



Computational
Propaganda
Research Project

The IRA, Social Media and Political Polarization in the United States, 2012-2018 (Appendices)

Philip N. Howard, *University of Oxford*
Bharath Ganesh, *University of Oxford*
Dimitra Liotsiou, *University of Oxford*

John Kelly, *Graphika*
Camille François, *Graphika*



UNIVERSITY OF
OXFORD

Contents

Appendix A: Monthly Timeline of IRA Activity.....	3
Appendix B: Top Posts from the IRA	4
Appendix C: Facebook Ads and Audience Segments	9
Appendix D: Networks of Trolls: Connectivity between the IRA Twitter Accounts.....	11
Appendix E: Russian Manipulation in Context.....	15
Appendix F: A Cross-Platform Case Study	17
Appendix G: Accounts Operated from the Same IP Address, by Category.....	22
Appendix H: Heterophily and Homophily of IRA Accounts on Twitter.....	25
Series Acknowledgements	27
Author Biographies	28

Tables

Table 1: The Five Most Shared Facebook Posts, with Most Shared First, 2015-17.....	4
Table 2: The Five Most Liked Facebook Posts, with Most Liked First, 2015-17.....	5
Table 3: The Five Most Liked Instagram Posts, with Most Liked First, 2015-17.....	7
Table 4: Classification of Twitter Accounts	12
Table 5: Further Classification of Twitter Accounts	13
Table 6: Top Ten Email Providers in Registration Email Addresses	14
Table 7: Vignette of BlackMattersUS Activity Timeline	19
Table 8: Category Overlap by Shared IP Addresses (Create IP).....	23
Table 9: Category Overlap by Shared IP Addresses (Last IP).....	24
Table 10: Heterophily and Homophily by Source and Target	26

Figures

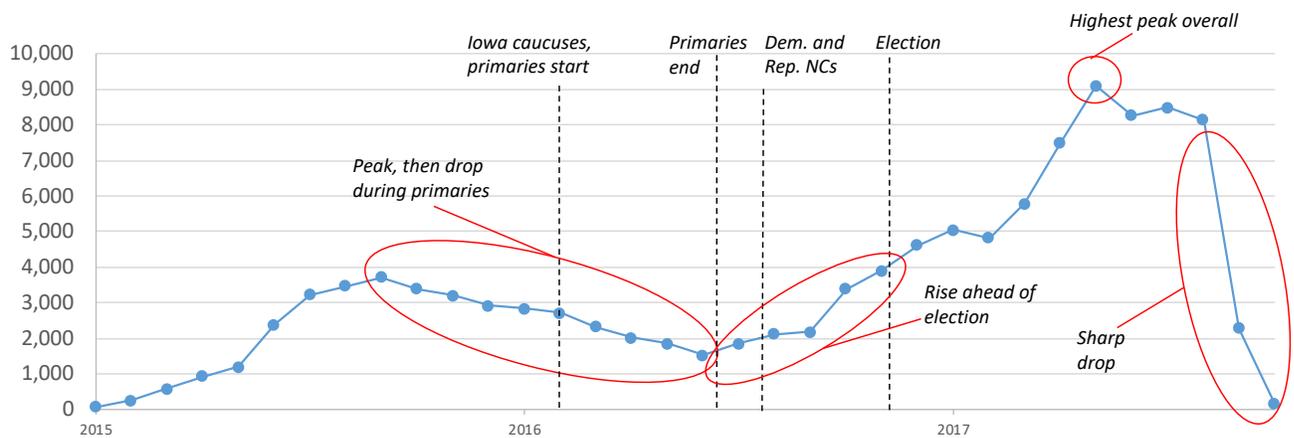
Figure 1: The Volume of Instagram Posts, Monthly Totals.....	3
Figure 2: Graph of Facebook Ads that Linked Different Audience Segments	9
Figure 3: A Facebook Event Created by the Black Matters Account.....	18
Figure 4: An Image from Black Matters US.....	20
Figure 5: Black Matters Cross Platform Activity.....	21

Appendix A: Monthly Timeline of IRA Activity

This section presents the monthly timelines of Instagram organic posts, in Figure 1.

Figure 1 shows that monthly Instagram post volume started rising in 2015, but dropped before and during the primaries. From the end of the primaries it started rising again, in the months leading up to the election. After the election it continued rising, reaching its highest overall peak in 2017, demonstrating that the IRA's activities did not stop, but rather intensified, after the 2016 elections. In the summer of 2017, there is a sharp drop in volume, likely because Instagram detected and deactivated these IRA accounts.

Figure 1: The Volume of Instagram Posts, Monthly Totals



Source: Authors' calculations based on data provided by the SSCI.

Appendix B: Top Posts from the IRA

This Appendix presents the full content of the most shared and most liked posts on Facebook, and the most liked posts on Instagram (which does not offer a built-in share capability). As stated earlier, these most liked and most shared posts are important because we have evidence that hundreds of thousands of people not only saw them, but they actively engaged with them on Facebook, and even shared them with their friends, spreading them even further.

The Five Most Shared Facebook Posts

In Table 1, for each Facebook post, we provide the image, the number of times shared, the name of the page, the text of any accompanying message, the date of post, and a summary of the themes.

Table 1: The Five Most Shared Facebook Posts, with Most Shared First, 2015-17

1		956,007 shares, posted by South United, on Sept. 3, 2017. Themes: pro-gun, the South
2		640,390 shares, posted by Being Patriotic, with message: “At least 50,000 homeless veterans are starving dying in the streets, but liberals want to invite 620,000 refugees and settle them among us. We have to take care of our own citizens, and it must be the primary goal for our politicians!”, on Sept. 8, 2016. Themes: anti-immigrant, nationalist, frames, immigration as veterans versus immigrants
3		539,012 shares, posted by Blacktivist, with message “PLEASE HELP THIS VIDEO GO VIRAL. THESE COPS CAN BEAT UP INNOCENT MAN, GET AWAY WITH MURDER EVEN WHEN IT’S CAUGHT ON CAMERA JUST BECAUSE THEY WEAR A 'MAGICAL' BADGE AND GUN. But it doesn` t make what they do right.”, on Nov. 23, 2016. Themes: police violence against African Americans. <i>Video: three white men (one seems armed and seems to have a police badge) holding down, beating up and hand scuffing a black man</i>
4		363,679 shares, posted by BM, with message “Please help this go viral. When the color of your skin is seen as a weapon, you will never be seen as unarmed even when you are handcuffed in a police custody.”, on Sept., 30, 2016. Themes: police violence against African Americans. <i>Video: seems to show one and then two white men (police?) repeatedly hitting a black man</i>

5



312,667 shares, posted by Stop A.I., with message “On June 27, 1952, American government passed a law, called “1952 McCarran Walters act”, that actually outlawed Sharia, but Obama never intended to enforce it or even let you know about it at all. Instead he started to import thousands of aggressive Muslim “refugees”, who refused to integrate and demanded to be allowed to live under sharia law instead of American constitution., Now it’s within Donald Trump’s authority to enforce that law and ban Sharia in every state across America., Do you want him to do that?!”, on Jan. 23, 2017. Themes: anti-Muslim, anti-refugee, anti-Obama, pro-Trump,

Source: Authors’ calculations based on data provided by the SSCI.

This table shows the content of the five most shared Facebook posts, including the message and the media (picture or video) that went with each, the number of total shares, the name of the Facebook page that posted it, the date it was posted, and the key themes and issues the posts discusses and targets. These posts have been shared at least 300,000 times, with the top three having been shared more than have a million times.

Posts 2 and 4 are the only ones posted before the 2016 presidential election. The former post was shared more than 600,000 times, and contains an anti-immigrant message, framing immigration as veterans versus immigrants. The latter post, shared more than 360,000 times and posted just a few days before the election, decries police violence against African Americans.

The Five Most Liked Facebook Posts

Table 2 shows that the top five most liked posts on Facebook contain three non-political posts, which are all around African American culture. The only two political posts were also in the top five most shared Facebook posts, presented in the previous sub-section in Table 1. The only post that was posted before the election is number 5 which shows anti-immigrant sentiment, and was discussed in the previous sub-section.

Table 2: The Five Most Liked Facebook Posts, with Most Liked First, 2015-17



250,644 likes, posted by Blacktivist, with message “This video will make you cry. Husband surprises her wife. This is how black man appreciates and loves his wife”, on Dec. 12, 2016. Themes: meme, no political content

2



162,862 likes, posted by Stop A.I., with message “On June 27, 1952, American government passed a law, called “1952 McCarran Walters act”, that actually outlawed Sharia, but Obama never intended to enforce it or even let you know about it at all. Instead he started to import thousands of aggressive Muslim “refugees”, who refused to integrate and demanded to be allowed to live under sharia law instead of American constitution. Now it’s within Donald Trump’s authority to enforce that law and ban Sharia in every state across America. Do you want him to do that?!”, on Jan. 23, 2017. Themes: anti-Muslim, anti-refugee, anti-Obama, pro-Trump

3



107,998 likes, posted by Blacktivist, with message “Listen to these young Kings.”, on Mar. 09, 2017. Themes: meme, no political content. *Video: Two young African American boys teaching another how to box, helping him train*

4



92,993 likes, posted by Blacktivist, with message “Never trade you brotherhood for this world”, on Dec. 22, 2016. Themes: seems to not be (directly) political. *Video: a row of African American men, dressed formally, in what starts like military routines and seems to turn into a synchronized dance, context not clear.*

5



73,167 likes, posted by Being Patriotic, with message “At least 50,000 homeless veterans are starving dying in the streets, but liberals want to invite 620,000 refugees and settle them among us. We have to take care of our own citizens, and it must be the primary goal for our politicians!”, on Sept. 8, 2016. Themes: Anti-immigrant, nationalist, frames immigration as veterans versus immigrants

Source: Authors’ calculations based on data provided by the SSCI.

The Five Most Liked Instagram Posts

Table 3 shows that the top five most liked Instagram posts were each liked by hundreds of thousands of Instagram users. None of them is directly political or related to the 2016 election, rather they tend to relate primarily to African American culture and rights, as well as to feminism and LGBT rights. All of them are from after the election, from summer 2017, with the top four from June 2017.

Table 3: The Five Most Liked Instagram Posts, with Most Liked First, 2015-17

1	<p>All the tones are nude! Get over it!</p> 	<p>254,179 likes, posted by blackstagram__, with message “What is your color? @expressiontees @kahmune #blackexcellence#blackpride#blackandprod#blackpower#africanamerican#melanin#ebony#panafrican#blackcommunity#problack#brownskin#unapologeticallyblack#blackgirl#blackgirls#blackwomen#blackwoman”, on Jun. 11, 2017. Themes: not related to the election</p>
2	<p>SIZE 5</p> 	<p>168,998 likes, posted by _blacktivistt_, with message “What a great idea for kids. #Blacktivist#hotnews #black#africanamerican#blacklivesmatter#blackpride#blackandproud#dreamchasers#blackgirls#blackwomen#blackman#westandtogether#proudtobeblack#blackbusiness#blackunity#blackis#melanin#icantbreathe#neverforget #sayhername #blackhistorymonth”, on Jun. 17, 2017. Themes: Not directly political, but some hashtags are (in bold)</p>
3	<p>When people say fashion is only for slim women show them this 🍑🍑🍑🍑</p> 	<p>129,400 likes, posted by feminism_tag, with message “@expression_tees #feminismtag #feminism #female #humanrights #equality #freedom #love #support #feminist #yesallwomen #womensrights #heforshe #femen #equality #freedom #feminist #woman #riot #support #protest #gay #lesbian #equality #motherhood #sextremism”, on Jun. 10, 2017. Themes: Not directly political. Some hashtags somewhat political (in bold). Feminism, African American and LGBT themes.</p>

4

👑 Best cake ever 👑



124,936 likes, posted by blackstagram__, with message “Repost from @sincerely_black_ Yesss, girl #lastchemo #blackpower #pride #power”, on Jun. 19, 2017. Themes: Not directly political

5



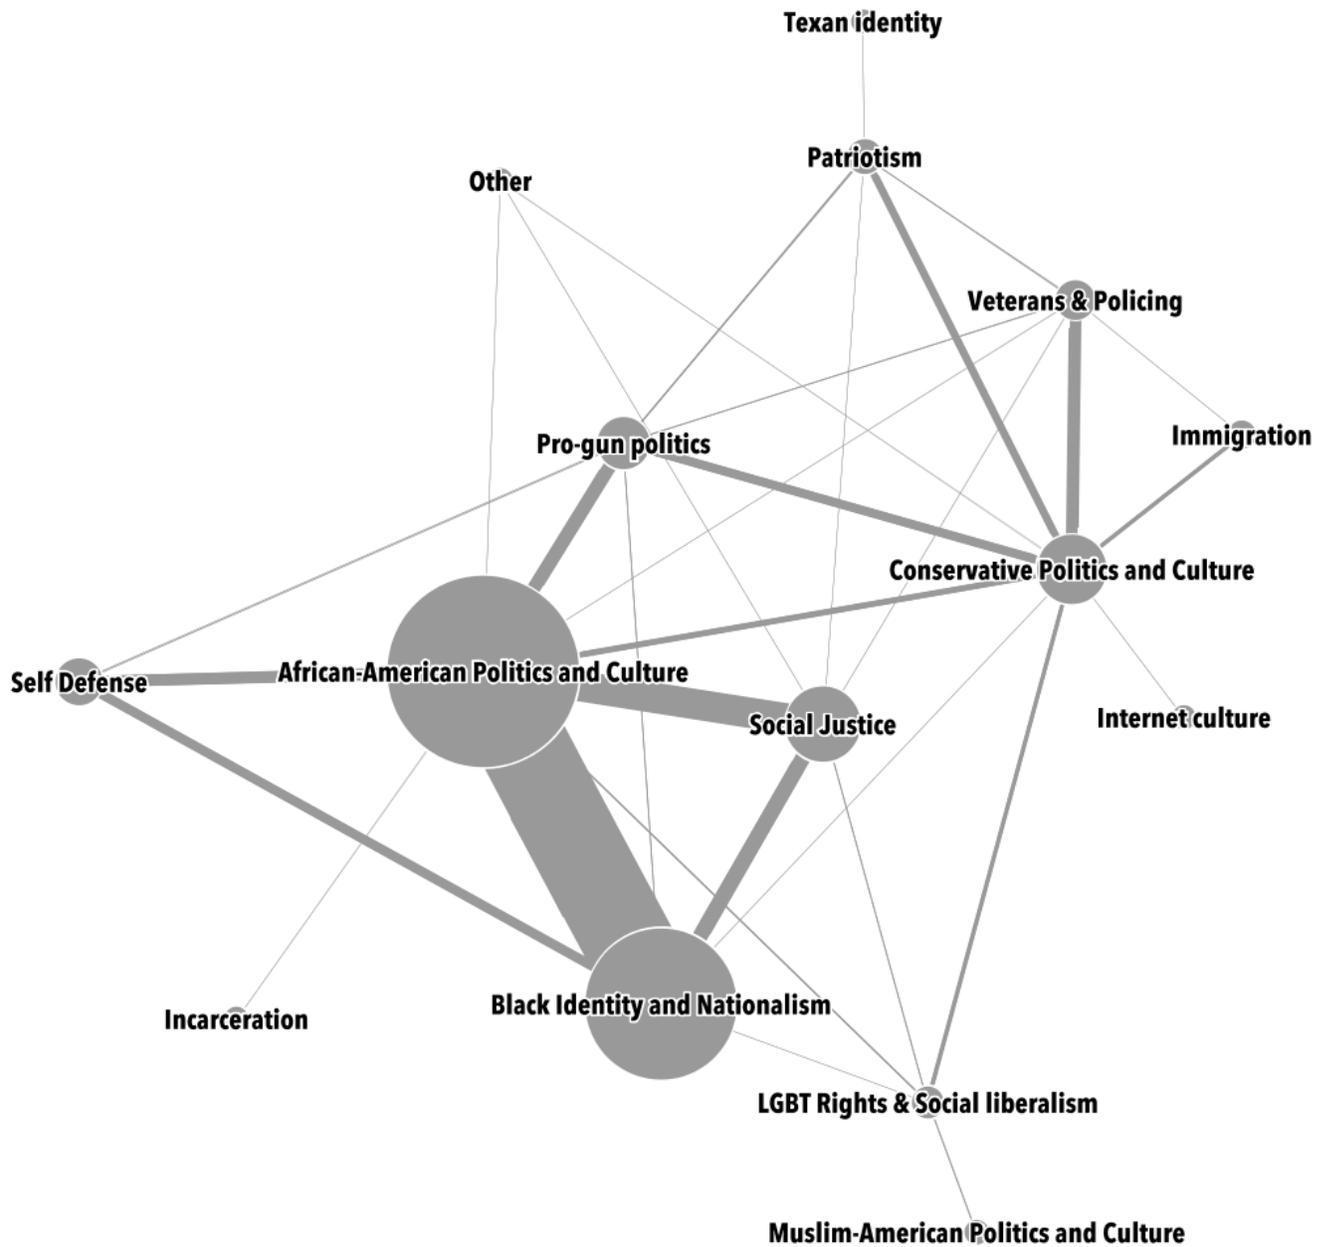
111,102 likes, posted by blackstagram__, with message “Beautiful in every shade. @afrokingdom_ #blackexcellence#blackpride#blackandproud#blackpower#blackbeauty#blackisbeautiful#blackgirlmagic#blackgirlsrock#naturallyshesdope#blackgirl#blackgirls#blackwomen#blackwoman#blackout#blackqueens#blackmodel#blackmodels#blackgirlskillingit#melaninonfleek# melaninpoppin”, on Aug. 19, 2017. Themes: Not directly political.

Source: Authors’ calculations based on data provided by the SSCI.

Appendix C: Facebook Ads and Audience Segments

Through audience segmentation, the IRA sought to bring together and push apart US voter groups in different ways, which we can see by examining the 378 ads that targeted multiple audience segments (Figure 2).

Figure 2: Graph of Facebook Ads that Linked Different Audience Segments



Note: This analysis based on 378 ads. The figure only includes ads that targeted multiple segments, so (for example) “Latin American Politics and Culture” are not included. The figure does not reflect the IRAs overall activity.

Source: Authors’ calculations based on data released by the HPSCI.

Figure 2 only shows ads that targeted multiple segments, so does not reflect the IRA’s overall activity. Nevertheless, we can see two main clusters of segments: one centering on issues facing African Americans, and the other focusing on interests that are common in conservative politics. A couple of crossovers between these groups should be noted, seeking to link “African American Politics and Culture” and “Pro-gun politics”, as evidenced in ads for a campaign to “Protect Black Gun Owners”, whose slogans included statements such as “Black Guns Matter” (Ad ID: 3274) and “Black American racial experience is real. We support the 2nd ammendment [sic] for our safety” (Ad ID: 3194), as well as advertisements for free self-defense classes (Ad ID: 3355). Ads that had interests which crossed over between “African American Politics and Culture” and “Conservative Politics and Culture” segments were part of the Williams & Kalvin campaign (now known to be a set of bloggers co-opted by the IRA to produce political content for black audiences). These interests were primarily about Christianity, such as “Jesus”, “Bible”, “Faith”, and “Worship”. As part of their campaign, Williams & Kalvin often posted videos deriding Hillary Clinton, likely producing messages that could have resonated with both conservatives and African Americans.

Appendix D: Networks of Trolls: Connectivity between the IRA Twitter Accounts

The Twitter IRA data set contains account metadata (including creation date, profile descriptions, language code, and self-reported location) and tweet histories for 3,841 accounts controlled by the IRA. The accounts in this dataset were active over a period of about eight years, from 2009 to 2017. Using a series of methods (including metadata analysis, text analysis and qualitative examination of account names, profile descriptions and sample tweets) we constructed a classification of the IRA Twitter accounts. Our classification included language groups (Russian, English, Italian, etc.); political ideology groups for evidently US accounts (Liberal vs. Conservative); and non-political groups like marketing (which includes accounts practicing techniques for online audience building, such as “hash tagging”). Several subgroups (Table XY) were evident within accounts in the Liberal group, including African American and LGBT, as well as accounts that appeared generally Liberal/Progressive. There were also subgroups on the Conservative side, but these were based on behavioral rather than thematic differences. There was also a large group of mixed accounts, which included a variety of accounts without clear or stable intended identities, such as those that messaged on different sides of the political spectrum at different times, tweeted motivational sayings, or amplified hashtag campaigns such as #ColumbianChemicals and #Fukushima2015. The IRA accounts differ across a number of thematic and behavioral features in ways that will support a range of additional categorization schemas.

Table 4: Classification of Twitter Accounts

Category	Tweets	Accounts	Tweets per Account	Start Date	End Date
Russian	5,231,688	1,384	3,810	September 9, 2009	May 31, 2018
Conservative	959,951	745	1,294	April 18, 2014	March 22, 2018
Local/News	747,147	53	14,097	December 27, 2012	December 24, 2017
Liberal	591,403	247	2,394	June 18, 2013	May 31, 2018
SMM	348,367	105	3,415	June 6, 2014	September 20, 2017
German	105,278	117	908	October 8, 2015	October 21, 2017
Mixed	415,367	1,002	416	May 9, 2009	December 9, 2017
Arabic	37,714	52	725	November 27, 2013	January 5, 2018
Food Hoax	33,040	106	312	May 8, 2014	October 4, 2016
Italian	18,413	9	2,046	March 6, 2017	October 21, 2017
French	1,409	11	128	October 9, 2015	March 8, 2017
Spanish	57	9	6	September 19, 2014	May 26, 2017
Portuguese	1	155	155	April 15, 2017	November 14, 2017
Totals	8,489,834	3,841	2,210*	n/a	n/a

Note: (*) denotes average.

Source: Authors' calculations based on data provided by the SSCI.

Table 5: Further Classification of Twitter Accounts

Category	Accounts	Tweet Count	Tweets per Account	Earliest	Latest
Arabic	52	37,714	725	November 27, 2013	January 5, 2018
French	11	1,409	128	October 9, 2015	March 8, 2017
German	117	105,278	900	October 8, 2015	October 21, 2017
Italian	9	18,413	2,046	March 6, 2017	October 21, 2017
Food Hoax	106	33,040	312	May 8, 2014	October 4, 2016
Liberal - Mix	136	231,021	1,699	June 18, 2013	May 31, 2018
Liberal - Black	103	357,979	3,476	June 18, 2013	March 22, 2018
Liberal - LGBT	3	1,319	440	February 26, 2016	September 27, 2016
Liberal - Muslim	5	1,084	217	January 7, 2016	May 31, 2017
Local/News	53	747,147	14,097	December 27, 2012	December 24, 2017
Portuguese	1	155	155	April 15, 2017	November 14, 2017
Conservative - Group 1	227	613,291	2,702	April 18, 2014	March 22, 2018
Conservative - Group 2	75	204,647	2,729	May 29, 2017	March 15, 2018
Conservative - Group 3	163	60,255	370	August 26, 2014	March 22, 2018
Conservative - Group 4	280	81,758	292	May 23, 2014	October 20, 2017
Russian	1,384	5,231,688	3,780	September 9, 2009	May 31, 2018
SMM/Hashtagger	105	348,367	3,318	June 6, 2014	September 20, 2017
Spanish	9	57	6	September 19, 2014	May 26, 2017
Mixed	1,002	415,367	415	May 9, 2009	December 9, 2017
<i>Totals</i>	3,841	8,489,989	2,210	n/a	n/a

Source: Authors' calculations based on data provided by the SSCI.

Appendix E: Platforms beyond Social Media: A Glimpse into the IRA’s Infrastructure
 Beyond their use of social media platforms to share and target content to specific audiences, the IRA also used US technology companies to support their operations. Metadata included in one of the data contributions gives a window into this finding, by highlighting the IRA’s use of US email providers. The set included a sample of 1,902 social media account registration email addresses. In this set, the top 3 email providers are managed by US technology companies (Yahoo, Microsoft and Google).

Table 6: Top Ten Email Providers in Registration Email Addresses

Email Provider	Number of Email Accounts
yahoo.com	1,002
hotmail.com	510
gmail.com	379
mail.ru	240
yandex.com	236
yandex.ru	221
golemico.com	216
inbox.ru	190
bk.ru	184

Note: These addresses associated with the sample of 1,902 social media personas managed by the IRA.

Source: Authors’ calculations based on data provided by the SSCI.

Appendix E: Russian Manipulation in Context

The story of Russian manipulation of US social media does not begin or end in 2016. The Russian government is heir to a long Soviet propaganda tradition stretching back nearly a century, with sophisticated techniques deployed successfully against countries around the world throughout the 20th Century. In the 21st Century, these techniques have been unleashed with a new power that combines the ability to bypass traditional media gatekeepers and reach democratic publics directly, with the scalability of automated amplification to target specific audiences in ways never possible with mass media. 21st Century social media is proving much easier to weaponize than 20th Century mass media.

The Russian government first developed its methods of manipulating online discourse in order to control the Russian public. Beginning in 2009, Harvard's Berkman Center began a multi-year project studying "RuNet," Russia's online public sphere. Two Graphika research scientists (John Kelly and Vlad Barash) helped lead the Harvard team. In a series of papers, the team observed the beginning of a distinct transformation from an early state of open and free political discourse, reflecting a great diversity of organic political opinions among the Russian people, to a narrow and polarized political environment featuring organized pro-government amplification.

In 2009, the primary online nexus of Russian political discussion was LiveJournal, a US-built blogging platform with features that many social media platforms now have. There were a number of Russian-created blogging platforms, but Russians had for years preferred to use LiveJournal. This was in part because the servers were located in the US and user data was presumably out of reach of the Russian Government, even after the platform was bought by a Russian media company in 2007. Political speech flourished on LiveJournal, and was largely absent from the native Russian platforms. Only a small portion of this speech supported the Russian government.

After 2009, the LiveJournal ecosystem began to change. Some of the company's operations were moved from California to Russia, and other platforms, such as Twitter and Facebook, attracted a younger audience and fostered new forums for political discussion. Russia slowly brought LiveJournal under its control, culminating with the movement of LiveJournal's servers to Russia in December 2016 and the ban of political and LGBTQ-related posts on the platform in April 2017, in conformity with Russian Internet laws.

In the meantime, between 2011 and 2014, Russian pro-government forces slowly built an audience on the Twitter platform. Harvard and Graphika's work on RuNet demonstrated that the space of Russian users on Twitter moved from a politically diverse one in 2010-2011 to one dominated by accounts that supported President Putin and the ruling United Russia party in 2013-2014.

Beginning in 2014, the Russian-focused Twitter data in the IRA dataset featured three periods of sustained intense activity that focused on propaganda to build domestic support for Russia's near-abroad aggressions:

- Q2-Q4 2014: Ukrainian invasion
- Q1 2015: Crimea and Donbass conflicts
- Q3-Q4 2015: Additional regional conflicts, e.g. Belarus, Slutsky

The top hashtags promoted during these three periods by Russian-focused Twitter are all Russian words; the accounts promoting these hashtags have Russian profiles with self-reported locations inside Russia. The one exception to this pattern consists of a few IRA accounts with self-reported locations in the Donbass region, currently the site of a separatist pro-Russia movement within Ukraine. This exception proves the rule: the goal of IRA's Russia-focused Twitter activity during 2014 and 2015 was not to sow propaganda amongst hostile or neutral foreign populations, but to cheerlead Russia's foreign policy for a domestic (or already friendly foreign) population.

The evolution of the topics confirms this pattern. The most popular hashtags promoted by IRA accounts in 2014 focused on aspects of the Ukrainian invasion specifically relevant to a Russian audience, specifically, anti-sanctions hashtags such as #противсанкций ("against sanctions") and hashtags alleging that Ukrainian forces shot down flight MH17 (#киевсбилбоинг, "Kiev shot down the Boeing"). These hashtags were used to support the narrative of Russia as framed and unfairly maligned, both by Kiev and by the sanctioning governments. This narrative advances the idea that any economic hardship that Russians might suffer as a result of foreign sanctions is due to unjustified aggression by foreign powers, not to any mistakes on the part of the Russian government. Indeed, the narrative has been at least partially effective, given the broad level of support Putin enjoyed among the Russian populace until recently despite worsening economic conditions.

In the first quarter of 2015, Russia-focused IRA accounts continued to cover the Ukrainian conflict, but their focus shifted to Crimea, where a massive public works project -- a construction of an auto and rail bridge connecting Crimea to the Russian mainland -- had just received a two hundred billion ruble price tag. The IRA's propaganda efforts (крымпутьна родину, -- "a path [from] Crimea to the homeland") were aimed at framing the bridge as a success for Putin's government, and not as a massive waste of resources. Other IRA hashtags (#судьбадонбасса, "fate of the Donbass", #минскиедоговоренности, "the Minsk accords") focused on the conflict in the Donbass region. