

Understanding RT's Audiences: Exposure Not Endorsement for Twitter Followers of Russian State-Sponsored Media

The International Journal of Press/Politics
2022, Vol. 27(1) 220–242
© The Author(s) 2021



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/1940161220980692
journals.sagepub.com/home/hij



Rhys Crilley¹ , Marie Gillespie², Bertie Vidgen³,
and Alistair Willis²

Abstract

The Russian state-funded international broadcaster RT (formerly Russia Today) has attracted much attention as a purveyor of Russian propaganda. To date, most studies of RT have focused on its broadcast, website, and social media content, with little research on its audiences. Through a data-driven application of network science and other computational methods, we address this gap to provide insight into the demographics and interests of RT's Twitter followers, as well as how they engage with RT. Building upon recent studies of Russian state-sponsored media, we report three main results. First, we find that most of RT's Twitter followers only very rarely engage with its content and tend to be exposed to RT's content alongside other mainstream news channels. This indicates that RT is not a central part of their online news media environment. Second, using probabilistic computational methods, we show that followers of RT are slightly more likely to be older and male than average Twitter users, and they are far more likely to be bots. Third, we identify thirty-five distinct audience segments, which vary in terms of their nationality, languages, and interests. This audience segmentation reveals the considerable heterogeneity of RT's Twitter followers. Accordingly, we conclude that generalizations about RT's audience based on analyses of RT's media content, or on vocal minorities among its

¹University of Glasgow, Glasgow, UK

²The Open University, Milton Keynes, UK

³The Alan Turing Institute, London, UK

Corresponding Author:

Rhys Crilley, Department of Politics and International Relations, University of Glasgow, The Adam Smith Building, Glasgow G12 8RT, UK.

Email: Rhys.crilley@glasgow.ac.uk

wider audiences, are unhelpful and limit our understanding of RT and its appeal to international audiences.

Keywords

RT (Russia Today), Twitter, Russian propaganda, audience research, computational analysis

Introduction

The Russian state-funded international broadcaster RT (formerly Russia Today) purportedly aims to provide a Russian perspective on events and issues to international audiences. In doing so, RT blurs the lines between news reporting, propaganda, and soft power (Wright et al. 2020: 2), and has, since its creation in 2005, become the subject of much controversy—especially in the “Western” states in which it operates and seeks to exert influence. The controversies vary in different parts of the world. In America, the intelligence services stated that, throughout the 2016 presidential elections, RT was central to Russian “influence efforts to denigrate Secretary Clinton” (Office of the Director of National Intelligence 2017). As a result, RT has had to register there as a foreign agent and has also been banned from using paid advertisements on some social media platforms. In France, RT has been accused of promoting protests by the “Yellow Vests,” leading President Macron to draft legislation to ban media that are considered to “de-stabilise the country” (Newman et al. 2019: 84). In the United Kingdom, RT was fined £200,000 by the communication regulator Ofcom for impartiality breaches in their reporting on the Skripal poisoning and Syrian conflict (Elliot 2019).

Most previous academic work has focused on the content that RT produces and its role in the international media environment. Studies have found that RT courts conspiracy theories (Yablokov 2015), propagates mis/disinformation (Cull et al. 2017; Ramsay and Robertshaw 2019), disseminates anti-Semitism (Rosenberg 2015), promotes Islamophobia (Lytvynenko and Silverman 2019), and has produced and amplified media content that supported the election of Donald Trump (Jamieson 2018). Far less attention has been paid to RT’s online audiences, such as the users who consume and are exposed to its content (cf. Crilley and Chatterje-Doody 2020a; Crilley et al. 2019; Fisher 2020a). More broadly, Szostek (2018: 117) notes that studies of Russian state-funded international news media are only “rarely informed by substantive research into the behaviour and thought processes of news consumers” themselves. Indeed, few analyses have empirically investigated how RT’s audiences engage with and interpret their content.

Our study examines the validity of many existing assumptions about RT’s audience through a data-driven analysis of a large Twitter dataset. The paper proceeds in five parts. First, we review the literature on RT and characterize it as commuting across three main operational paradigms: propaganda, alternative media, and soft power. Second, after noting the limited number and scope of audience research

studies on RT, as well as a tendency to assume that audiences and RT mirror each other, we raise some key questions about RT's audiences and Twitter followers. Third, we present our data on RT's main international Twitter account @RT_com. Fourth, we proceed to a discussion of the results and suggestions for future research before. Fifth, we draw conclusions.

Our study makes two interlinked arguments that contribute toward a better understanding of RT's audiences. We find that, first, RT Twitter followers rarely engage with RT content, and therefore claims that RT has a large audience that supports its "anti-Western" worldview are misguided (Cohen 2014). RT's Twitter followers may be exposed to its content, but rarely do they appear to endorse it. Second, through computational network analysis, we find that RT's Twitter followers do not appear to be "a niche audience of activists" (Orrtung and Nelson 2019: 14). Rather, the majority of RT's Twitter followers follow RT alongside many other major sources of international news. RT's Twitter audiences are fragmented mainly along lines of national, linguistic, and cultural interests rather than by political identities or ideologically extreme views.

RT's Operational Paradigms: Reviewing the Literature

Despite the political attention it has received, research on RT is patchy, including academic studies (Hutchings et al. 2015; Miazhevich 2018), think-tank reports (Bodine-Baron et al. 2018; Richter 2017), and journalistic pieces (Dowling 2017). Most research has focused on the role and objectives of RT, with some suggesting that it is a key component of Putin's propaganda machine (Cull et al. 2017). Others argue that it is a public diplomacy resource, only mobilized to augment Russian soft power when needed (Rawnsley 2015). RT is also understood to operate as an "alternative" media channel which lacks the legitimacy of legacy broadcasters such as the BBC World Service, Deutsche Welle, and France Médias Monde (Rawnsley 2015). RT is generally held in widespread disrepute and mainly criticized as a nefarious tool of the Russian state. Yet, as Galeotti (2018) convincingly argues, there is "no single organizing principle, let alone controlling agency" that manipulates RT as an organization or dictates the agenda for the journalists who work for it. Just because RT positions itself at times as an alternative to, and an opponent of, "mainstream media," this does not mean that it operates exclusively as "Putin's puppet" or a simplistic propaganda agent for Russia's neo-authoritarian regime (Chatterje-Doody and Tolz 2020; Dajani et al. 2019; Galeotti 2018; Tolz et al. 2020).

In much of the literature, RT is judged to be overtly anti-American, aiming to undermine the U.S.-led global order (Elsawah and Howard 2020) in an attempt to "justify Russian government policies and create an image of Russia as the leader of global resistance to the US" (Yablokov 2015: 312). RT, it is argued, thrives through its constant attacks against the hypocrisy of the "West," and support for Putin's leadership (Miazhevich 2018: 6), while projecting an image of Russia as "a rapidly advancing nation threatening (though not succeeding) to disturb the hegemonic balance of power" (Hutchings et al. 2015: 641). Similarly, Kragh and Åsberg (2017) find that RT portrays

Russia as a major global power which enjoys the support and cooperation of other international actors.

As an international news provider, RT's greatest success has been on social media, particularly YouTube, where RT was the first news channel to reach over a billion views. RT's most successful YouTube videos rarely touch on political or ideological matters. Rather, their focus has been on sensational, eye-witness reports of catastrophes and disasters, for example, user-generated footage of the 2011 Japanese Tsunami (Chatterje-Doodly and Crilley 2019b; Mickiewicz 2018). In fact, according to figures quoted by Mickiewicz, political videos count for only 1 percent of RT's YouTube content. Perhaps exposure to RT is limited not only in size but also in content—a key focus of our empirical investigations.

