the actual substance of the tweet as well. To the extent that gender difference matter in the presentation of disinformation, we would expect to see that content of tweets to vary significantly by gender. As a first pass at differences in content, consider the wordclouds presented in Figure 3, which represent the top 50 most frequently tweeted words[7].

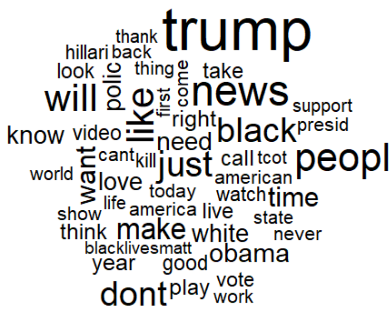

(a) Male-Assigned Accounts

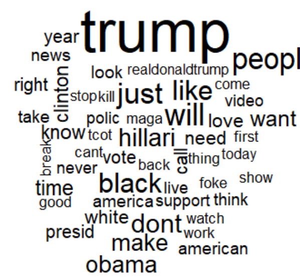

(b) Female-Assigned Accounts

Figure 3: **Male- and Female-Assigned Accounts Tweet about the Same Things with Different Frequencies** Comparing the wordclouds, we see similarities in the overall vocabulary of words used, but differences in the frequencies with which those words are used.

As evidenced in the wordclouds, the vocabulary used by both male- and female-assigned accounts is actually very similar. Both use many of the same words, with President Trump being mentioned the most.[8] However, female-assigned accounts are more likely to Hillary Clinton, who appears 8th most frequently as opposed to 30th most-frequently for male-assigned accounts. Obama, on the other hand, is tweeted about with the same relative frequency, 10th place for Females and 12th place for Males. Interestingly, the hashtag "#tcot," short for "top conservative on twitter," also occurs with the same relative frequency for male- and female-assigned accounts. This is counter to gendered expectations that women are expected to be less ideological than men. Another interesting difference is for the word stem "polic," which is connected to police and policing. This word occurs much more frequently for male-assigned

[7] Common Natural Language Processing practices were used to generate this list, including stemming, removing stopwords, and removing punctuation.

[8] The appendix contains frequency tables for the top 50 words.

accounts (16th) than for female assigned accounts (33th). This suggests a law and order dimension that may be being pushed along gender lines.

Words that appear in the top 50 for male-assigned accounts and not for female-assigned accounts include “play,” “life,” “thank,” “state,” “world,” and “blacklivematt.” Words that appear for the top 50 for female-assigned accounts but not for female-assigned accounts include “clinton,” “maga,” “break,” “foke,” “realdonaldtrump,” and “stop.” It is interesting that female-assigned accounts were more likely to tweet at Trump than male-assigned account, and this may suggest greater appeals to authority or attempts at establishing stronger female-affect for Trump.

Moving beyond word frequencies, we can also consider longer phrases. While words are singular units, tri-grams consist of three words, and allow for individual words to be considered in a greater context. In order to evaluate the different word frequencies in context, we can compare the tri-gram frequency presented in Figure 4.

As was the case when it came to word frequencies, male- and female-assigned accounts do frequently tweet about similar things, including the White House or the United States. However, there are distinctive phrases used by each group. Male-assigned accounts, for example, are much more likely to use phrases referring to ISIS. Female-assigned accounts on the other hand, are much more likely to use phrases such as MAGA or POTUS that refer Trump. This suggests that male-assigned accounts may have been used to call attention to terrorism or Islamophobia as target while, once again, female-assigned accounts are used to create affinity for Donald Trump.

Gender Differences in Category Assignment

Finally, when looking at a corpus as a whole, it is useful to be able to classify tweets into intelligible groups. Based on a mixed methods approach, Linvill and Warren (2020) have developed a taxonomy for IRA tweets based on a qualitative evaluation of the content of tweets. These groups include Fearmonger, Hastag Gamer, Left Troll, News Feed, Non-English, Right Troll, and a residual unknown category. Based on data provided by the authors, we tallied the gender composition of each of these groups.

In Table 3 below, we present the gender breakdown of each of the handle categories defined by Linvill and Warren (2020).

