International Journal of Communication 15(2021), 4662–4689 1932–8036/20210005

Trollfare: Russia's Disinformation Campaign During Military Conflict in Ukraine

LARISSA DOROSHENKO
Northeastern University, USA

JOSEPHINE LUKITO[1]
The University of Texas at Austin, USA

In this study, we explore online informational warfare by the Russian Internet Research Agency (IRA) against Ukraine during the military conflict in Donbass. Introducing a digital dimension to the long-standing Russian disinformation strategy of reflexive control as a historic and theoretical framework, we investigate how the IRA combined online news and social media platforms to promote propaganda to its growing number of followers. Combining computational and qualitative content analyses with time series modeling, we demonstrate how the IRA blurs distinctions between fact and fiction through interlinks among digital platforms, and we expose its successful strategies for follower growth on Twitter. We conclude with implications for understanding and promptly identifying modern hybrid warfare strategies, with a focus on coordinated multiplatform efforts that spread disinformation through the hybrid media ecosystem.

Keywords: disinformation, information warfare, Internet Research Agency, Russia, Ukraine, Donbass, topic modeling, time-series analysis

ColdWar 2.0, a website concocted by the Russian Internet Research Agency (IRA) during the 2014 conflict in Eastern Ukraine, epitomizes Cold War disinformation strategies in the digital era. Feigning the appearance of a news organization (see Figure 1), the website was linked to social media platforms based in both Russia and the United States, and spread the Kremlin's propaganda to dismay and distract its readers and social media users. Twitter accounts affiliated with ColdWar 2.0 posted links to its stories, crafted hashtag campaigns, and interacted with other IRA-connected accounts. The website's affiliations came to light only after Twitter suspended IRA-associated handles following an investigation into Russian meddling during the 2016 U.S. elections (*U.S. v. Internet Research Agency LLC*, 2018). As it turns out, the Russian and English-language ColdWar2.0 Twitter handles were among the most popular accounts in the aftermath of the

Larissa Doroshenko: l.doroshenko@northeastern.edu
Josephine Lukito: jlukito@utexas.edu
Date submitted: 2020-11-17

[1] The authors would like to thank Brooke Foucault Welles and the anonymous reviewers for their valuable advice on revising this article.

Euromaidan revolution in Ukraine. How did these disinformation actors gain traction to amplify their message and increase traffic to their propaganda website? Using a 10% Twitter Gardenhose archive, we retroactively inspected the digital strategies employed by prominent IRA accounts and identified those that successfully grew their accounts' popularity. Our analysis revealed that they did not just spread falsehoods; rather, they engaged in reflexive control—a Cold War–era strategy to alter key factors in an adversary's perception of the world, thereby encouraging that adversary to make decision that were favorable to a controlling agent (Giles, Sherr, & Seaboyer, 2018; Snegovaya, 2015).

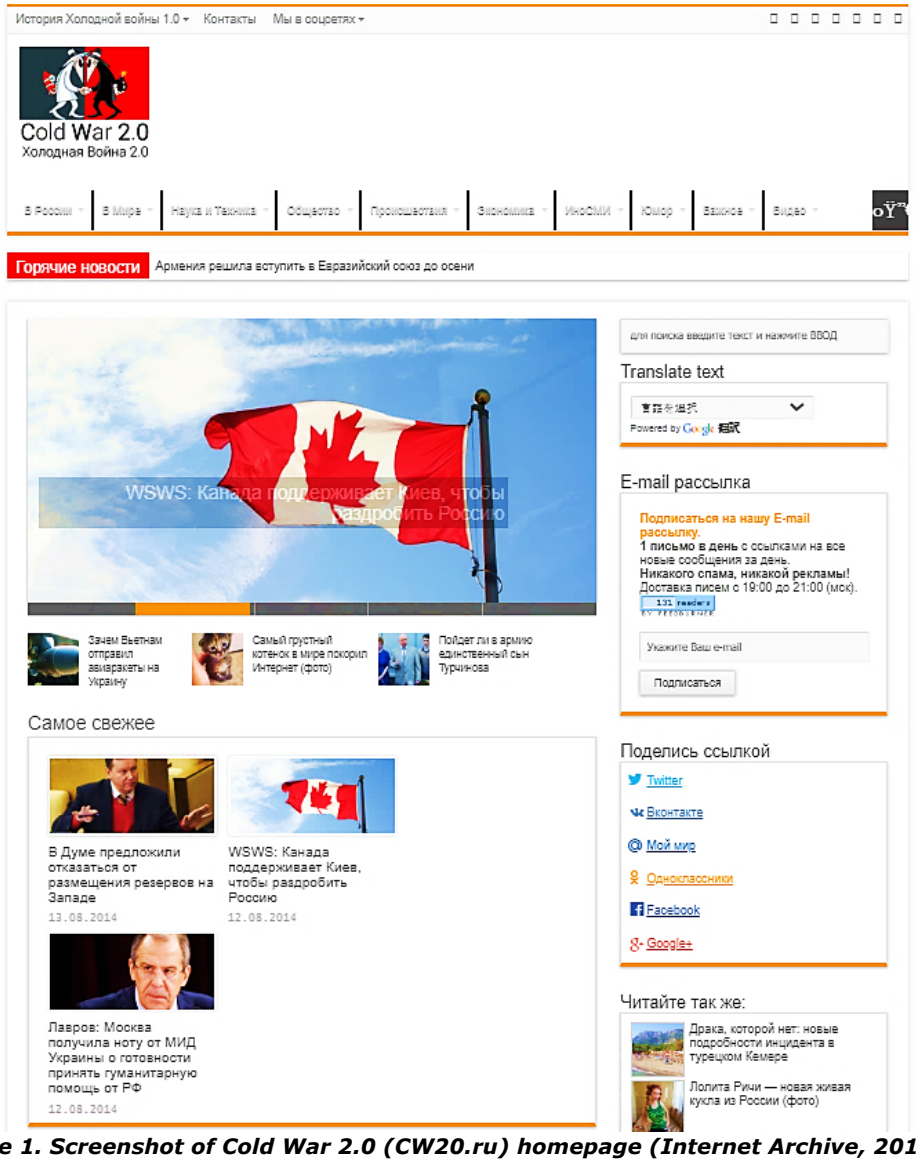

Figure 1. Screenshot of Cold War 2.0 (CW20.ru) homepage (Internet Archive, 2014; https://web.archive.org/web/2020*/cw20.ru).

This article makes several timely contributions to the study of disinformation campaigns. First, it enables us to see how Russia improved its digital disinformation campaign within its geopolitical sphere of influence before exporting the most effective strategies to the West. As other studies have shown, the IRA's disinformation campaign targeting the United States in 2016 employed fake Twitter accounts that impersonated local news aggregators (Bastos & Farkas, 2019) and accumulated followers through retweets (Zhang et al., 2021). Russia first developed these strategies in Ukraine. By revealing other successful tactics employed in Donbass, this study provides insights on the increasing sophistication of these campaigns and highlights ways to detect disinformation strategies in future IRA campaigns.

Second, this article uses a historic and theoretical framework of reflexive control (Thomas, 2004) to better understand disinformation campaigns. We demonstrate how the Soviet disinformation toolkit has been adopted to the digital realm and emphasize that the primary goal of falsehoods is not just to deceive an adversary, but to engage that adversary in poor decision making. Third, we highlight Russia's transition from traditional to cyber propaganda, where both news media and their associated media accounts work in tandem to engage in information warfare. We describe how the IRA used the hybrid media ecosystem of online news aggregators and their social media handles to distract from reality, distort it, and dismay both Russians and Ukrainians.

We use a theoretical framework of reflexive control, demonstrating how its 4D strategies were used on digital platforms to amplify propaganda and gain the attention of social media users. Through a combination of computational and qualitative content analyses, we offer insight into this coordinated disinformation campaign by studying the interlinks across several platforms (which blurred the distinction between fact and fiction) and exposing the IRA's successful strategies for follower growth. We conclude with implications for understanding modern hybrid warfare strategies, with a focus on multiplatform digital efforts among IRA-affiliated accounts, which we expect to continue into the 2020s.

Reflexive Control as a Weapon of Russian Information Warfare

Heightened attention to the Russian disinformation campaigns since 2016 facilitates the illusion that disinformation is a modern issue brought about by the interactive Web and social media. However, to understand these strategies and recognize them proactively, it is imperative to understand the history of Russian disinformation strategies, honed for over four decades (Thomas, 2004). Acknowledging the variations in defining disinformation, in this article we describe it as the intentional and coordinated spread of false, inaccurate, or misleading information designed, presented, and promoted to achieve a political communication goal (European Commission, 2018; Fetzer, 2004). Disinformation messages are munitions and nonlethal weapons in modern information warfare, which scholars believe are intended to subdue adversaries rather than reason with them" (Freelon & Wells, 2020, p. 146). However, this is where Russian disinformation strategies are often misunderstood and underestimated: These campaigns engage with their opponents' reasoning and lead to erroneous decision making, a strategy called reflexive control.

The concept of reflexive control was first developed at the height of the Cold War in the 1960s (Levefr & Smolyan, 1968) as "a process by which one enemy transmits the reasons or bases for making decisions to another" (Thomas, 2004, p. 238). According to Russian generals, American use of information weapons did

more damage to the Soviet Union's defeat that any other weapon (Prokhozhev & Turko, 1996, as cited in Thomas, 2004). Reflexive control happens when the controlling actor presents an enemy with information that leads the enemy to a desired decision (Leonenko, 1995). This "control" is described as "reflexive" because a key component of this strategy involves the backtracking process from the desired outcome to the enemy's current reasoning or possible behavior (Thomas, 2004). The chief task of reflexive control is to find and exploit weak links in information assessment during decision making. Russian disinformation strategies are not meant to just present falsehoods and confuse adversaries. Rather, the goal is to spread disinformation that would lead adversaries to make erroneous decisions favoring Russia, the controlling agent.