Following RT on Twitter

Increasingly, there is recognition that attention needs to be shifted away from RT as a broadcaster to its online audiences. In particular, Mickiewicz (2018) urges scholars to undertake more robust and detailed audience research—which is much needed given the lack of robust evidence informing current debates. Audience research is particularly important for developing a better understanding of RT and its distinctive operational paradigms. However, investigating RT's audiences is a difficult task. Hutchings et al. (2015: 653) argue that “we still know precious little about the way in which international audiences interact with news organisations like RT” and Yablokov similarly contends that it is “hard to define [RT's] audience/s and the efficiency of its message” (Hutchings et al. 2015: 310).

Notably, RT's audiences are frequently viewed with the same disdain directed at RT itself and are (at least implicitly) considered to be ideologically closely aligned with RT's political framing on issues and events (Birrell 2018). For example, according to the journalist Nick Cohen (2014), RT has a “huge western audience that wants to believe that human rights are a sham and democracy a fix.” Equally, Turner (2016) suggests that RT has “a large and disillusioned western audience.” Yet the assertion that RT has a “huge” or “large” audience is not borne out by empirical research. Mickiewicz (2018: 1–13) draws on data from the Neilson reports to suggest that the size of RT's TV audiences is small to negligible in America, while research from IPSOS suggests that RT's television audiences are also extremely small in Western European countries and are not growing except in the Middle East, and in Syria and Iraq particularly (Ipsos Connect 2018). Online, RT has a larger following than on TV, but it is still relatively small compared with other state-backed broadcasters (Al-Rawi 2017). Part of the challenge for researchers of international broadcasting, propaganda, and soft power is that estimating RT's true audience is difficult. For instance, Yablokov (2015: 311) argues that “it is virtually impossible to measure the channel's success and influence.”

Understanding the size of RT's audience is only one part of the puzzle—of equal importance is understanding who those audiences are. A small number of previous

studies have attempted to address this problem. Miazhevich (2018: 3) suggests that if one is to go by its narrative frames, RT aims “to appeal to audiences who have an anti-establishment, anti-corporation and anti-western (particularly anti-American) predisposition,” an argument also expressed by others (Elliot 2019; Newman et al. 2019; Ramsay and Robertshaw 2019: 24). Orttung and Nelson’s (2019: 1) analysis of YouTube videos published by RT supports this view and suggests that RT has a “tripartite strategy,” underscoring its distinctive operational paradigms: “primarily targeting an audience outside the West; working to complement local media in target countries; and pushing narratives to promote a positive image of Russia.” They also argue that “RT seems to have found a niche audience of activists. Unsurprisingly, its viewers, like for most international news channels, tend to skew toward highly educated males” (Orttung and Nelson 2019: 14). A study of YouTube comments made in response to RT videos about the Syrian conflict found that the sample of audiences analyzed tended to support the broadcaster’s characterization of the conflict and expressed support for Russia alongside anger and mistrust of the “West” (Chatterjee-Dooddy and Crilley 2019a: 174). Richter (2017: 13) characterizes RT’s audience as comprised of people who support conspiracy theories because they are “distrustful” of “mainstream media.” Similarly, Birrell (2018) suggests that RT’s online audiences are vulnerable to the “malign force” of its media content, and are situated on the fringes of mainstream politics, attracted to niche interests and incendiary ideas, as well as anti-American and anti-“Western” discourse.

Part of the problem here is that relatively little work has focused directly on RT’s audiences and, instead, has attempted to infer the interests, outlooks, and traits of RT’s audience based on the content that it broadcasts. This is a fundamentally flawed approach as there is no intrinsic reason why the perspectives of RT’s audiences would necessarily mirror those of the broadcaster itself. As Mickiewicz (2018: 3) notes, “content is one big piece of the puzzle, but one cannot argue from content alone back to the audiences’ intake.” Furthermore, we caution that the results of many previous studies are based on the behavior of only a very vocal minority (i.e., those who actually engage with RT’s content), and therefore their attitudes and beliefs are unlikely to be representative of their entire audience. There is, therefore, a pressing need to understand RT’s audiences more holistically, especially to understand whether they are truly as extreme and niche as is often suggested. In particular, it is unclear from previous work whether the large audiences that RT’s content can potentially reach have any genuine interest in supporting RT’s “anti-Western” perspective. Here, data-driven approaches which can analyze large numbers of followers, as opposed to smaller scale qualitative research, can help to understand the nuances of RT’s broader audience. Subsequently, in this study, we investigate the extent to which the average RT Twitter follower engages with RT’s content. Alongside this, we use computational network science to ascertain the shared demographic features and interests of RT’s Twitter followers to understand whether they are as politically extreme as others have claimed.

Researching RT's Audience on Twitter

The need to better understand RT's audiences raises the more conceptual question of what an audience is. The changing aspects of how, when, where, and by whom international news is consumed are re-shaping audience configurations, and subsequently, social science research on audiences is undergoing change too (Gillespie and Webb 2012; Rogers 2013). We can no longer assume that audiences consume news in conventional ways or places. News is now often consumed on a smartphone, and attention paid to any one news item can be short and superficial. Indeed, audiences often receive news from multiple platforms and sources in a "hybrid media system" where the cultural power to produce, interpret, and respond to news shifts more fluidly between users, producers, and political actors (Chadwick 2017).

Online technologies have radically altered social networks and modes of engagement, leading to a proliferation of weak ties and informal modes of interaction. On Twitter, following an account is financially cost-free and logistically easy. It can also be done without reciprocation, which means that it is not dependent upon anyone else (as with two-way ties on platforms such as Facebook). Twitter followers are not equivalent to a traditional broadcast audience, nor even to the engaged "active" audiences who comment on YouTube videos (Chatterje-Doody and Crilley 2019a; Crilley and Chatterje-Doody 2020a, 2020b). Indeed, a user might follow an account but never engage with it—and simply forget to unfollow it. Just because someone may follow an account (and are therefore likely to be exposed to its content), they might not endorse it (Marwick 2018).

Yet following an account is a useful indication of interest because, although it does not entail a financial cost, users have to choose which accounts they follow, and, in the so-called "attention economy," following an account involves "spending" attention (Harsin 2015). As each user's timeline is algorithmically populated, partly based on which accounts they follow, following an account means a user (that is not a bot or an inactive account) is likely to be exposed to that account's content which could, in turn, not only resonate with prior worldviews but also inform and sustain their beliefs, preferences, and outlook by appealing to deep stories (Marwick 2018). Such processes of algorithmic content surfacing have been criticized for creating and reinforcing biases and injustice online, and for potentially encouraging users to engage in harmful behaviors (Caplan et al. 2018). We also note that which accounts users follow has been successfully used in studies to infer ideological positions, indicating that followership contains important political and social signals (Barberá 2015). Furthermore, researching RT's Twitter followers provides insight into the "big and broad" audiences who are likely to be exposed to its content, extending research beyond the narrow group of highly vocal and highly engaged, but also most likely atypical, users who comment on or reply to its content. This is particularly important given that many online users are "lurkers" who passively consume content, rarely if ever engaging with it (Sun et al. 2015).

Overall, insufficient attention has been paid in previous work to who is exposed to, and consumes, RT's content. Current accounts are either very narrow, lack nuance and

detail, or are not routed in empirical evidence. To test assumptions about audiences outlined in the literature above, and to better understand RT's audiences, we conduct empirical analyses to answer three key questions:

- **Research Question 1:** To what extent does the average Twitter follower of RT engage with its content?
- **Research Question 2:** What are the key demographic features of RT's Twitter followers?
- **Research Question 3:** What are their main interests of RT's Twitter followers?

Data

We study Twitter, due to its major significance as a news sharing platform. Twitter allows information about its users to be accessed through their Application Programming Interface (API). However, restrictions are also imposed on the amount of data and the frequency by which it can be collected.¹ Our analysis is therefore based on a sample of the 2.6 million users who follow RT's main international Twitter account (@RT_com). We selected a random sample of users, as well as information about the other accounts they follow. Our data collection proceeded in several steps. First, we collected a list of every user which follows RT ($n = 2.6$ million). Second, we randomly selected a sample of users from this list of 2.6 million. Third, for each user in our sample, we collected all of the accounts which they follow.

Power tests indicate that for a two-sided effect in a one-sample t -test, 10,000 users would be capable of detecting effects as small as 0.042, with alpha of .01 and power of 0.95. We use this as a heuristic for sampling users from Twitter, given that this would allow for even very small effects to be detected. We initially sampled 12,151 followers of RT; 1,490 accounts (12.3 percent) were set to private/protected (and so data were not available) and 270 accounts (2.22 percent) had been deleted during the period of data collection. The rate of 12.3 percent is toward the higher end of estimates of private accounts on Twitter; Liang et al. found that the rate of private accounts varied by country, with a range of between 1.4 and 16.9 percent and an overall mean of 5.4 percent (Liang et al. 2017). In total, 85.5 percent of the accounts we collected were available for analysis, leading to a sample of 10,391 users (above our target size of 10,000).

For the 10,391 users in our sample, we collected their profile information and the accounts that they follow, creating a followership network with 6.48 million connections in total. Each user in our sample follows, on average, 624 other accounts. We also collected all of their tweets on November 25, 2019. In line with Twitter's API constraints, we collected the last 3,200 tweets for each user. However, while for most users this provides full historical coverage, for a small number of highly active users it only covers a small window of time. To account for this, we limited the earliest date of the tweets to November 25, 2018, thereby giving a one-year period of data

coverage. Less than 5 percent of the users in the dataset hit the 3,200 tweet limit, which indicates our data collection has strong overall coverage. The dataset comprises 1.87 million tweets.

In addition to this, we collected a further random sample of 1,000 Twitter users to serve as a control group, which enables us to understand how our sample of RT followers differs from other users more generally, with a high level of statistical power. The random sample was identified by collecting 100,000 tweets from the Twitter stream, searching for users who used the term “News,” on November 1, 2019. We then randomly sampled 1,000 unique users from these tweets. We use a “general” sample of users so as to understand RT’s followers in the broader context of Twitter, rather than biasing interpretation of the results by focusing on a purposive sample, such as followers of another news provider. We collected metadata for 934 of these users (93.4 percent) as the remainder had protected their accounts. None of the users in this random sample engage with or follow RT.

Method

Social network analysis has been widely used to understand social media users. For Twitter, the audience is typically represented as a graph, or network, with nodes representing user accounts, and directed links representing follower relationships. To identify different segments of users within our 10,391 sample, we use the concept of modularity. In the case of a Twitter network, a modular collection of nodes in the network is a set of users which have more links between each other than links with other users (i.e., they have more links within the set than leading out of it). In practice, this means that they tend to follow many accounts, which are the same. A widely used algorithm for identifying modular groups is the Louvain algorithm (Blondel et al. 2008). We applied the Louvain algorithm to the sampled Twitter network using the `python-louvain`² library to identify distinct segments of users.

To characterize each of the user segments identified by the Louvain algorithm, we then applied a surfacing methodology, which let us identify the “most characteristic” accounts followed by each segment. These are accounts which are followed by an unexpectedly large number of users from each segment. The intuition behind this approach is that it is not particularly surprising that in every segment a large number of users follow Barack Obama because he is one of the most followed accounts on Twitter. However, the fact that in some segments many users follow some fairly obscure Twitter accounts is much more informative. This methodology lets us surface these unexpectedly well-followed accounts for each segment, providing a way of differentiating between them in a very granular way. The methodology is based on previous work in network science to characterize online communities, as well as insights about the long-tailed distribution of social media followership in general (Cha et al. 2010), and it is a novel application of established statistical procedures in this context (Fog 2007).

There are five steps to our method. First, the 10,391 users in our sample follow 3.03 million unique accounts in total. For each of these 3.03 million unique accounts, we

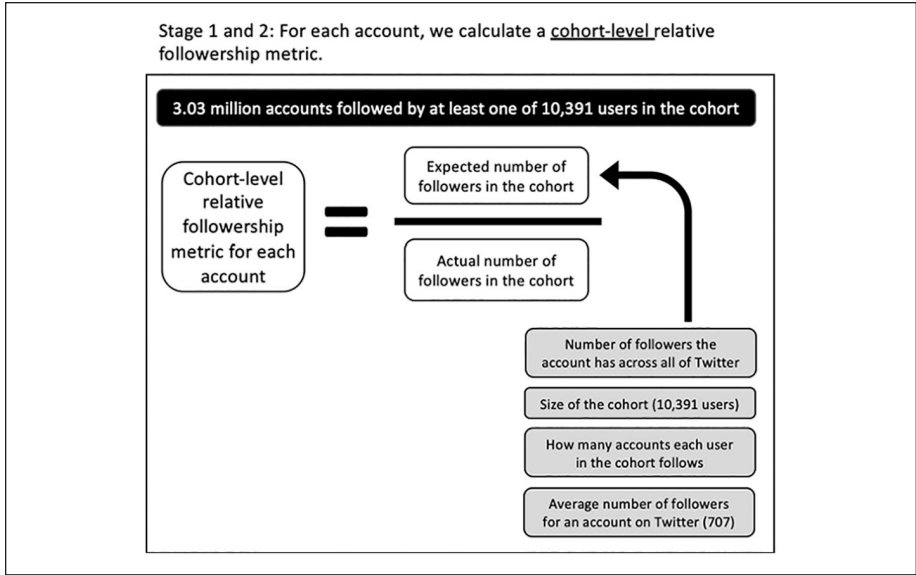


Figure 1. Cohort-level relative followership metric.

calculate the expected number of followers from our sample using Wallenius’s non-central hypergeometric distribution. This takes into account (1) the number of followers the account has across all of Twitter, (2) how many accounts each user in the sample actually follows (i.e., how many “follows” they have to allocate), (3) the size of the cohort ($n = 10,391$), and (4) the average number of followers for all accounts on Twitter, which we estimate is 707, based on figures provided by the company Brandwatch.³ Wallenius’s distribution is a variation of the hypergeometric distribution (a discrete probability distribution which is similar to the binomial distribution), which describes sampling of items without replacement (Fog 2007). Here, the “items” are followers. Wallenius’s explicitly accounts for the fact that a randomly selected user is statistically more likely to follow an account with many followers on Twitter than an account with very few followers. Second, the expected number of followers for each of the 3.03 million unique accounts is divided by their actual number of followers to calculate a cohort-level “relative followership” metric. This gives us a better understanding of the popularity of accounts within our cohort given their popularity on Twitter overall. At this stage, all calculations are for the cohort, and the segmentation is not taken into account. This is shown in Figure 1.

Third, the same relative followership metric is calculated for each of the followed accounts in each segment. To calculate the segment-level relative followership metric for each of the followed accounts, the numerator is the expected number of followers *within that segment* and the denominator is the actual number of followers *within that segment*. Note that the expected number of followers for an account in each segment can be a very small fraction, such as 0.001 followers. This is shown in Figure 2.