Gender Assignment Breakdown by Handle Category					
Account Type	Fear Monger	Hashtag Gamer	Left Troll	Non-English	Right Troll
Female-Assigned	37	38	50	39	147
	41.6%	49.3%	34.5%	16.7%	46.5%
Male-Assigned	52	39	95	21	169 49
	58.4%	50.6%	65.5%	83.3%	53.5%

$\chi^2 = 23.49, df = 6, p \leq 0.001$

Table 3: **Male- and Female-assigned Accounts Tweet Different Topics** Comparing the gender distribution for different handle categories, male-assigned accounts are disproportionately grouped as Fear Monger or Left Troll. There is more balance among the Hashtag Gamer and Right Troll groups.

Generally speaking, the number of male-assigned accounts in each category tends to outnumber the number of female-assigned accounts. This is to be expected given that there were more male-assigned accounts than female-assigned accounts. What is interesting is where gender imbalances fell among the categories. For example, female-assigned accounts are slightly over represented in the Hashtag Gamer category and the Unknown category. An alternative interpretation would be that male-assigned accounts are over-represented in more traditional political categories of Right Troll and especially Left Troll. Since women are traditionally thought of as more liberal than men[9], these pattern defies conventional wisdom. This may be an attempt by disinformation agents to undermine Leftist causes by misrepresenting salient voices.

It is also interesting that women are over-represented in the "Unknown" category. This may highlight a shortcoming in the rules for taxonomy, but it may also represent a less targeted approach to disinformation when tweeting from female-assigned accounts. Relying on traditional gender norms, for example, politics is typically a masculine domain (among many others Sapiro 1983). Additionally, more traditional countries with respect to gender roles, like Russia, are more likely to carry those norms into international relations (Karim and Beardsley 2016). Thus, the gender imbalance seen

[9] see, e.g. Box-Steffensmeier, De Boef and Lin (2004)

in the "Unknown" category may be a function of Russian agents being limited in how they portray women in disinformation campaigns.

Tweet Activity per Handle Category by Gender					
Account Type	Fear Monger	Hashtag Gamer	Left Troll	Non-English	Right Troll
Female-Assigned	4072	86,519	104,006	595	180,386
	51.3%	49.6%	33.0%	66.7%	57.4%
Male-Assigned	3871	88,030	211,153	298	133,933
	48.7%	50.4%	67.0%	33.3%	42.6%

Table 4: **Female-Assigned Accounts are More Likely to Tweet Many Things** Whereas female-assigned accounts do not make up they majority of accounts in any account category, they make up the majority of tweets in several categories, including Right Troll and Fear Monger.

As a final consideration, it may also be the case that tweet activity differs in important ways from the number of accounts. Table 4 shows the breakdown of tweets per account category. Surprisingly, the number of tweets does not tell the same story as the number of accounts in an account category. Indeed, while female-assigned accounts were not a majority in any substantive category by account number, they are leaders in activity for the Right Troll and Fear Monger categories.

The predominance of tweets from female-assigned accounts in these two categories suggests a number of possibilities. Perhaps the volume of these tweets suggests that female voices in these categories were more credible when they were very active. It may also suggest that female voices were amplified when they did tweet, which would explain why fewer accounts could generate more activity. Finally, the use of female voices in the Right Troll and Fear Monger categories may be an attempt to reach out beyond traditional gender norms to reach a broader audience.

Conclusion

Is there a gendered pattern to efforts of disinformation from the IRA? Substantial evidence suggests yes. Twitter accounts with female-assigned names tweet less and tweet on different topics when compared to their male-assigned counterparts. Given that only one intelligence agency, the IRA, had been attributed to these tweets, it is especially compelling that an effort was made to send gendered messages. Depending on the issue, female-assigned accounts may potentially have louder and more influential voices. This evidence, though exploratory, suggests many additional questions about how gender can be leveraged in disinformation campaigns.

While we have tried to present the best evidence possible, future research can take many approaches to digging deeper into this research question. From a technical standpoint, future research may be better able to identify the presented gender of Twitter accounts. Recent developments in computer vision suggest that higher identifications are possible using account pictures, for example. While not available in this data set, using both the account profile picture or analyzing the pictures tweeted by the account may give additional signal when classifying the gender of accounts.

As for more substantive questions, this research has been limited just the release of IRA tweets from 2012 - 2018. Future research may very well expand this data set, or attempt at streaming classification of disinformation campaigns. While it is clear that gender plays a role in disinformation campaigns, it is not necessarily clear the motivations behind certain content, and further qualitative research may identify these motivations. Finally, the patterns here raise interesting questions about cross-national uses of gender in disinformation, and examining how cultural context shapes these behaviors would greatly contextualize the results presented here.