The Soviet toolkit of reflexive control strategies nowadays is described along the *4D dimensions*: dismiss, distort, distract, and dismay (Snegovaya, 2015). *Dismissing* presents evidence in a way that obfuscates objectives of the controlling agent or denies presented evidence. *Distortion* alters one's perception of reality by presenting various falsehoods: from made up "facts" to characteristics of institutions and people. *Distraction* creates a real or imaginary threat or reveals new evidence, which forces an adversary to reconsider a decision. *Dismay* buffs and dramatically escalates the situation to discourage an opponent from taking an action. With the advent of digital technologies, reflexive control has been adapted for computational propaganda and cyberwarfare.

Digital Information Warfare and Its Cyber Soldiers

Recent scholarship has modified the traditional definition of disinformation and propaganda to account for its digital and programmable component. Computational propaganda is used to describe the assemblage of social media platforms, autonomous agents, and big data, which are collectively employed to manipulate public opinion (Woolley & Howard, 2016). The most well-known source of Russian computational propaganda is the IRA, Russia's "troll army," which was officially registered in 2013 (Garmazhapova, 2013).

Cyberwarfare, like other military operation, needs its soldiers, which can include bots and sock puppets. Bots are programs that automate the online activity of social media accounts, including sending mass messages and retweeting content (Snegovaya, 2015; Zannettou et al., 2019). Sock puppets, meanwhile, are fake personas managed by real people. A sock puppet's persona adds meaning and credibility to social media discourse (Cook, Waugh, Abdipanah, Hashemi, & Rahman, 2014). In this project we focus on sock-puppet content specifically because their influence is harder to detect compared with automated accounts, making insights obtained with the ground-truth data particularly useful for developing better detection strategies. We pay special attention to the success of influential sock puppets to understand how they gained popularity with false personas.

We first focus on the context of informational warfare against Ukraine to better understand disinformation strategies that sock puppets have employed on Twitter to advance the Kremlin's geopolitical goals during military conflict in Donbass. Secondly, we shift our attention to specific tactics that successful sock puppets might have used for growing their army of followers on Twitter and, as a result, boosted their credibility across media ecosystem.

Ukrainian Military Conflict and the IRA's Reflexive Control

The IRA's reflexive control strategy during the Ukrainian military conflict, which began after the Euromaidan revolution, is a useful case for understanding digital information warfare. Initially, the Kremlin denied their involvement in the annexation of Crimea and support of insurgents in Donbass. This was coupled with a disinformation campaign in state-controlled Russian media against the new Ukrainian government, which portrayed the annexation of Crimea as the will of local people to unite with Russia and the conflict at Donbass as a civil war (Golovchenko, Hartmann, & Adler-Nissen, 2018). This portrayal aligned with the conceptualization of Euromaidan by protesters. Contrary to Western coverage of the revolution, which presented it as a geopolitical struggle, Ukrainians on Facebook discussed Euromaidan in terms of domestic issues and as an antiregime protest (Surzhko-Harned & Zahuranec, 2017).

Social media metrics reveal that there was more attention toward Ukraine after the revolution, despite the fact that protesters used these platforms for mobilization and coordination during Euromaidan (Bohdanova, 2014). The three largest spikes in Twitter and Facebook post activity about Ukraine occurred during 2014: when the ex-president of Yanukovych fled, when Crimea was annexed, and when the Malaysia Airlines flight was shot down (Onuch, 2015). These metrics align with information warfare tactics—to the Kremlin, the key strategic moment began after the protesters ousted the pro-Russian government in Kyiv. The IRA aspired to distort and change the narrative about the consequences of the revolution and the interim Ukrainian government, using disinformation "soldiers" to spread messages across social media.

Along with attempts to discredit Euromaidan's outcomes, the IRA sought to keep the audience distracted. It developed and supported various conspiracy theories, using the principle "if nothing is true, then anything is possible" (Pomerantsev, 2015). This distorted perception of reality encouraged actions that aligned with Russian interests—a goal of reflexive control. Similar to the Cold War period, the IRA kept the West from engaging in an open confrontation with Russia by fueling denial of Russian troop presence in Donbass and dismissing Moscow's involvement in Ukraine (Snegovaya, 2015).

IRA disinformation was not only supported by offline Russian news; it was perceived as organic content in the traditional media environment. Russian media supported the IRA's disinformation campaign by characterizing the Ukrainian government as illegitimate and brutal, using derogatory terms such as "the fascist junta" and "Banderites," referring to Ukrainian WWII independence movement leader Stepan Bandera (Snegovaya, 2015). Certain sock puppets also acted as news feed accounts (Bastos & Farkas, 2019; Linvill, Boatwright, Grant, & Warren, 2019) or anonymous websites simulating news media (Alexander, 2015).

Summarizing previous research about Russian information warfare and reflexive control, we strive to improve our scholarly understanding of computational propaganda tactics, posing and answering the following research questions:

RQ1: What disinformation tactics were used by the IRA accounts during information warfare against Ukraine?

RQ2: How were these disinformation tactics manifested in messages spread by the IRA accounts during information warfare against Ukraine?

Influence Building on Twitter: Possible Tactics of the IRA Sock Puppets

Studies have noted that the IRA relies heavily on audience metrics, which are often considered "objective indicators" of newsworthiness (Webster, 2014). Audience metrics can also signal an account's value and status (Marwick, 2013). Follower count, a specific metric reflecting an account's popularity, is often considered by journalists when assessing the value of tweets (Chorley, Colombo, Allen, & Whitaker, 2015; Lukito et al., 2020). Unsurprisingly, accounts with many followers amplify word-of-mouth-effects (Chen, Haber, Kang, Hsieh, & Mahmud, 2015), creating a perception of credibility (Jin & Phua, 2014) and making follower count valuable social media capital. It, therefore, behooves political actors, including disinformation agents, to grow their follower count (Saxton & Guo, 2014). Though following an account is a low-labor act, people are selective in who they follow, especially political accounts (Wang, Luo, & Zhang, 2017).

Previous research offers a useful classification of different user groups based on their number of followers; we apply this typology to sock puppets. *Mass-media* are extremely well-connected users with more than 100,000 followers, *grassroots* are the least connected users with fewer than 200 followers, and *evangelists* are the remaining well-connected small group of users (Cha, Benevenuto, Haddadi, & Gummadi, 2012). Grassroots and evangelists usually reciprocate their followers, while mass media does not; however, when mass media forms a link, it is reciprocated with a high probability of 88.6%; contrastingly, grassroots get this reciprocation only 16.2% of the time. Though follower counts can be artificially inflated (Aggarwal, Kumar, Bhargava, & Kumaraguru, 2018), this audience metric still greatly influences social media users' perceptions of an account (Walther, Van Der Heide, Kim, Westerman, & Tong, 2008).

Studies have pointed to several "organic" tactics for increasing one's follower count. The most extensive of these is Hutto, Yardi, and Gilbert (2013), who investigated 17 variables potentially predicting follower growth. They found that various content and user attributes, such as the frequency of activity, use of hashtags, retweet activity, and the use of URLs, all contribute to an account's follower growth.

#Hashtags identify and engage conversation around trending topics. Users include them to maximize the chances of being noticed and to increase engagement with their content (Lahuerta-Otero & Cordero-Gutierrez, 2016). A popular hashtag can generate thousands of tweets, as people come up with creative ways to use it (Parker, 2016). Research has noted that, in 2016, IRA Twitter sock puppets frequently employed political and non-political hashtags to attack civil institutions (Linvill et al., 2019). We hypothesize that similar tactics were employed in information warfare against Ukraine:

H1a: Including relevant hashtags in tweets helped prominent IRA-linked accounts to gain followers when using information warfare against Ukraine.

The @mention is another mechanism to increase message diffusion. The more an account is mentioned, the longer distance travels information it spreads (Yang & Counts, 2010). Users employ @mentions to engage in conversation with other users and to expose them to opposite views or start a discussion (Conover et al., 2011). To increase social media prominence, users seek being mentioned and mention others, leading to the next hypothesis:

H1b: Including @mentions in tweets helped prominent IRA-linked accounts to gain followers when using information warfare against Ukraine.

URLs embedded in tweets correlated with higher credibility scores for the tweet and its author (Gupta & Kumaraguru, 2012) and enabled these messages to reach farther across networks (Yang & Counts, 2010). A study about the dissemination of ISIL propaganda on Twitter observed that bots strategically shared URLs to sites and blogs of this organization (Al-khateeb & Agarwal, 2015), suggesting that accounts used URLs for self-promotion. Based on these findings, we propose our last hypothesis:

H1c: Including URLs in tweets helped prominent IRA-linked accounts to gain followers when using information warfare against Ukraine.

Methods

Data

We used a ground-truth list of 2,752 Twitter handles published by the U.S. House Intelligence Committee (2017), which Twitter identified as affiliated with the IRA. Using the account's ID numbers, we searched a 10% Twitter Gardenhose archive for tweets and retweets from each account from November 1, 2013, to December 31, 2014. This time frame was selected because it encompassed the Euromaidan revolution, which lasted from November 2013 to February 2014, as well as subsequent confrontations with Russia, including the annexation of Crimea, the military conflict in Donbass, the Minsk Protocol signed in September 2014, and their aftermaths. Unlike the data set created by FiveThirtyEight (Roeder, 2018), our data (collected in real time) include6 the number of followers each account had during the study period, enabling the analysis of follower growth.

Our search yielded 193,495 tweets. To align the data temporally, we converted the time stamp of each tweet to the GMT+3 (Moscow time zone). We disaggregated all shortened URLs in our corpus, resulting in 94,068 links, which were inductively split based on their main URL domains into nine categories in no particular order (see Table 1).

Table 1. Categories and Counts of URLs, #Hashtags, and @Mentions Used by IRA Handles, 2013–14.

	URLs		#Hashtags		@Mentions
1	RT (Russia Today) (*N* = 1,640)	1	Donbass conflict (*N* = 4,647)	1	Political and public figures (*N* = 1,109)
2	Other Russian mass media (*N* = 7,381)	2	Boeing shutdown (*N* = 1,568)	2	Western mass media & social media platforms (*N* = 117)
3	Western mass media (*N* = 132)	3	Western mass media (*N* = 7)	3	IRA accounts (*N* = 3,644)
4	Self-referencing links (*N* = 55,146)	4	Russian mass media (*N* = 53)	4	Journalists (*N* = 259)
5	Ukrainian mass media (*N* = 349)	5	Crimea & Odessa (*N* = 850)	5	Regular Twitter users (*N* = 3,070)
6	Russian unverified news (*N* = 4,878)	6	Euromaidan (*N* = 109)	6	Russian news media (*N* = 296)
7	Russian government websites (*N* = 554)	7	Kyiv & Ukraine (*N* = 1,850)	7	Russian unverified news (*N* = 356)
8	Livejournal (*N* = 8,751)	8	Western states & alliances (*N* = 236)	8	Organizations & institution (*N* = 399)
9	Social media platforms (*N* = 7,072)	9	Russia, Kremlin, & Moscow (*N* = 528)	9	Euromaidan (*N* = 155)
		10	Politicians & Parties (*N* = 829)		
		11	Hashtag campaigns (*N* = 1,654)		
		12	#новости (news) and #политика (politics) (*N* = 482)		

We were unable to read or categorize 3,629 links, which were excluded from later analyses. Websites were labeled as nonverified news if they presented themselves as a mass media outlet, but failed to provide registration information, which is required for all Russian mass media (Roscomnadzor, 2017). Likewise, we

constructed and qualitatively assigned typologies for two other content features, which might have been strategically used by the IRA: hashtags and @mentions (see Table 1). An "other" category—3.8% of the URLs and 14.4% of the hashtags—was excluded from the time series modeling. In our data set, there were 14,974 tweets with hashtags and 21,235 tweets with @mentions. Table 1 includes only categories and counts for all unique @mentions used by four selected IRA accounts (see Table 2), while URLs and hashtags are classified for the entire data set.

Table 2. Description of Four IRA Handles.

Handle	Date account created	# of tweets in the data set	# of RT in the data set	# of followers at start of series	# of followers at end of series
coldwar20_ru	02/19/2014	1,406	41,008	2	21,876
coldwar20_en	04/14/2014	268	2,895	3	3,326
KadirovRussia	12/29/2011	534	27,793	16,829	76,525
LavrovMuesli	07/21/2014	125	1,281	0	15,339

In addition to social media content, we also considered the impact of four types of external events: mass protests, violent protests, military actions, and political events. This event data, comprising 139 events, was constructed using the English-language Wikipedia time lines of Euromaidan, the Crimea annexation, and the conflict in Eastern Ukraine, and by checking references in reputable international news media (e.g., AP, BBC, Euronews). Mass protests included demonstrations of more than 100 people without violent outcomes, while violent protests included demonstrations of more than 100 people that resulted in bodily injures and/or causalities among protesters. Military events ranged from attacks and fights between Ukrainian army and insurgents to sieges of buildings and/or cities and the launching of missiles. Political events encompassed resignations, proclamations, declarations, negotiations, agreements, and exchange of prisoners.

Relevant Tweets Corpus

We start our analysis by separating our data into tweets relevant to our project and control corpus comprised of other tweets posted at the same time. It is needed to identify whether type of content included in IRA tweets is significantly different from other Twitter users (RQ2). To identify tweets related to Ukrainian revolution and the ensuing conflict in its South-East region, two native Russian speakers coded 2,000 randomly sampled tweets for their relevance. This was coded as a binary variable (0 = not relevant, 1 = relevant); tweets comprising only of URLs were excluded. This coding was independently verified by another native speaker of Russian and Ukrainian who coded 10% of that random sample. The agreement between the two coders was very high (Krippendorff's alpha = 0.94; Hayes & Krippendorff, 2007). Based on that sample, a keyword list of 80 Russian and 10 English words and abbreviations was created to construct a relevant tweet corpus using machine learning (all grammatical cases of the Russian keywords were included).[2] Using a combination of the human-coded random sample and the keywords list, a supervised machine-learning algorithm was able to code the entire data set with 80% recall and 95% precision. The sample of relevant tweets included 41,272 tweets or 21.3% of the initial sample.

[2] A list of these keywords is available on the first author's GitHub repository.

Control Corpus

To compare the strategies of IRA users with typical accounts tweeting in Russian, we used a control corpus from the aforementioned Twitter archive. Russian-language tweets within this archive were sampled using a constructed week method. In this stratified sampling strategy, the researcher randomly selects one date per day-of-the-week to construct a "full week" (Hester & Dougall, 2007). Previous studies have employed constructed week samples in their study of social media, particularly Twitter, because of this platform's role within the journalism profession (Armstrong & Gao, 2010). We build a two constructed week sample, with all dates falling within the time frame of our corpus of IRA tweets. We then took a 193,495-tweet sample of this content to ensure this corpus was balanced (i.e., the same size) with the IRA tweet corpus.

LDA Topic Modeling

To explore disinformation strategies used by the IRA during information warfare against Ukraine (RQ1), we used a Latent Dirichlet Allocation (LDA) model to construct a 10-topic model. Perplexity tests consistently indicated that the optimal number of topics for our data was 10. LDA topic modeling is a common strategy for exploring media content because of its flexibility in handling text documents, few tuning parameters, and computing ease (Maier et al., 2018). This strategy has been previously used to study Russian social media (Koltsova & Koltcov, 2013), demonstrating its applicability to Slavic languages.

To construct the 10-topic model, we pre-processed the data by removing numbers, punctuation marks, html marks, URLs, and Russian stop words. We then stemmed each word, removing flexions but preserving suffixes. The corpus of tweets posted by the IRA was converted into a document-term matrix and trimmed to exclude terms appearing in fewer than 15 and more than in 150 tweets. The script for this analysis is available on the first author's GitHub repository.

Time Series

To test our hypotheses, we conducted a series of time series models to determine what type of IRA tweet content helped them gain followers. One advantage of time series is the ability to account for internal autocorrelation: the degree to which a variable at time t is explained by the variable at time t-1 (Wells et al., 2019). For this analysis, we use a technique known as time-series regression modeling (Catalano, Dooley, & Jackson, 1983), which is done in two steps. First, we "prewhiten" the time series variables by taking the residual of an autoregressive integrated moving average (ARIMA) model. This process removes the internal data-generating process explained by the variable's own temporal dynamics. Second, we perform an ordinary least squares (OLS) regression to understand how multiple prewhitened time series are related to one another.

To test how the IRA accounts increased their follower count, we focused on tweets posted by four users: @coldwar20_ru, @coldwar20_en, @KadirovRussia, and @LavrovMuesli. We selected these accounts based on two criteria: the number of retweets they received and their correspondence to classification of user groups based on the number of followers: mass-media, evangelists, and grassroots (Cha et al., 2012). Table 2 summarizes several key descriptives of these four handles.

Acting as a news aggregator, @coldwar20_ru was the most retweeted account in our sample. It was created after violent clashes between police and protesters during Euromaidan and was described as the only account paid by the Kremlin. A sister English-language account, @coldwar20_en, was launched after the occupation of local government buildings in Donbass by pro-Russian insurgents. The accounts @KadirovRussia and @LavrovMuesli were selected because they were among the most retweeted in our data and acted as opinion leaders—or evangelists—by associating with prominent politicians: the head of the Chechen Republic Ramzan Kadyrov and the Russian Minister of Foreign Affairs Sergei Lavrov. In 2011, @KadirovRussia began after the presidential elections and subsequent protests in Russia, and @LavrovMuesli (the youngest) was created after the MH17 crash.

To perform the time-series analysis, we constructed four data sets. Each data contained daily counts of the following variables: the number of tweets posted by the account, the retweets of the account, the use of #hashtags, @mentions, URLs, and the number of followers. To construct these counts, duplicate tweets were removed. Our dependent variable was follower count; we treated the other features as independent variables.

Results

LDA Topic Modeling

Our first research question asked what type of disinformation strategies were used by the IRA accounts. We start answering this question with a computational content analysis. Our 10-topic model demonstrated a presence of reflexive control strategies: *distortion, distraction, dismissal,* and *denial* of involvement. *Distortion* tactics targeted and belittled Ukrainian citizens, the government, and the army—labeling them as fascists, Banderites, betrayers, a mob, a circus, a junta and even "UkroWehrmacht." This tactic portrayed the Ukrainian side in a negative light by playing on far-right sentiments and xenophobia, exacerbating already existing divisions based on language and ethnicity. The IRA also used *distraction* to skew attention toward certain topics, such as the missing Russian photojournalist Andrey Stenin or the assassination of the Ukrainian political far-right activist Oleksandr Muzychko. While these issues were important during the conflict, they were mentioned in two of the 10 topics, showing that the IRA accounts disproportionally emphasized these accidents to distract from other news. Other distraction tweets discussed solidarity with Donbass separatists in Serbia, the humanitarian convoy from Russia to Donbass, and speculations about potential connections between Ukraine's Euromaidan and Syria. The *dismissing* tactic was expressed through the frequent use of terms such as "fake," "lie," and "unclear" when discussing topics pertaining to government actions and information. Lastly, *denial* was evident through the frequent co-occurrence of the words "MH17" and "plane" in tandem with "provocation" and "lie," refuting Russia's involvement in the Boeing crash. The IRA tweets also denied Russia's involvement in Donbass conflict, portraying it solely as a civil war across several topics.

Comparison of IRA and Control Corpuses

Our second research question sought to understand how disinformation strategies were manifested in IRA accounts' messages. To address it, we first compared IRA tweets with the control corpus to identify unique behaviors of Russian sock puppets on Twitter. In the IRA corpus, there were 2,226 unique hashtags

used a cumulative total of 27,269 times. Compared with it, the control corpus used both more hashtags overall (*n* = 69,861) and more unique hashtags (*n* = 8,671). The IRA corpus tweeted more with @mentions (*n* = 139,741) compared with the control corpus (*n* = 71,727). However, the number of unique accounts that the IRA corpus mentioned was 2,226; this was substantively fewer than the control corpus (*n* = 45,502). In addition to using @mentions often, IRA accounts also posted more links in their tweets (*n* =111,901) compared with the control corpus (*n* = 104,621). These results suggest that the core strategy for IRA disinformation accounts on Twitter focused on targeted interactions via @mentions and posting URLs, suggesting the presence of coordinated effort to spread disinformation via multiple accounts and digital platforms.

Hashtag Analysis

To further explore content analysis and validate unsupervised machine learning findings, we turn to the analysis of hashtags used by IRA accounts (see Figure 2). Of the nearly 13,000 (12,926) categorized hashtags, a third (35.6%) were devoted to the military conflict in Ukraine. Like the findings of our topic modeling analysis, many of these hashtags favored Russia's position through *distorting* and *denying* their involvement. For instance, #StopUkrainianArmy, #ХватитБомбитьДонбасс ("enough bombing of Donbass"), and #помощьукраине ("help to Ukraine") blamed the conflict on Ukraine and presented actions of Russia as help to its neighbor. Few hashtags (15%) were devoted to the MH17 crash, following strategies of *dismissing* alternative views, *denying* involvement, and *distorting* information: #UkrainianLie, #КиевОтветьЗаБоинг ("Kyiv take responsibility for the Boeing") #КиевСбилБоинг ("Kyiv shot down the Boeing").

Additionally, IRA accounts also employed hashtag campaigns that were in line with the *distracting* strategy, constituting the third largest group of tweets (12.8%). For instance, a bilingual hashtag campaign #freeAndrew and #свободуАндрею focused attention on the missing photojournalist Andrey Stenin, while #SAVELENIN advocated against dismantling Lenin's statues in Ukraine. Ukrainian struggles to find alternative supplies of natural gas were mocked with another hashtag #СекторБезГаза ("sector without gas"), which is based on a world play with the Russian for Gaza Strip. Other hashtags that were related to politics, Euromaidan, or were explicitly anti-Ukraine constituted a minor fraction of IRA hashtags.

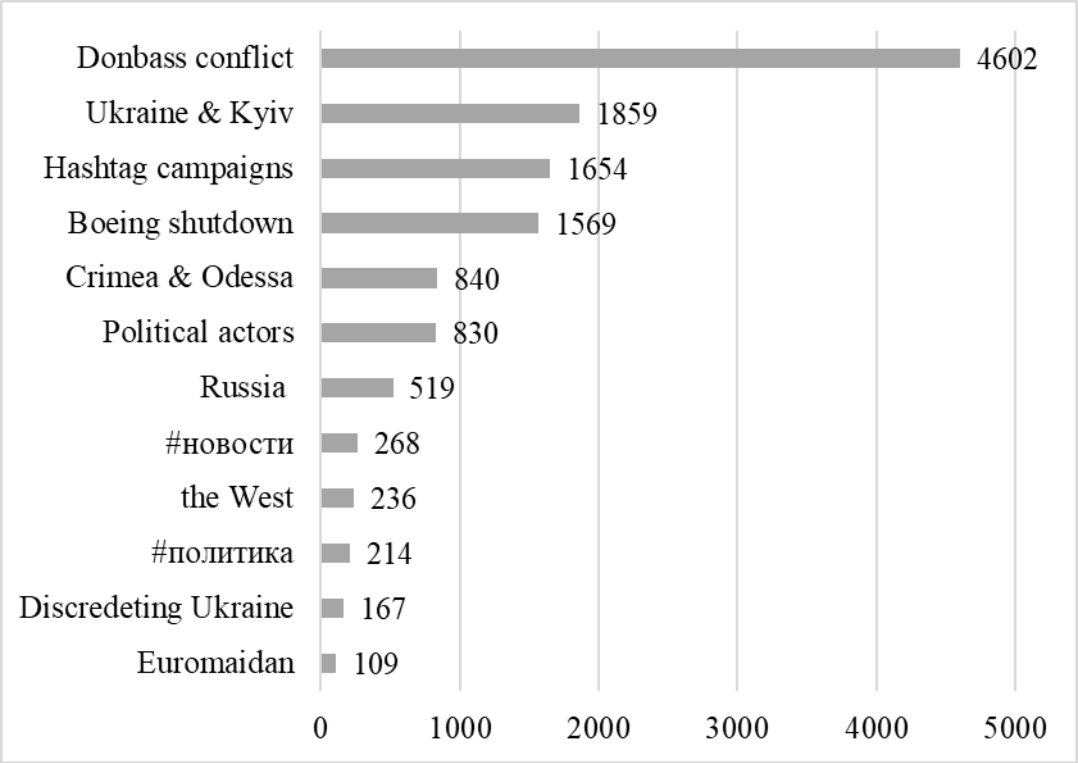

Figure 2. Hashtag categories used by IRA accounts.

Thus, with the content and hashtag analyses, we show how the IRA employed reflexive control 4D strategies in their production of social media messages during and especially after Euromaidan revolution in Ukraine. While the mode of using these strategies has changed with technology, the goal remained the same and supported the continuity of disinformation tactics employed by Russia since the Cold War.

URL Analysis

As mentioned earlier, the major weapon of the IRA information warfare were not hashtags, but links. A majority of the URLs tweeted by IRA sock puppets (55,313, or 58.6% of all the links) directed to their own content, either on Twitter, other social media platforms, LiveJournal blogs, or their own websites (see Figure 4). On examining domains of these self-referential links, we found that most of them included websites associated with three popular IRA accounts: @coldwar20_ru, @nevnov_ru, and @riafanru. At the time this article was written, two of these websites were still live (riafan.ru and nevnov.ru), but the third one (cw20.ru/coldwar20.ru) was inactive. However, all Western social media accounts of the two live websites are suspended.

Nevnov.ru positioned itself as St. Petersburg's news aggregator, while Riafan.ru presented itself as the Federal News Agency, whose website domain uncannily resembles that of RIA News, a reputable Russian

news agency. Unlike the established media outlet, its IRA-connected counterpart included nontraditional news sections such as "Syria," "Ukraine," and "Novorossia"[3] (see Figure 3). According to the mail.ru website traffic statistics, riafan.ru was the third most popular Russian News Agency site, with more than 65,000 visitors daily. CW20.ru, suspended and only available through the Wayback Machine, also acted as a news aggregator with a self-telling name: "Cold War 2.0." This website was devoted to clickbait propaganda news stories about Russia's relationships with two neighboring countries: Ukraine and Belarus, and with the West, mainly the EU, United States, and NATO.

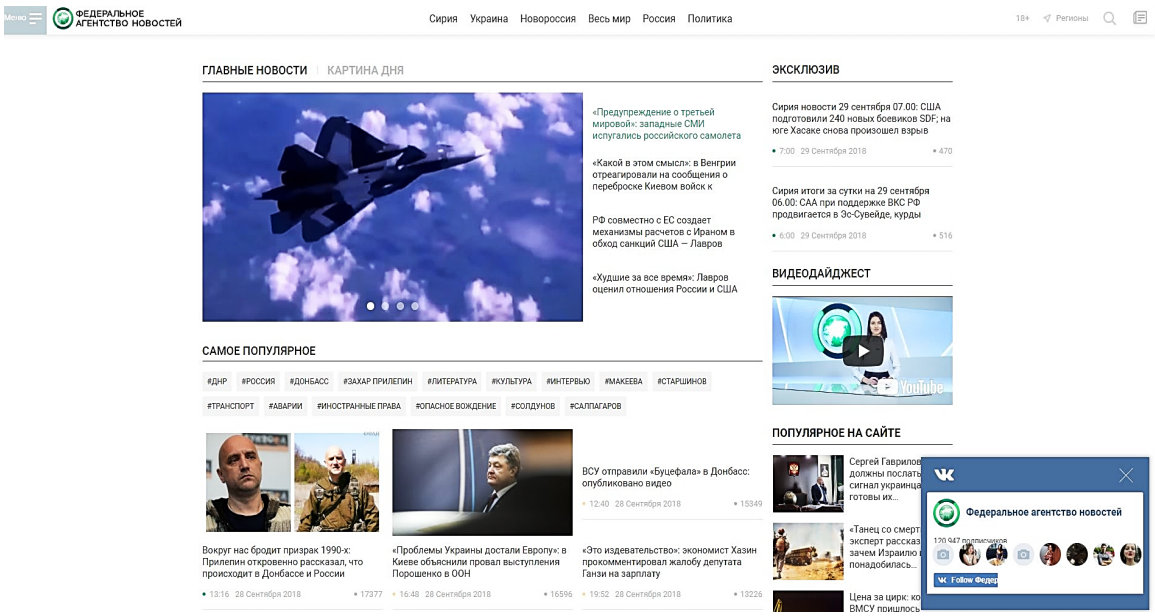

Figure 3. Screenshot of Riafan.ru from September 29, 2019 (https://riafan.ru/).

Collectively, these websites provided IRA sock puppets with additional noncensored, unlimited space to develop their arguments in more detail and to grow a more consistent followership that read both Twitter and their websites.

The second largest group of links tweeted by the IRA accounts, just over 9.5 thousand, led to traditional media outlets, including RT (Russia Today), TV channel Life News, and the RIA news agency. This strategy not only helped IRA accounts increase their legitimacy, but also revealed a two-pronged strategy between IRA and government-controlled mass media to obfuscate reality and disorient citizens.

[3] The name given by Donbass' separatists to their captured territory, which literally translates as "new Russia."

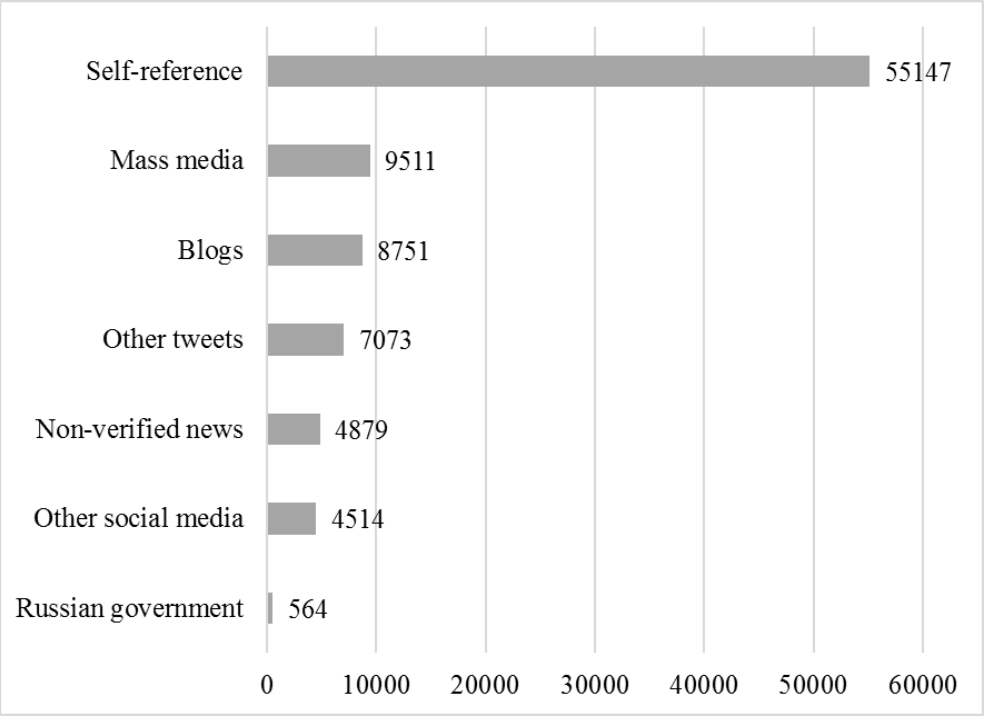

Figure 4. URLs categories used by IRA accounts.

Overall, this URL analysis reveals yet another *distraction* and *distortion* strategy of Kremlin's sock puppets: coordinated action to repeat the same message on different platforms (social media and websites) across hybrid media ecosystem to alter audience's perception of reality and to blur the distinction between fact and fiction. Through computational content analysis of IRA tweets, along with more qualitative analysis of hashtags and links included in those tweets, we demonstrate that IRA engaged in reflexive control strategy of dismissal, distraction, distortion, and denial. Our URL analysis also contributes an important addition to this online toolkit: interlinks between news websites and social media platforms to transmit the same message and increase the perceived legitimacy of disinformation.

Followers Count Analysis

Now that we determined what disinformation strategies were used by the Kremlin sock puppets, we study what types of content helped these accounts to gain more followers, thereby increase their visibility and credibility among other users. We predicted that including relevant hashtags, @mentions, and URLs helped prominent IRA-linked accounts gain followers. To test it, we first selected four of the most successful accounts in our data set based on the previous classification of user groups (Cha et al., 2012) and our content analysis of tweets and URLs. Two of these accounts are *mass-media,* or online news aggregators, and two others are *evangelists,* or opinion leaders. All these sock puppets quickly assembled or expanded an impressive army of followers, which visibly grew in the beginning of summer 2014 when the conflict in Donbass intensified (see Figure 5). Before the formal test of the hypothesis, let us first look at what type of hashtags, @mentions, and URLs were included in tweets of these four most successful accounts.

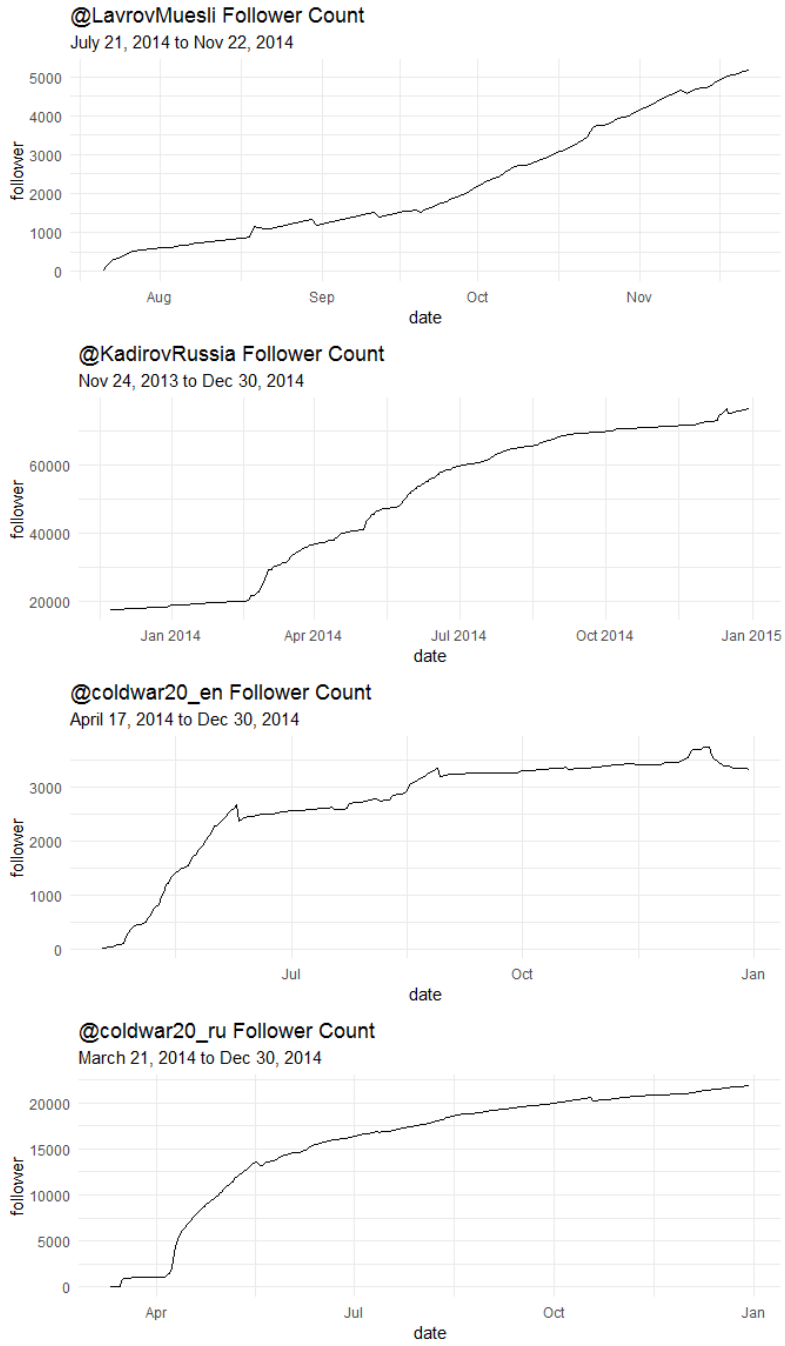

Figure 5. Follower growth of four IRA-affiliated accounts.

A Twitter account of the news aggregator Cold War 2.0, @coldwar20_ru, appeared in our data 454 times and was launched right after the annexation of Crimea. It often mentioned other nonverified news aggregators (e.g., @Pravdiva_pravda, @NOVORUSSIA2015) and Russian politician Konstantin Rykov (@rykov). These mentions suggest engagement with like-minded accounts and influential pro-Kremlin political figures, but also with international institutions like the UN. Almost half of shared URLs were directed toward its own tweets, and this account used 229 hashtags, mostly devoted to military actions in Ukraine.

The English-version @coldwar20_en mimicked strategies of its Russian-language counterpart: It was mentioned 80 times in our data. It also tried to engage with the UN institutions and BBC news by @mentioning them. Unlike the Russian version, @coldwar20_en included only 215 URLs, half of them leading to its own tweets, and only 58 hashtags, often related to the Eastern Ukrainian conflict.

Table 3. Top @mentions, URLs, and #Hashtags Used by Four IRA Accounts, 2013–14.

Handle Content features (count)	@mentions	URLs	#hashtags
coldwar20_ru	*Itself* (*N* = 401)	*Itself* on Twitter (*N* = 704)	Conflict in Donbass (*N* = 69)
	coldwar20_en (*N* = 57)	*Itself* on other social media platforms (*N* = 167)	Crime & Odessa (*N* = 40)
	Nonverified news (*N* = 218)	*Its* own webpage (*N* = 68)	Political actors (*N* = 29)
	Rykov (*N* = 58)	Russian news media (*N* = 146)	Ukraine & Kyiv (*N* = 28)
	UN (*N* = 168)		
coldwar20_en	UN (*N* = 49)	*Itself* on Twitter (*N* = 108)	Conflict in Donbass (*N* = 32)
	UNICEF (*N* = 35)	YouTube (*N* = 34)	
	Itself (*N* = 20)	Russian news media (*N* = 5)	
	BBCBreaking (*N* = 19)		
	coldwar20_ru (*N* = 19)		
KadirovRussia	Ru. opposition actors (*N* = 488)	*Itself* on Twitter (*N* = 32)	Ukraine & Kyiv (*N* = 2)

	TVRain (*N* = 119)	Russian news media (*N* = 20)	Donbass conflict (*N* = 2)
	Euromaidan (*N* = 80)	YouTube (*N* = 8)	
LavrovMuesli	Alexey Pushkov (*N* = 7)	*Itself* on Twitter (*N* = 28)	Ukraine & Kyiv (*N* = 1)
	Itself (*N* = 4)	Russian news media (*N* = 4)	# campaign (*N* = 1)
		YouTube (*N* = 2)	Political actors (*N* = 1)

In our corpus, @KadirovRussia was mentioned 74 times. Unlike the two other mass-media oriented accounts, this opinion leader acting as the head of the Chechen Republic engaged with Russian oppositional politicians. It mentioned three popular leaders: Alexey Navalny, Ksenia Sobchak, and Boris Nemtsov. This account also tagged the Russian oppositional TV Rain and Euromaidan Twitter. This account did not include many URLs—109 in our sample—and most of these links were to itself or to government-controlled Russian news media, like RT and LifeNews.

The account @LavrovMuesli was an opinion leader specializing in foreign policy at a nascent stage in our data. It was only mentioned six times and was mostly engaging with prominent politician Alexey Pushkov, Chairman of the Foreign Affairs Committee in Russian parliament. This account tweeted 54 URLs and used few hashtags.

Time-Series Analysis

To analyze how these content features influenced the follower count of individual accounts and directly test our hypothesis, we used univariate ARIMA models to identify autoregressive, integrated, and moving average components. These self-generating processes are removed from the time series ("prewhitened"), and the residuals are then used in an OLS regression.[4] This strategy is superior to using lagged variables (Reikard, 2009). The Bayesian information criterion (BIC) statistic was used to find optimal ARIMA models for each variable. Several follower count time series of individual handles had double unit roots, a relatively rare phenomenon (Harvey & Mills, 2002).

@coldwar20_ru

The analysis of @coldwar20_ru used a 295-day time series with 20 independent variables. Of these 20 series, eight had a unit root and were first-differenced. In this time, @coldwar20_ru posted 411 tweets. Table 4 displays the results of this analysis. Several key variables appear to influence @coldwar20_ru's

[4] We also tested a Prais–Winston model, which accounts for a 1–unit autoregressive component. Results from this analysis showed no improvement compared with the OLS model.

followership, including mentions of handles related to Euromaidan, other IRA accounts, and unverified news aggregators (see Figure 6). Tweeting links to Russian and Ukrainian media, and links to IRA content and other unverified news aggregators, also increased followers. Using hashtags increased follower count, particularly when they referred to the Donbass conflict and the annexation of Crimea, or to Kyiv, Ukraine, and Euromaidan. This model explained a fair amount (41%) of variance.

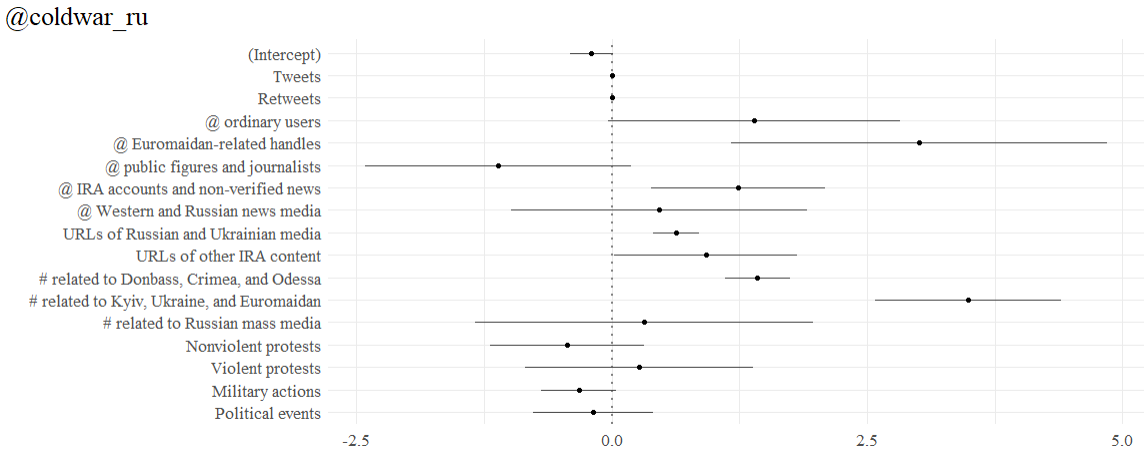

Figure 6. Factors influencing @coldwar20_ru's follower count.

@coldwar20_en

The analysis of @coldwar20_en relied on a 258-day time series with 14 independent variables (see Table 5). During this time, the account posted 57 original tweets. Two independent variables were first-differenced. The follower count of @coldwar20_en exhibited a fractionally integrated and a fully integrated component, producing a (0, 1.36, 1) ARFIMA model; both the integrated and fractionally integrated component were pre-whitened out. We found that little in our model explained @coldwar20_en's follower count (see Figure 7). The use of URLs from other IRA accounts increased follower count, but no other variable was statistically significant. As a result, the model's explanatory power was insignificant (1%).

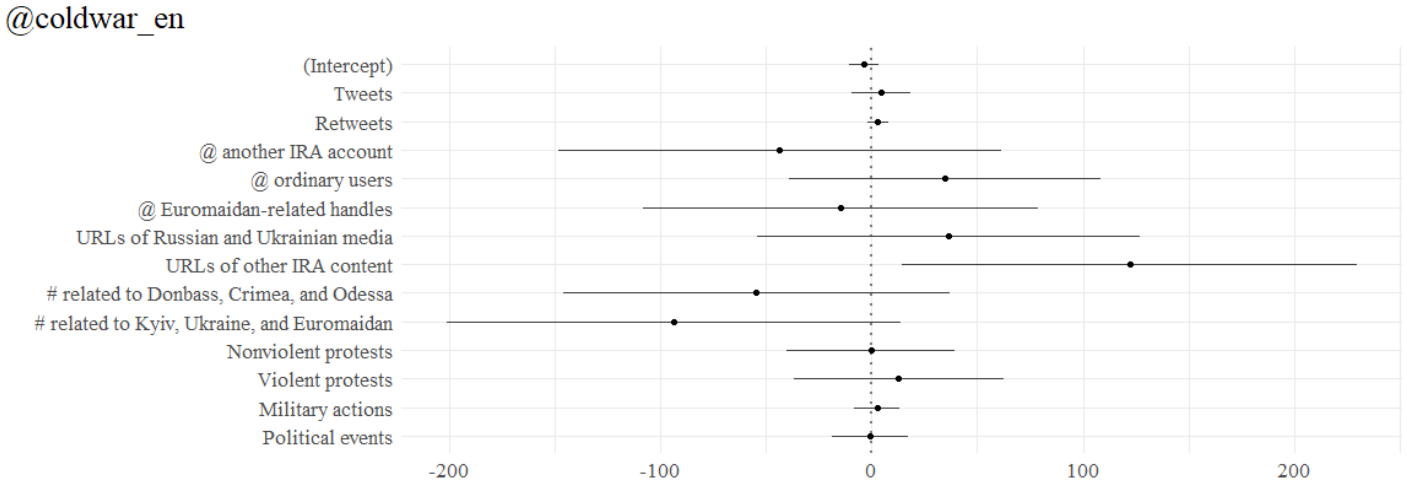

Figure 7. Factors influencing @coldwar20_en's follower count.

@KadirovRussia

The analysis of @KadirovRussia contained 403 day-units in which the account posted 190 tweets. In addition to a double-unit-root follower count (0, 2, 1), three independent variables were also integrated to the first order. Only mentions of other IRA accounts and tweets predicted follower growth (see Figure 8). This is reflected in the model's low explanatory power (6%).

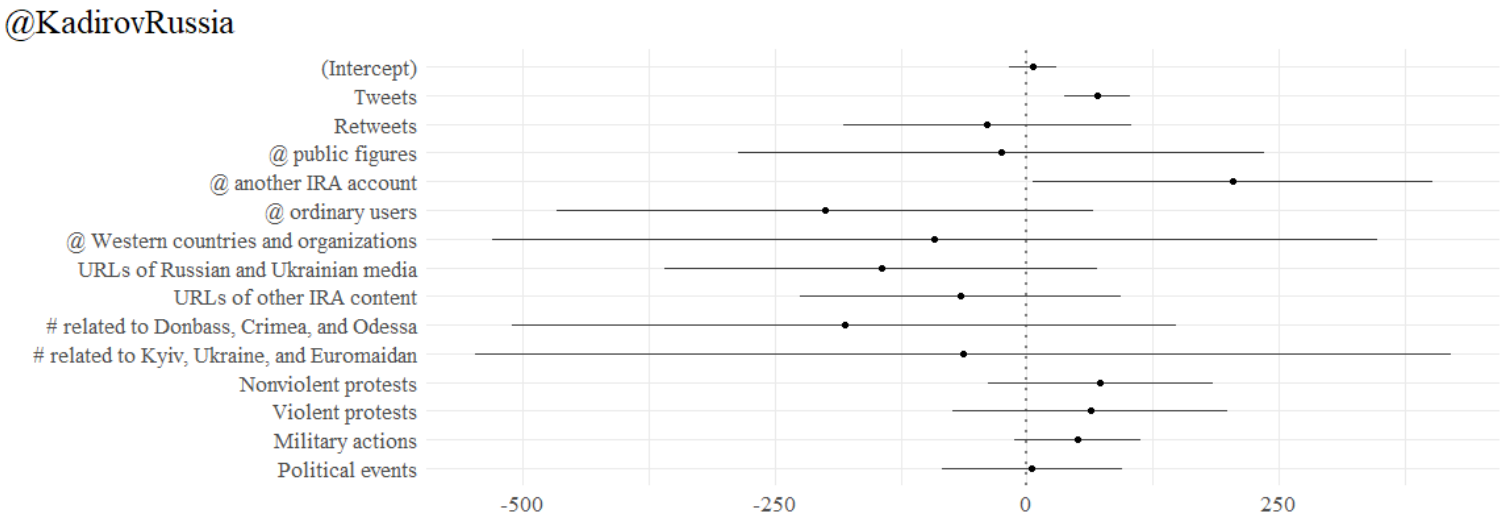

Figure 8. Factors influencing @KadirovRussia's follower count.

@LavrovMuesli

In the @LavrovMuesli model, only the follower count and three event variables had unit roots. This series was the shortest, with 125 day-units. During this time, the IRA handle posted 30 tweets. Here, tweets appeared to have a positive, but small, relationship. However, when military actions occurred, follower count decreased (see Figure 9). These variables accounted for much of the variance (69%).

These results overall provide support for our hypothesis by showing that the inclusion of relevant hashtags, @mentions, and URLs helped IRA-linked accounts grow their number of followers over our timeframe. The reflexive control strategies we discovered through content analysis did not only distract and dismay Twitter users—they helped IRA accounts gain attention and prominence. Notably, though, these strategies did not work equally for all accounts: Of the ones studied, only @coldwar20_ru was able to harness them to their full ability.

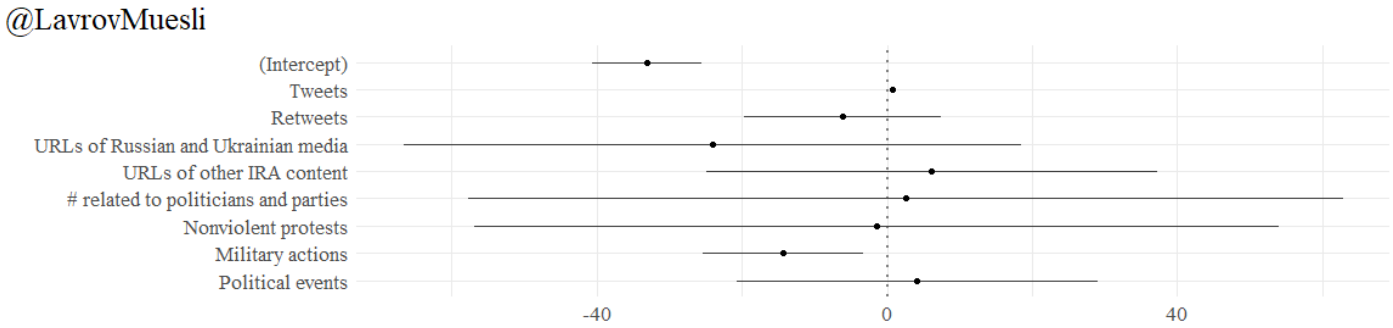

Figure 9. Factors influencing @LavrovMuesli's follower count.

Discussion

Our findings revealed that the computational propaganda tactics employed by the IRA against Ukraine was a multiplatform, reflexive control strategy employing news websites and social media to distort, distract, dismiss, and deny Russia's involvement in the conflict. Our analysis show how these tactics are an extension of traditional reflexive control strategies used during the Cold War to alter adversaries' key perceptions of reality and promote bad decision making. The IRA and its sock puppets managed to create a perception of active online news aggregators and opinion leaders, blurring the distinction between fact and fiction, news, and propaganda. In the predigital era, the Kremlin could only rely on government-controlled media for influencing citizens, but the Internet has made it easy to create propagandized websites promoted by social media accounts and blogs. This skillful combination of online platforms exploited the vulnerabilities of hybrid media system. This finding offered a more nuanced understanding to the finding that the IRA flooded Twitter and social media platforms with repeated messages (Mejias & Vokuev, 2017). As we show, it was not just mindless "flooding," but rather a deliberate and strategic campaign to gain attention while simultaneously distracting, disorienting, dismissing, and denying. The combination of this 4D approach and follower growth on Twitter should alarm communication scholars, as the further blurring of the borders

between professional fact-checking and conspiratorial thinking (Linvill et al., 2019) may have strategic benefits for malicious actors.

The content of IRA tweets not only contributed to reflexive control, but it also helped to grow the IRA accounts' army of followers. Our confirmed hypotheses revealed that tweeting more often, mentioning IRA-associated handles, linking to IRA-associated websites and blogs, and including hashtags about the Donbass conflict and Kyiv helped popular IRA accounts gain followers. Notably, we were able to explain a considerable amount of variance in @coldwar_ru's followership, which grew by over 20 thousand between 2013 and 2014. All three content features—URLs, #hashtags, and @mentions—mattered. Coordinated campaign with other IRA sock puppets and across several digital platforms helped this account to become a legitimate news source for its readers.

Taken together, these content strategies are characteristic of a hybrid warfare, with both media system and military implications. As perceptions of warfare are contingent on its portrayal, disinformation campaigns now consider social media, blogs, news aggregators, and all aspects of the media ecology system. Using reflexive control tactics to deny military involvement in Donbass and portraying it as a domestic issue, Russia was able to avoid direct Western interference and turned both the Crimea annexation and Donbass conflict into a stalemate, which the Kremlin can reactivate as needed for its foreign policy. Our findings may be helpful to others studying hybrid disinformation campaigns, including both "hot" ones, such as Ukraine or Syria, or "cold" one, like the 2016 U.S. election or Brexit.

There are several limitations worth noting. First, while most tweets collected were about Ukraine, our data set does have some noise: Russian tweets posted by IRA sock puppets that were not about Ukrainian conflict. However, we argue that keeping as close to a census as possible is important, as filtering by keyword would have eliminated other potential strategies. Our data also rely on Twitter's ability to distinguish IRA and non-IRA content. Though Twitter released the names of these accounts, it did not describe the processes used to identify them. Lastly, we do not consider the roles of bots, which might have been used to amplify messages from disinformation actors. While it is beyond the scope of analysis, which focuses on human-produced disinformation, there is also a reason to believe that this activity functions more as white noise than a systematic process. Furthermore, information about the followers of IRA accounts have not been made public by Twitter (and this data cannot be explored further as the accounts have been suspended). Finally, as human users themselves cannot determine whether a majority of a tweet's retweets are by bots, we therefore argue that bots are a natural part of social media communication networks.

Limitations notwithstanding, the analysis of Russian disinformation during Ukrainian conflict in 2013–14 provides key insights for scholars seeking to understand how the IRA adopted reflexive control strategies to the digital realm. Our work connects ongoing disinformation literature to the historic use of reflexive control (Thomas, 2004), particularly the 4D dimensions (Snegovaya, 2015). Furthermore, by selecting a case that Russia is both militarily and politically invested in, we can identify IRA strategies during an early "hot" campaign (Broniatowski et al., 2018). Our analysis reveals that not only is IRA able to translate reflexive control tactics in a digital ecosystem, but that the use of these tactics had benefits for its accounts' audience size, revealing insightful dynamics between content and follower count.

References

Aggarwal, A., Kumar, S., Bhargava, K., & Kumaraguru, P. (2018, April 9–13). The follower count fallacy: Detecting Twitter users with manipulated follower count. In *Proceedings of the 33rd Annual ACM Symposium on Applied Computing* (pp. 1748–1755). New York, NY: ACM.

Alexander, L. (2015, December 24). *Massive LiveJournal troll network pushes pro-Kremlin narratives.* Retrieved from www.stopfake.org/en/massive–livejournal–troll–network–pushes–pro–kremlin–narratives

Al-khateeb, S., & Agarwal, N. (2015, March 31–April 5). Analyzing deviant cyber flash mobs of ISIL on Twitter. In N. Agarwal, K. Xu, & N. Osgood (Eds.), *8th International Conference Social Computing, Behavioral–Cultural Modeling, and Prediction* (pp. 251–257). Cham, Switzerland: Springer.

Armstrong, L., & Gao, F. (2010). Now tweet this: How news organizations use Twitter. *Electronic News, 4*(4), 218–235. doi:10.1177/1931243110389457

Bastos, M., & Farkas, J. (2019). "Donald Trump is my president!": The Internet Research Agency propaganda machine. *Social Media + Society*, *5*(3), 1–13. doi:10.1177/2056305119865466

Bohdanova, T. (2014). Unexpected revolution: The role of social media in Ukraine's Euromaidan uprising. *European View, 13*(1), 133–142. doi:10.1007/s12290-014-0296-4

Catalano, R. A., Dooley, D., & Jackson, R. (1983). Selecting a time-series strategy. *Psychological Bulletin, 94*(3), 506–523. doi:10.1037/0033-2909.94.3.506

Cha, M., Benevenuto, F., Haddadi, H., & Gummadi, K. (2012). The world of connections and information flow in Twitter. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans, 42*(4), 991–998. doi:10.1109/TSMCA.2012.2183359

Chen, J., Haber, E., Kang, R., Hsieh, G., & Mahmud, J. (2015, May 26–29). Making use of derived personality: The case of social media ad targeting. In D. Quercia & B. Hogan (Eds.), *Ninth International AAAI Conference on Web and Social Media* (pp. 51–60). Palo Alto, CA: AAAI.

Chorley, J., Colombo, B., Allen, M., & Whitaker, M. (2015). Human content filtering in Twitter: The influence of metadata. *International Journal of Human–Computer Studies, 74*, 32–40. doi:10.1016/j.ijhcs.2014.10.001

Conover, M., Ratkiewicz, J., Francisco, M., Goncalves, B., Flammini, A., & Menczer, F. (2011, July 17–21). Political polarization on Twitter. In *Proceedings of the 5th International AAAI Conference on Weblogs and Social Media* (pp. 89–96). Palo Alto, CA: AAAI.

Cook, D., Waugh, B., Abdipanah, M., Hashemi, O., & Rahman, S. (2014). Twitter deception and influence: Issues of identity, slacktivism, and puppetry. *Journal of Information Warfare, 13*(1), 58–71.

European Commission. (2018, March 12). *Final report of the High Level Expert Group on Fake News and Online Disinformation.* Retrieved from https://digital-strategy.ec.europa.eu/en/library/final-report-high-level-expert-group-fake-news-and-online-disinformation

Fetzer, J. (2004). Disinformation: The use of false information. *Minds and Machines, 14*(2), 231–240. doi:10.1023/B:MIND.0000021683.28604.5b

Freelon, D., & Wells, C. (2020). Disinformation as political communication. *Political Communication, 37*(2), 145–156. doi:10.1080/10584609.2020.1723755

Garmazhapova, A. (2013, September 9). Gde zhivut troll: I kto ih kormit [Where trolls live: And who feeds them]. *Novaya Gazeta*. Retrieved from http://novayagazeta.spb.ru/articles/8093/

Giles, K., Sherr, J., & Seaboyer, A. (2018). *Russian reflexive control.* Kingston, Ontario: Royal Military College of Canada.

Golovchenko, Y., Hartmann, M., & Adler-Nissen, R. (2018). State, media and civil society in the information warfare over Ukraine: Citizen curators of digital disinformation. *International Affairs, 94*(5), 975–994. doi:10.1093/ia/iiy148

Gupta, A., & Kumaraguru, P. (2012, April 17). Credibility ranking of tweets during high impact events. In P. Kumaraguru & V. Almeida (Eds.), *PSOSM'12 Proceedings of the 1st Workshop on Privacy and Security in Online Social Media* (pp. 2–15). New York, NY: ACM.

Hayes, A. F., & Krippendorff, K. (2007). Answering the call for a standard reliability measure for coding data. *Communication Methods and Measures, 1*(1), 77–89. doi:10.1080/19312450709336664

Harvey, D., & Mills, T. (2002). Unit roots and double smooth transitions. *Journal of Applied Statistics, 29*(5), 675–683. doi:10.1080/02664760120098739

Hester, B., & Dougall, E. (2007). The efficiency of constructed week sampling for content analysis of online news. *Journalism & Mass Communication Quarterly, 84*(4), 811–824. doi:10.1177/107769900708400410

House Intelligence Committee. (2017). Exhibit B [List of IRA-linked Twitter accounts]. Retrieved from https://intelligence.house.gov/uploadedfiles/exhibit_b.pdf

Hutto, J., Yardi, S., & Gilbert, E. (2013, April 22–27). A longitudinal study of follow predictors on Twitter. In R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries, & G. Olson (Eds.), *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 821–830). New York, NY: ACM.

Internet Archive. (2014, August 13). *Web.Archive.org*. Retrieved from https://web.archive.org/web/2020*/cw20.ru

Jin, S.-A., & Phua, J. (2014). Following celebrities' tweets about brands: The impact of Twitter-based electronic word-of-mouth on consumers' source credibility perception, buying intention, and social identification with celebrities. *Journal of Advertising, 43*(2), 181–195. doi:10.1080/00913367.2013.827606

Koltsova, O., & Koltcov, S. (2013). Mapping the public agenda with topic modeling: The case of the Russian Livejournal. *Policy & Internet, 5*(2), 207–227. doi:10.1002/1944-2866.POI331

Lahuerta-Otero, E., & Cordero-Gutierrez, R. (2016). Looking for the perfect tweet: The use of data mining techniques to find influencers on Twitter. *Computers in Human Behavior, 64*, 575–583. doi:10.1016/j.chb.2016.07.035

Leonenko, S. (1995). Refleskivnoye upravlenie protivnikom [Reflexive control of an adversary]. *Armeyskiy Sbornik, 8*, 27–32.

Levefr, V., & Smolyan, G. (1968). *Azbuka Konflikta* [ABCs of conflict]. Moscow, Russia: Znanie.

Linvill, L., Boatwright, C., Grant, J., & Warren, 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

Lukito, J., Suk, J., Zhang, Y., Doroshenko, L., Kim, S. J., Su, M.-H., . . . Wells, C. (2020). The wolves in sheep's clothing: How Russia's Internet Research Agency tweets appeared in U.S. news as vox populi. *International Journal of Press/Politics, 25*(2), 196–216. doi:10.1177/1940161219895215

Maier, D., Waldherr, A., Miltner, P., Wiedemann, G., Niekler, A., Keinert, A., . . . Adams, S. (2018). Applying LDA topic modeling in communication research: Toward a valid and reliable methodology. *Communication Methods and Measures*, *12*(2/3), 93–118. doi:10.1080/19312458.2018.1430754

Marwick, E. (2013). *Status update: Celebrity, publicity, and branding in the social media age.* New Haven, CT: Yale University Press.

Mejias, A., & Vokuev, E. (2017). Disinformation and the media: The case of Russia and Ukraine. *Media, Culture & Society, 39*(7), 1027–1042. doi:10.1177/0163443716686672

Onuch, O. (2015). Euromaidan protests in Ukraine: Social media versus social networks. *Problems of Post-Communism, 62*(4), 217–235. doi:10.1080/10758216.2015.1037676

Parker, J. (2016). *Hashtag games make Twitter users get creative*. Retrieved from https://www.cnet.com/news/hashtag–games–make–twitter–users–get–creative/

Pomerantsev, P. (2015). *Nothing is true and everything is possible: The surreal heart of the new Russia.* New York, NY: Public Affairs.

Reikard, G. (2009). Predicting solar radiation at high resolutions: A comparison of time series forecasts. *Solar Energy, 83*(3), 342–349. doi:10.1016/j.solener.2008.08.007

Roeder, O. (2018, July 31). *Why we're sharing 3 million Russian troll tweets*. Retrieved from https://fivethirtyeight.com/features/why-were-sharing-3-million-russian-troll-tweets/

Roscomnadzor. (2017). *Registracia saitov v informacionno-telekommunikacionnoi seti Internet v kachestve sredstv massovoi informacii* [Registration of websites on Internet as mass media] [PowerPoint slides]. Retrieved from https://rkn.gov.ru/docs/SEMINAR_po_setevym_izdanijam_na_sajt.pdf

Saxton, D., & Guo, C. (2014). Online stakeholder targeting and the acquisition of social media capital. *International Journal of Nonprofit and Voluntary Sector Marketing, 19*(4), 286–300. doi:10.1002/nvsm.1504

Snegovaya, M. (2015). *Putin's information warfare in Ukraine: Soviet origins of Russia's hybrid warfare.* Washington, DC: Institute for the Study of War.

Surzhko-Harned, L., & Zahuranec, A. J. (2017). Framing the revolution: The role of social media in Ukraine's Euromaidan movement. *Nationalities Papers, 45*(5), 758–779. doi:10.1080/00905992.2017.1289162

Thomas, L. (2004). Russia's reflexive control theory and the military. *Journal of Slavic Military Studies, 17*(2), 237–256. doi:10.1080/13518040490450529

U.S. v. Internet Research Agency LLC, 18 U.S.C. (2018).

Walther, B., Van Der Heide, B., Kim, Y., Westerman, D., & Tong, T. (2008). The role of friends' appearance and behavior on evaluations of individuals on Facebook: Are we known by the company we keep? *Human Communication Research, 34*(1), 28–49. doi:10.1111/j.1468-2958.2007.00312.x

Wang, Y., Luo, J., & Zhang, X. (2017, September 13–15). When follow is just one click away: Understanding Twitter follow behavior in the 2016 U.S. presidential election. In G. Ciampaglia, A. Mashhadi, & T. Yasseri (Eds.), *Proceedings of the 9th International Conference on Social Informatics* (pp. 409–425). Cham, Switzerland: Springer.

Webster, G. (2014). *The marketplace of attention: How audiences take shape in a digital age.* Cambridge, MA: MIT Press.

Wells, C., Shah, D. V., Pevehouse, J. C., Foley, J., Lukito, J., Pelled, A., & Yang, J. (2019). The temporal turn in communication research: Time series analyses using computational approaches. *International Journal of Communication, 13*, 4021–4043.

Woolley, S., & Howard, P. (2016). Political communication, computational propaganda, and autonomous agents. *International Journal of Communication, 10*, 4882–4890.

Yang, J., & Counts, S. (2010, May 23–26). Predicting the speed, scale, and range of information diffusion in Twitter. In M. Hearst, W. Cohen, & S. Gosling (Eds.), *Proceedings for the 4th International AAAI Conference on Weblogs and Social Media* (pp. 355–358). Menlo Park, CA: AAAI.

Zannettou, S., Caulfield, T., De Cristofaro, E., Sirivianos, M., Stringhinin, G., & Blackburn, J. (2019, May 13–17). Disinformation warfare: Understanding state-sponsored trolls on Twitter and their influence on the Web. In L. Liu & R. White (Eds.), *Companion Proceedings of the 2019 World Wide Web Conference* (pp. 218–226). New York, NY: ACM.

Zhang, Y., Lukito, J., Su, M.-H., Suk, J., Xia, Y., Kim, S. J., . . . Wells, C. (2021). Assembling the networks and audiences of disinformation: How successful Russian IRA Twitter accounts built their followings, 2015–2017. *Journal of Communication, 71*(2), 305–331. doi:10.1093/joc/jqaa042