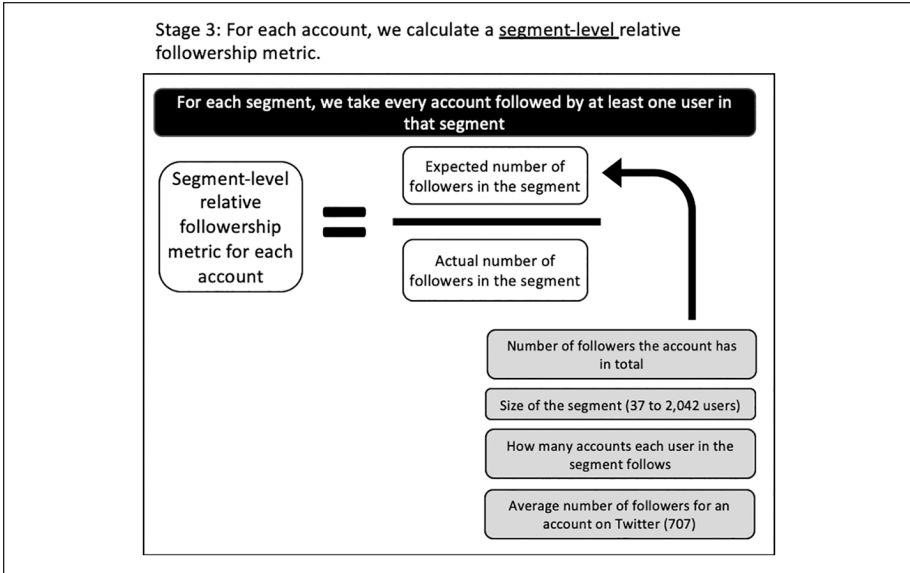


Figure 2. Segment-level relative followership metric.

Fourth, a scaling factor is then calculated for each followed account within each segment by dividing the segment-level relative followership metric by the cohort-level relative followership metric. This scaling factor gives us a better understanding of the popularity of accounts within each segment given their popularity within the cohort overall. Finally, for each segment, we take the accounts with the largest scaling factors. One limitation of the method is that accounts with very few followers can have incredibly high scaling factors, depending on how those followers are distributed. To ensure that the method returns accounts which are typical of the segment, we set a minimum threshold: Only accounts which are followed by at least 10 percent of users in each segment can be considered characteristic for that segment.⁴ For our analyses, we take the 10 accounts with the highest scaling factors which are above the 10 percent limit (Figure 3).

Results

Audience Engagement with RT

We examine how users in our 10,391 sample engage with RT. Our findings indicate that engagement with RT is low: There are only 2,806 engagements with RT, which we count as a retweet (2,371 instances) or a reply/mention (913 instances); this accounts for just 0.15 percent of the 1.87 million tweets sent by our sample. The distribution of the number of engagements with RT per user is long-tailed which means that a small number of users drive most of the engagement, as shown in Figure 4. This

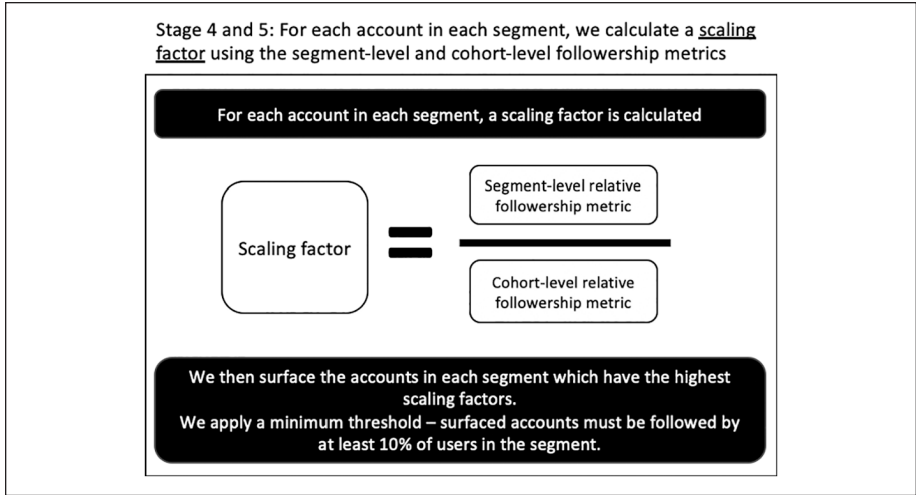


Figure 3. Scaling factor for each account in each segment.

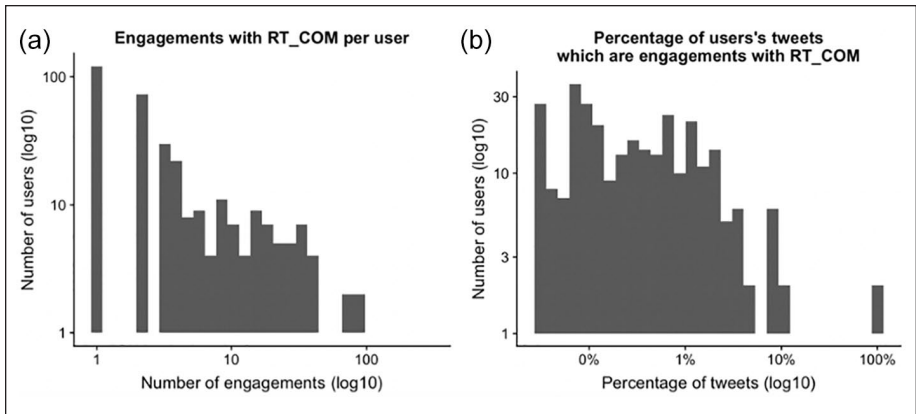


Figure 4. (a) Distribution of engagements with RT_COM per user and (b) distribution of the percentage of user’s tweets which are engagements with RT_COM.

is typical of online behavior in general. The most engaged user accounted for 284 (10 percent) of all the 2,806 engagements with RT and the 10 most engaged users accounted for 1,133 (40.4 percent) of all the engagements. Only 325 users of the 10,391 users engaged with RT at least once (3.23 percent). However, it should be noted that during the year we study, only 3,395 of the 10,391 users in our sample send any tweets at all (33 percent). Of these users, 9.87 percent engage with RT, which is a larger but still small proportion. Figure 4(b) shows that while engagements with RT

are a small percentage of most user's tweets, for two users it accounts for 100 percent of their activity. Both are low-volume users who send very few tweets.

The Demographics of RT's Twitter Audience

To better understand the makeup of RT's audience, we examine the demographics of the users in our sample. As Twitter's API does not make such information available, we use two inference tools to uncover users' demographics. The first tool is the "botometer," which uses over 1,000 signals to calculate the probability that accounts are bots (Davis 2016). It has been widely used in many empirical studies of misinformation, although it has some well-established limitations, such as the fact that many bots are now increasingly sophisticated and mask obviously automated behavior. Previously reliable signals such as having a "stock photo" profile picture or tweeting at the same time each hour have now been addressed and many bots are more human-like. A further problem is that many genuine Twitter users deploy automated or semi-automated software to help manage their tweeting and followership patterns, which can make them appear bot-like. Notwithstanding these limitations, botometer is a scalable and accurate way of identifying bots and represents state-of-the-art for this task. Given the risk of false positives, we identify accounts as bots if they have a botometer score equal to or greater than 0.8.