References

Bail, Christopher A. 2016. "Emotional Feedback and the Viral Spread of Social Media Messages about Autism Spectrum Disorders." *American Journal of Public Health* 106(7):1173–1180.

Box-Steffensmeier, Janet M, Suzanna De Boef and Tse-Min Lin. 2004. "The Dynamics of the Partisan Gender Gap." *American Political Science Review* pp. 515–528.

Carpenter, Jeffrey P and Peter Hans Matthews. 2012. "Norm Enforcement: Anger, Indignation, or Reciprocity?" *Journal of the European Economic Association* 10(3):555–572.

Karim, Sabrina and Kyle Beardsley. 2016. "Explaining Sexual Exploitation and Abuse in Peacekeeping Missions: The Role of Female Peacekeepers and Gender Equality in Contributing Countries." *Journal of Peace Research* 53(1):100–115.

Linvill, Darren L and Patrick L Warren. 2020. "Troll Factories: Manufacturing Specialized Disinformation on Twitter." *Political Communication* pp. 1–21.

Mason, Lilliana. 2018. *Uncivil Agreement: How Politics Became our Identity.* University of Chicago Press.

Sapiro, Virginia. 1983. *The Political Integration of Women: Roles, Socialization, and Politics.* University of Illinois Press.

Tucker, Joshua A, Andrew Guess, Pablo Barberá, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal and Brendan Nyhan. 2018. "Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature." *Political Polarization, and Political Disinformation: A Review of the Scientific Literature* .

Webster, Steven W. 2020. *American Rage: How Anger Shapes Our Politics.* Cambridge University Press.

Weeks, Brian E. 2015. "Emotions,Partisanship, and Misperceptions: How Anger and Anxiety Moderate the Effect of Partisan Bias on Susceptibility to Political Misinformation." *Journal of Communication* 65(4):699–719.

West, Candace and Don H Zimmerman. 1987. "Doing Gender." *Gender & Society* 1(2):125–151.

Appendix

Rank	Word	Freq	Rank	Word	Freq
Top 50 Words for Male-Assigned Accounts					
1	trump	28037	26	good	7505
2	news	19612	27	video	7323
3	peopl	17435	28	kill	6893
4	will	17141	29	vote	6863
5	like	17134	30	hillari	6811
6	just	15862	31	watch	6761
7	black	15119	32	show	6736
8	dont	15023	33	america	6706
9	make	12567	34	presid	6671
10	want	10938	35	work	6667
11	time	10896	36	come	6632
12	obama	10573	37	never	6515
13	love	10198	38	back	6484
14	know	10068	39	today	6432
15	need	10022	40	thing	6389
16	polic	9604	41	american	6285
17	white	9385	42	cant	6277
18	right	8749	43	life	6224
19	year	8540	44	tcot	6196
20	call	8080	45	thank	6000
21	live	7975	46	first	5979
22	take	7898	47	state	5974
23	play	7781	48	support	5885
24	think	7764	49	world	5884
25	look	7694	50	blacklivesmatt	5626

Table 5: Top 50 Word Frequencies for Male-Assigned Accounts

Top 50 Words for Female-Assigned Accounts					
Rank	Word	Freq	Rank	Word	Freq
1	trump	30285	26	take	5983
2	will	14017	27	think	5982
3	peopl	13777	28	american	5758
4	just	13231	29	live	5499
5	like	12372	30	video	5348
6	black	11722	31	never	5292
7	dont	11294	32	support	5263
8	hillari	10537	33	polic	5184
9	make	9522	34	watch	5098
10	obama	9513	35	good	5066
11	want	8397	36	maga	4969
12	time	8173	37	cant	4946
13	clinton	7358	38	work	4931
14	know	7347	39	show	4920
15	white	7231	40	break	4896
16	love	7207	41	foke	4843
17	need	7135	42	back	4813
18	year	6616	43	tcot	4679
19	call	6607	44	first	4671
20	vote	6589	45	kill	4670
21	presid	6310	46	thing	4610
22	right	6273	47	realdonaldtrump	4544
23	news	6105	48	come	4471
24	america	6072	49	today	4464
25	look	6054	50	stop	4439

Table 6: Top 50 Word Frequencies for Female-Assigned Accounts

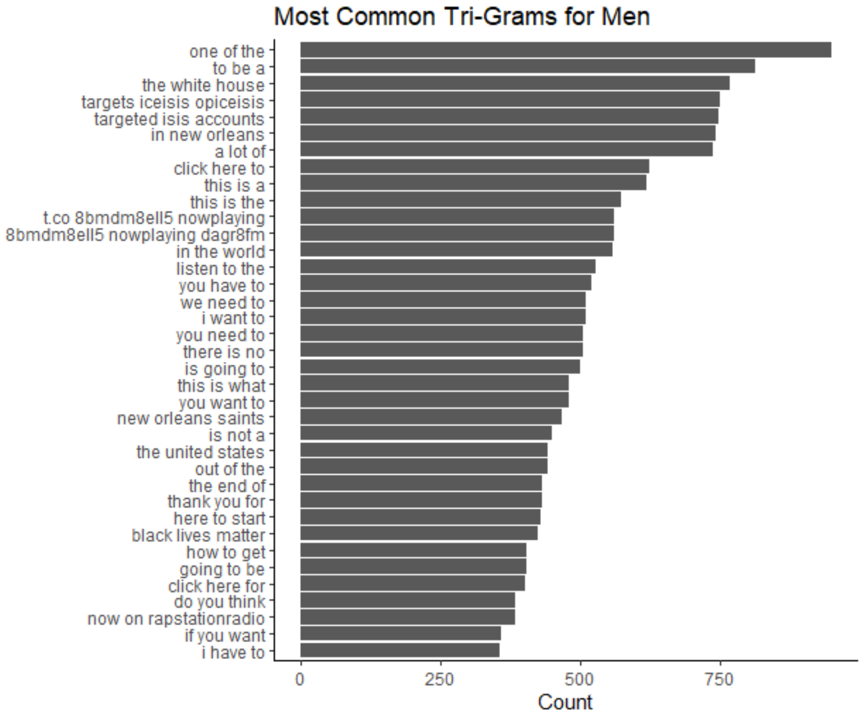

(a) Male-Assigned Accounts

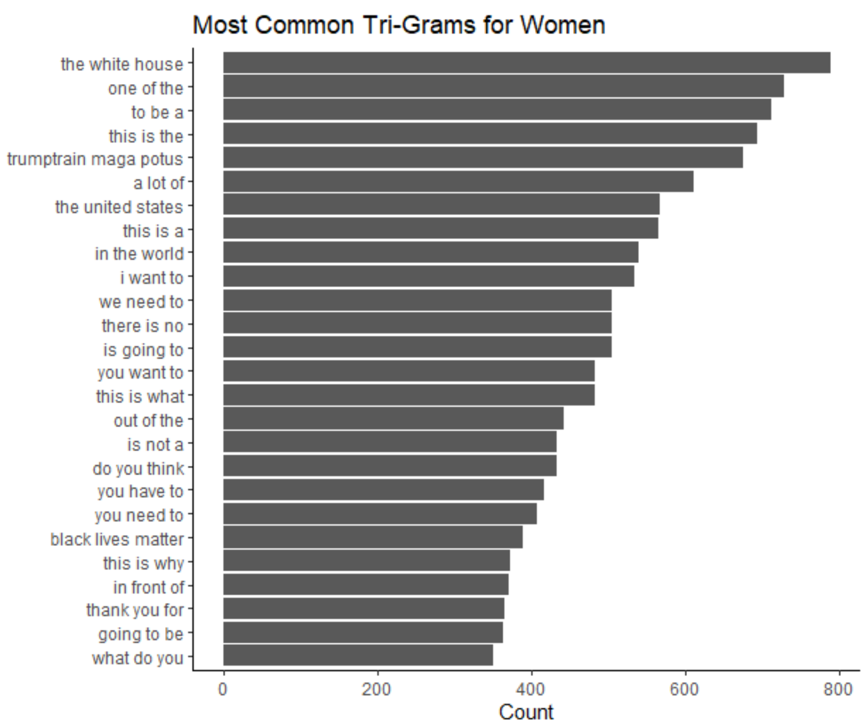

(b) Female-Assigned Accounts

Figure 4: **Male- and Female-Assigned Accounts Use Different Phrases to Tweet** Comparing the frequency of different tri-grams, male- and female-assigned accounts use different language overall, even if they use similar vocabulary at the word level.