The second tool is the "M3" inference tool, which provides probabilistic demographic information about users (Wang et al. 2019). It estimates their age and gender, and whether or not they are an organization. This is a notoriously difficult classification task (Hinds and Joinson 2018). The M3 tool is trained on 37 million profiles in total, of which 24 million were used to identify organizations, 15 million were used for gender inference, and 3 million for age. Performance varies across demographics, with an F1 score of 0.90 for Organizations, 0.92 for Gender, and only 0.52 for Age. This substantially outperforms other similar demographic inference tools, although we advise caution when interpreting the results for Age, which are notably weaker.

With any automated computational tool, there is a degree of error; the scale and speed afforded by such methods are counterbalanced by imperfect performance. However, they have been shown to work well when applied in aggregate to large datasets—even if some of the individual predictions they make are wrong. Although these errors are a reasonable constraint for social research, they nonetheless raise the risk of unfair outcomes and even social injustice when tools are used without consideration of their biases. Biased tools may perform unequally across different social groups; for instance, many content classifiers have different error rates for people from different ethnicities, which undermine their scientific validity and trustworthiness, and can have negative social effects (Davidson et al. 2019). These problems are often reflected in tools' design; the M3 inference tool only makes binary gender classifications, potentially misgendering individuals who do not identify as male or female, and it is likely to misidentify non-cis-gendered Twitter users. However, to the best of our knowledge, M3 and the botometer have not been found to contain large biases and are appropriate for empirical use, so long as they are used to assess online users (as is the case here) rather than to *target* or make decisions about them.

Table 1. Demographics and Activity of Users in Our Sample and a Control Group.

Variable	RT Sample	Control Group	Difference	Significance (Permutation t-test) with Holm's Correction	Significance (Wilcoxon) with Holm's Correction
Bot	0.6	0.21	0.39	$p < .001$	$p < .001$
Age (18 and below)	0.23	0.30	-0.07	$p < .05$	$p < .001$
Age (19–29)	0.3	0.33	-0.03	$p < .05$	$p < .001$
Age (30–39)	0.2	0.15	0.05	$p < .05$	$p < .001$
Age (40 and above)	0.27	0.22	0.05	$p < .05$	$p < .001$
Gender (male)	0.75	0.62	0.13	$p < .05$	$p < .001$
Gender (female)	0.25	0.38	-0.13	$p < .05$	$p < .001$

Note. RT = Russia Today.

Our results show several important differences between followers of RT and Twitter users more generally. First, the botometer classifies 39 percent of RT’s followers as bots (the average score is 0.6, as shown in the table, from which 39% of users have a score that is equal to or greater than 0.8). This is far larger than the 1.5 percent it classifies for our random sample (from an average score of 0.21), as well as Varol et al.’s (2017) estimate that between 9 and 15 percent of active Twitter accounts are bots. Second, followers of RT have different demographics to typical Twitter users (Table 1). They are far more likely to be male (.75 average probability compared with .62) and are likely to be slightly older (their probability of being 30–40 is .2 compared with .15 for the random sample, and their probability of being over 40 is .27 compared with .22). These values are all statistically significant measured using two nonparametric tests: (1) a permutation *t*-test and (2) a Wilcoxon test. To avoid Type I errors (a “false positive”), we apply a Holm’s correction for the multiple pairwise comparisons and in all cases results are still significant. The samples can be considered independent for the purposes of statistical testing (McDonald 2014). Results are shown in Table 2. We advise caution when interpreting the differences in age given the lower performance of the M3 inference tool on this demographic, as discussed above.

The Interests of RT’s Twitter Audience

To better understand the interests of RT’s Twitter followers, we examine the ten most followed accounts in our sample, as shown in Table 3. Users identified as bots were removed.⁵ The ten most followed include six news sources (NY Times, BBC Breaking News, BBC World, CNN Breaking, CNN, and Reuters), two social media platforms (YouTube and Twitter), and two American presidents (Barack Obama and Donald Trump). These are all mainstream Twitter accounts, which have a large number of followers. The fact that these are the most followed accounts suggests, initially, that RT’s followers are quite typical of broader Twitter populations. However, the ten most followed accounts offer only a very broad view of the interests of RT’s Twitter audience.

Table 2. Ten Most Followed Accounts for the Cohort of Users.

Rank	First	Second	Third	Fourth	Fifth	Six	Seventh	Eighth	Ninth	Tenth
Account name	Barack Obama	Nytimes	BBC Breaking	BBC World	YouTube	cnbrk	CNN	Reuters	realDonaldTrump	Twitter
Number of followers (with bots excluded)	2,592 (41%)	2,304 (36%)	2,214 (35%)	2,212 (35%)	2,203 (35%)	2,189 (35%)	2,091 (33%)	2,031 (32%)	1,918 (30%)	1,839 (29%)

Table 3. The Thirty-Five Main Segments of RT's Twitter Audience.

Segment	Number of Users	Label	Segment	Number of Users	Label
1	2,042	Indian (film) celebrities	19	88	Spam (sales accounts)
2	1,745	Anglophone international news	20	84	U.S. rappers/music
3	984	Argentinian/Hispanic celebrities	21	76	Korean pop music
4	760	Arabic religious/political actors	22	68	Computer game players
5	551	Brazilian/Lusophone media and left-wing political actors	23	67	Malay celebrities
6	548	U.K. and U.S. left-wing activists	24	61	Spanish satirical news/pro-Catalan politics
7	396	Politicized U.K. celebrities	25	59	Serbian/Croatian politics/news
8	367	Russian news/politics	26	57	Dutch news/left politics
9	366	U.S. right-wing and alt-right politics	27	56	Italian news/left politics and celebrities
10	266	Nigerian news/politics	28	55	Australian news/right and left politics
11	217	Indonesian political and business news	29	52	Greek news/right and left politics
12	179	French political and comic actors and news	30	50	Nepalese celebrities/politics
13	161	Turkish media and comic actors and news	31	49	Dance music
14	149	South African news, media, political and religious actors	32	48	Thai media, entertainment, and political celebrities
15	113	Travel news	33	44	Filipino celebrities
16	107	Pornography	34	42	German news/Green politics
17	98	African social media and celebrity actors	35	37	Japanese news/politics
18	96	Cryptocurrency and blockchain news/actors			

Note. RT = Russia Today.

We applied the Louvain algorithm to identify different segments of users within our sample. The Louvain identified sixty-nine segments, of which the largest thirty-five account for 97.5 percent of all 10,391 users. We focus on these thirty-five segments for the remainder of the analysis. As the Louvain algorithm is a stochastic process, it can return different results on each run. This raises the question of whether the sets of users discovered by the algorithm are consistent (i.e., they remain the same each time). To check the stability of the segments, we ran the algorithm on the data ten times, each time with a different random seed, and compared the sets of users with the Adjusted Rand Index⁶ (ARI; Steinley 2004). The mean average ARI between collections on different applications of the algorithm came out as 81.4 percent. Following Steinley (2004), this value (greater than 80 percent) suggests good agreement between the outputs of the algorithm, which suggests it has reliably partitioned the users into distinct segments. We included bots in this analysis as they play an important role in the structure of the network. The proportion of bots in each segment varies substantially from 65 percent in one community (accounting for twenty-five out of thirty-nine users) down to 17 percent in another (accounting for ten out of sixty users). This potentially reflects findings from other research which suggests that some bots operate as coordinated networks, working in concert to infiltrate and manipulate online discussions (Shao et al. 2018).

To characterize the thirty-five segments, we examined the accounts that they follow. First, we examined the ten accounts most followed by users in each segment. Second, we implemented a surfacing methodology to identify the ten most *characteristic* accounts followed by each segment, as described above in our “Method” section. We assigned a label to all of the most characteristic and most followed accounts; across the thirty-five segments, there were 567 accounts combined. Every account was labeled by two members of the research team using an open coding “ground-up” framework, which we iteratively revisited as the coding schema was developed (Strauss and Corbin 1990). From each segment, we took the labels for the ten most characteristic and ten most followed accounts and agreed a single label. These are shown in Table 3.

Analyzing the “most characteristic” accounts amplifies the interpretative signals, which are otherwise only latent in the list of “most followed” accounts. In nearly all cases, the most characteristic accounts hone in on interests that were hidden in the lists of most followed accounts, thereby allowing us to differentiate segments in a more nuanced way. This is important for cutting through the noise; the most followed accounts are often (1) the same or similar across all of the segments, (2) similar to the most followed accounts in the cohort overall (both with bots removed and included), and (3) similar to the most followed accounts on Twitter in general, such as Barack Obama, YouTube, The BBC, and Donald Trump. For instance, in Segment 6, the most followed accounts include some left-wing/activist actors (such as WikiLeaks, Edward Snowden, and Bernie Sanders, and also the BBC, New York Times, and Barack Obama), and the list of most characteristic accounts in Segment 6 is primarily comprised of left-wing organizations. Similarly, for Segment 10, there are two Nigerian accounts in the most followed list (@MobilePunch and @DONJAZZY); in the most

characteristic list, all of the accounts have a connection with Nigeria. Using both signals lets us identify the most salient distinctions between the segments.

Most of the thirty-five segments are distinguished by language and/or nationality, with many different nationalities represented. Within these linguistically distinct segments, the majority of them pivot around “soft news,” featuring a mix of celebrity actors (such as film stars, social media influencers, sports personalities, comedians, and entertainers) and political actors (including both representatives and media pundits, covering right, left, green, and feminist positions). The predominance of soft news segments reflects RT’s efficacy in using entertainment stories to gain social media audiences (Orttung and Nelson 2019). Some segments also reflect “hard news,” primarily defined in relation to international news, business, and political representatives.

Discussion

This is the first study that takes a wide view on RT’s Twitter followers, providing insight into who they are, how active they are, and what their interests are. We challenge the idea that RT’s audience mirrors the characteristics of the broadcaster itself, showing that it is highly heterogeneous. Our analysis of RT’s Twitter audience puts into question several long-standing views, and assumptions, about the types of people who engage with RT—and, by extension, the nature of RT and its role in the international news ecosystem. Noticeably, we find that most of RT’s Twitter audience engages with it only very rarely. This suggests that great caution should be taken with overly “amplified” interpretations of RT’s audience, which examine only a small hardcore of active supporters. This is because (1) they are unlikely to be representative of all of RT’s online audience and (2) the broader audience displays internal diversity, with a mix of perspectives identified. Our analysis suggests that simplistic assessments of the people who are exposed to and consume RT’s content (and of the putative effects of RT on them) are likely, by their very nature, to be wrong. Exposure does not mean either engagement or endorsement.

Significantly, we found that RT’s Twitter audience is far more likely to be bots, with the estimated number of bots in our sample two to three times higher than among active Twitter accounts in general. This suggests that many of the accounts are either automated or semi-automated, or exhibit unusual tweeting patterns (Varol et al. 2017). Noticeably, the users in our sample of 10,391 RT Twitter followers were far more likely to be male and slightly older than the random sample of Twitter users. Although the difference is reasonably small for both characteristics, it is statistically significant and supports previous characterizations of RT’s audiences in the literature (Orttung and Nelson 2019).

From the most followed accounts across the whole cohort, we find that they follow a large number of news sources (accounting for six of the ten most followed). This suggests that users engage with RT as one way of consuming news and information in the context of many other sources of such content, rather than as a single way of accessing content. This resonates with a key finding of qualitative analyses, which

indicate that many consumers use RT as one news reference point among others to get an alternative perspective on events (Crilley et al. 2019). RT's Twitter account may be used by audiences as an alternative to "Western" mainstream media in some cases and to receive a Russian perspective in others. Following RT on Twitter is not necessarily indicative of an extremist, fringe or anti-Western perspective.

Using network science techniques, we identified sixty-nine audience segments, of which we focused on and labeled 35. We find considerable variation in terms of the interests of the different segments, and demonstrate that they vary in terms of their country/cultural background, political views, and general interests. Some segments are defined primarily by nationality (e.g., Argentine pop culture, Indonesian pop culture, and Malay pop culture), others by different range of interests within the English language (e.g., Computer game playing, cryptocurrency, and pornography), and some where interests and nationality overlap, for example, Spanish news/politics, Italian news/politics (left-leaning), and Nigerian politics. Surprisingly, and further strengthening our argument that the bulk of RT's Twitter audiences are not for the most part inherently extremist, we find little evidence of large segments of radical left or ultra right-wing groups. However, it should be noted that some, albeit small, amounts of conspiratorial content in the tweets of QAnon and Make America Great Again (MAGA) accounts in the U.S. right-wing segment are apparent. Many of the smaller segments are also overlaid with either left-wing or right-wing political orientations—not so much extremist, as groups marginal to broad centrist politics. However, these small disparate right-wing and left-wing groupings following RT also map onto other segmented interests, and the theme of "News" cuts across several of the segments, often serving as the defining interest of the national segments. This also supports our claim that RT's audience engages with its content alongside other news providers as a point of reference rather than as a singular main source of news.

Our research design focuses on the large number of users who are exposed to RT's content, rather than focusing solely on the active users who engage with it. This has enabled us to provide new insights and to show the need to reconfigure understandings of RT, putting into question many widely held assumptions. However, it also raises its own limitations, namely, the weak tie that "followership" constitutes and what the true depth of connection that followers of RT have with the broadcaster. This is an important limitation, and we caution that our results are intentionally situated at a very broad optic and that future work is needed to contrast them systematically with other assessments of RT. We also focus primarily on Twitter, and the users there are unlikely to be representative of users online "in general" or of other platforms.

Conclusion

RT's global Twitter audiences are best characterized as being diverse and fragmented, rather than as politically extreme. Broad brush stroke characterizations based on unfounded assertions are not useful to developing a more accurate understanding of RT and its impact. For example, the idea that RT's audience "believe that human rights are a sham and democracy a fix" (Cohen 2014) and that they are a

“disillusioned western audience” (Turner 2016) are only applicable to a very small but vocal minority. It is unhelpful to make generalizations based on vocal minorities; they impede progress toward a better understanding of how RT is positioning itself as a new player in the shifting landscape of international news, disrupting the hegemony and prestige of legacy broadcasters, and providing alternative worldviews that have appeal with audiences. This data-driven study has offered new insights on who follows RT, how much they engage with it, and what their interests are. Future work should address the question of who consumes and engages with RT’s content further, investigating across multiple platforms and directly contrasting results with the audiences of mainstream media broadcasters.

Acknowledgments

The authors thank Precious Chatterje-Doody, Stephen Hutchings, and Vera Tolz for the feedback on earlier iterations of this paper. Thanks also to the editors of this journal and the anonymous peer reviewers who provided exceptionally constructive and helpful comments on this work. Rhys Crilley would also like to thank Mariëlle Wijemars, Teemu Oivo, and other members of the Aleksanteri Institute at the University of Helsinki who gave feedback on a presentation of this research in its early stages.


Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was part of the Arts and Humanities Research Council funded project “Reframing Russia: From Cold War to ‘Information War’?”: [Grant Number AH/ P00508X/1].

ORCID iD

Rhys Crilley  <https://orcid.org/0000-0003-1816-5535>

Notes

1. The full terms and conditions are available at <https://twitter.com/en/tos>.
2. See <https://python-louvain.readthedocs.io/>.
3. See <https://www.brandwatch.com/blog/amazing-social-media-statistics-and-facts/>.
4. The 10 percent threshold was selected by evaluating the results for different thresholds. Our testing indicates that a range of between 5 and 20 percent gives useful results; a lower value means that the surfaced accounts tend to be more niche and less well known, while a higher value means that the surfaced accounts tend to be more well known but also less easily differentiated, approximating the most followed accounts when the threshold is high enough.
5. Analysis was rerun with bots included, and very similar results were reported.
6. Using the Adjusted Rand Index function of python’s scikit-learn 0.23.1 library.

References

- Al-Rawi, A. 2017. "News Organizations 2.0." *Journalism Practice* 11 (6): 705–20.
- Barberá, P. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23 (1): 76–91.
- Birrell, I. 2018. "RT's Propaganda Deserves Mockery But Not a Ban." *The Times*, April 20. <https://www.thetimes.co.uk/article/rt-s-propaganda-deserves-mockery-but-not-a-ban-pnh-85vr3b>.
- Blondel, V. D., J. L. Guillaume, R. Lambiotte, and E. Lefebvre. 2008. "Fast Unfolding of Communities in Large Networks." *Journal of Statistical Mechanics: Theory and Experiment*, 10: 1–12.
- Bodine-Baron, E. A., T. Helmus, A. Radin, and E. Treyger. 2018. "Countering Russian Social Media Influence." RAND Corporation. https://www.rand.org/content/dam/rand/pubs/research_reports/RR2700/RR2740/RAND_RR2740.pdf.
- Caplan, R., J. Donovan, L. Hansen, and J. Mathews. 2018. "Algorithmic Accountability: A Primer." *Data & Society*. <https://datasociety.net/library/algorithmic-accountability-a-primer/>
- Cha, M., H. Haddadi, F. Benevenuto, and K. P. Gummadi. 2010. "Measuring User Influence in Twitter: The Million Follower Fallacy." In *Proceedings of the 4th ICWSM*, 10–17. http://twitter.mpi-sws.org/icws2010_fallacy.pdf.
- Chadwick, A. 2017. *The Hybrid Media System: Politics and Power*. Oxford: Oxford University Press.
- Chatterje-Doody, P. N., and R. Crilley. 2019a. "Making Sense of Emotions and Affective Investments in War: RT and the Syrian Conflict on YouTube." *Media and Communication* 7 (3): 167–78.
- Chatterje-Doody, P. N., and R. Crilley. 2019b. "Populism and Contemporary Global Media: Populist Communication Logics and the Co-construction of Transnational Identities." *Populism and World Politics*, ed. F. Stengel, D. B. MacDonald, and D. Nabers, 73–99. London: Palgrave Macmillan.
- Chatterje-Doody, P. N., and V. Tolz. 2020. "Regime Legitimation, Not Nation-Building: Media Commemoration of the 1917 Revolutions in Russia's Neo-authoritarian State." *European Journal of Cultural Studies* 23 (3): 335–53.
- Cohen, N. 2014. "Russia Today: Why Western Cynics Lap Up Putin's TV Poison." *The Guardian*, November 8. <https://www.theguardian.com/commentisfree/2014/nov/08/russia-today-western-cynics-lap-up-putins-tv-poison>.
- Crilley, R., and P. N. Chatterje-Doody. 2020a. "Emotions and War on YouTube: Affective Investments in RT's Visual Narratives of the Conflict in Syria." *Cambridge Review of International Affairs* 33:713–33. doi:10.1080/09557571.2020.1719038.
- Crilley, R., and P. N. Chatterje-Doody. 2020b. "From Russia with Lols: Humour, RT, and the Legitimation of Russian Foreign Policy." *Global Society*. Published electronically November 18. doi:10.1080/13600826.2020.1839387.
- Crilley, R., M. Gillespie, and A. Willis. 2019. "Tweeting the Russian Revolution: RT's#1917LIVE and Social Media Re-enactments as Public Diplomacy." *European Journal of Cultural Studies* 22 (3): 354–73.
- Cull, N. J., V. Gatov, P. Pomerantsev, A. Applebaum, and A. Shawcross. 2017. "Soviet Subversion, Disinformation and Propaganda: How the West Fought against It. An Analytic History, with Lessons for the Present." *LSE Consulting*. <https://bit.ly/2L7HClk>.
- Dajani, D., M. Gillespie, and R. Crilley. 2019. "Differentiated Visibilities: RT Arabic's Narration of Russia's Role in the Syrian War." *Media, War & Conflict*. Published electronically December 12. doi:10.1177%2F1750635219889075.

- Davidson, T., D. Bhattacharya, and I. Weber. 2019. "Racial Bias in Hate Speech and Abusive Language Detection Datasets." Proceedings of the Third Workshop on Abusive Language Online, ACL, 1–11.
- Davis, C. 2016. "BotOrNot: A System to Evaluate Social Bots." Proceedings of the World Wide Web Conference, Montreal, Canada, Association for Computing Machinery, 273–274.
- Dowling, T. 2017. "24-Hour Putin People: My Week Watching Kremlin 'Propaganda Channel' RT." *The Guardian*, November 29. <https://www.theguardian.com/media/2017/nov/29/24-hour-putin-people-my-week-watching-kremlin-propaganda-channel-rt-russia-today>.
- Elliot, R. 2019. "How Russia Spreads Disinformation via RT Is More Nuanced than We Realise." *The Guardian*, July 26. <https://www.theguardian.com/commentisfree/2019/jul/26/russia-disinformation-rt-nuanced-online-ofcom-fine>.
- Elsawah, M., and P. N. Howard. 2020. "Anything that Causes Chaos": The Organizational Behavior of Russia Today (RT)." *Journal of Communication* 70 (5): 623–45.
- Fisher, A. 2020. "Demonizing the Enemy: The Influence of Russian State-Sponsored Media on American Audiences." *Post-Soviet Affairs* 36 (4): 281–96.
- Fog, A. 2007. "Calculation Methods for Wallenius' Noncentral Hypergeometric Distribution." *Communication in Statistics—Simulation and Computation* 37 (2): 258–73.
- Galeotti, M. 2018. "I'm Sorry for Creating the 'Gerasimov Doctrine.'" *Foreign Policy*, March 5. <https://foreignpolicy.com/2018/03/05/im-sorry-for-creating-the-gerasimov-doctrine/>.
- Gillespie, M., and A. Webb. 2012. *Diasporas and Diplomacy: Cosmopolitan Contact Zones at the BBC World Service*. London: Routledge.
- Harsin, J. 2015. "Regimes of Posttruth, Postpolitics, and Attention Economies." *Communication, Culture & Critique* 8 (2): 327–33.
- Hinds, J., and A. Joinson. 2018. "What Demographic Attributes Do Our Digital Footprints Reveal? A Systematic Review." *PLoS ONE* 13 (1): 1–40.
- Hutchings, S., M. Gillespie, I. Yablokov, I. Lvov, and A. Voss. 2015. "Staging the Sochi Winter Olympics 2014 on Russia Today and BBC World News: From soft Power to Geopolitical Crisis." *Participations: Journal of Audience Reception Studies* 12 (1): 630–58.
- Ipsos Connect (2018) RT's Audiences in Middle East and North Africa. Internal 17-Country survey report for France 24 shared with research team.
- Jamieson, K. H. 2018. *Cyberwar: How Russian Hackers and Trolls Helped Elect a President: What We Don't, Can't, and Do Know*. Oxford: Oxford University Press.
- Kragh, M., and S. Åsberg. 2017. "Russia's Strategy for Influence through Public Diplomacy and Active Measures: The Swedish Case." *Journal of Strategic Studies* 40 (6): 773–816.
- Liang, Hai, Fei Shen, and King-wa Fu. 2017. "Privacy Protection and Self-Disclosure across Societies: A Study of Global Twitter Users." *New Media & Society* 19 (9): 1476–97.
- Lytvynenko, J., and C. Silverman. 2019. "A Timeline of How the Notre Dame Fire Was Turned Into an Anti-Muslim Narrative." *Buzzfeed News*. <https://www.buzzfeednews.com/article/janalytyvnenko/notre-dame-hoax-timeline>
- Marwick, A. E. 2018. "Why Do People Share Fake News?." *Georgia Law Technology Review* 2:474–512. <https://georgetownlawtechreview.org/why-do-people-share-fake-news-a-socio-technical-model-of-media-effects/GLTR-07-2018/>.
- McDonald, J. 2014. *Handbook of Biological Statistics*. 3rd ed. Baltimore: Sparky House Publishing.
- Miazhevich, G. 2018. "Nation Branding in the Post-Broadcast Era: The Case of RT." *European Journal of Cultural Studies* 21 (5): 575–93.
- Mickiewicz, E. 2018. *No Illusions: The Voices of Russia's Future Leaders*. New York: Oxford University Press.

- Newman, N., R. Fletcher, A. Kalogeropoulos, and R. Nielsen. 2019. "Reuters Institute Digital News Report 2019." Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2019-06/DNR_2019_FINAL_0.pdf.
- Office of the Director of National Intelligence. 2017. "Assessing Russian Activities and Intentions in Recent US Elections." https://www.dni.gov/files/documents/ICA_2017_01.pdf
- Orttung, R. W., and E. Nelson. 2019. "Russia Today's Strategy and Effectiveness on YouTube." *Post-Soviet Affairs* 35 (2): 77–92.
- Ramsay, G., and S. Robertshaw. 2019. "Weaponising News RT, Sputnik and Targeted Disinformation." King's College London: The Policy Institute, Center for the Study of Media, Communication, and Power. <https://www.kcl.ac.uk/policy-institute/assets/weaponising-news.pdf>.
- Rawnsley, G. D. 2015. "To Know Us Is to Love Us: Public Diplomacy and International Broadcasting in Contemporary Russia and China." *Politics* 35 (3–4): 273–86.
- Richter, M. L. 2017. "What We Know about RT (Russia Today)." European Values. <https://www.europeanvalues.net/wp-content/uploads/2017/09/What-We-Know-about-RT-Russia-Today-3.pdf>.
- Rogers, R. 2013. *Digital Methods*. Cambridge: MIT press.
- Rosenberg, Y. 2015. "Russia Today Airs Bizarre Anti-Semitic Conspiracy Theory about Hillary Clinton." *Tablet Mag*. <https://www.tabletmag.com/scroll/194211/russia-today-air-bizarre-anti-semitic-conspiracy-theory-about-hillary-clinton>
- Shao, C., G. L. Ciampaglia, O. Varol, K.-C. Yang, A. Flammini, and F. Menczer. 2018. "The Spread of Low-Credibility Content by Social Bots." *Nature Communications* 9 (4787): 1–14.
- Steinley, D. 2004. "Properties of the Hubert-Arable Adjusted Rand Index." *Psychological methods* 9 (3): 386–96.
- Strauss, A., and Corbin, J. (1990). *Basics of qualitative research*. London: Sage.
- Sun, N., P. Rau, and L. Ma. 2015. "Understanding Lurkers in Online Communities: A Literature Review." *Computers in Human Behaviour* 38 (1): 110–17.
- Szostek, J. 2018. "Nothing Is True? The Credibility of News and Conflicting Narratives during 'Information War' in Ukraine." *The International Journal of Press/Politics* 23 (1): 116–35.
- Tolz, V., S. Hutchings, P. N. Chatterje-Doody, and R. Crilley. 2020. "Mediatization and Journalistic Agency: Russian Television Coverage of the Skripal Poisonings." *Journalism*. Published electronically July 16. doi: 10.1177/1464884920941967.
- Turner, M. 2016. "To All the Self-Identifying Liberals Cheering about Russia Today's Bank Accounts Being Frozen, Did You Ever Consider Your Own Bias?" *The Independent*, October 18. <https://www.independent.co.uk/voices/russia-today-rt-bank-accounts-frozen-uk-england-putin-liberal-twitter-cheering-occupy-wall-street-a7367261.html>.
- Varol, O., E. Ferrara, C. A. Davis, F. Menczer, and A. Flammini. 2017. "Online Human-Bot Interactions: Detection, Estimation, and Characterization." Proceedings of the 11th ICWSM, AAAI, 280–89.
- Wang, Z., S. A. Hale, D. Adelani, P. A. Grabowicz, T. Hartmann, F. Flöck, and D. Jurgens. 2019. "Demographic Inference and Representative Population Estimates from Multilingual Social Media Data." *Proceedings of the World Wide Web Conference*, Association for Computing Machinery, 2056–67.
- Wright, K., M. Scott, and M. Bunce. 2020. "Soft Power, Hard News: How Journalists at State-Funded Transnational Media Legitimize Their Work." *The International Journal of Press/Politics* 25:607–631. doi:10.1177/1940161220922832.
- Yablokov, I. 2015. "Conspiracy Theories as a Russian Public Diplomacy Tool: The Case of Russia Today (RT)." *Politics* 35 (3–4): 301–15.

Author Biographies

Rhys Crilly is a Leverhulme Trust funded Early Career Research Associate in the Department of Politics and International Relations at the University of Glasgow. Prior to this, he worked on the Arts and Humanities Research Council-funded “Reframing Russia” project at The Open University, UK. His research explores the intersections of popular culture, social media, and global politics with a specific interest in war and legitimacy. He has published articles in outlets such as *Journalism*, *International Studies Review*, *International Affairs*, and *Millennium*. He is currently working on writing his first monograph. He tweets at @rhyscrilly.

Marie Gillespie is a Professor of Sociology at The Open University, UK. Her teaching and research interests revolve around media and migration, diaspora cultures, politics, and social change. Her books include *Social Media, Religion, and Spirituality: Diasporas and Diplomacy: Cosmopolitan Contact Zones at the BBC World Service*; and *Drama for Development: Cultural Translation and Social Change*. Her recent research includes projects on forced migration and digital connectivity among Syrian and Iraqi refugees, and on culture and diplomacy in Egypt and Ukraine.

Bertie Vidgen is a Research Associate within the public policy program at the Alan Turing Institute. His main research is focused on detecting, analyzing, and countering online hate speech, examining it in the context of both news and social media. In his work, he primarily uses computational social science methods, including machine learning, natural language processing, and statistical modeling. Before he joined the Turing, he studied for a DPhil at the University of Oxford’s Oxford Internet Institute, where he researched Islamophobic hate speech among followers of U.K. political parties on Twitter.

Alistair Willis is a Senior Lecturer in Computing and leads the Artificial Intelligence group in the School of Computing and Communications at The Open University, UK. His research interests focus mainly on natural language processing, especially on computational models of linguistic ambiguity and large-scale text and data analytics.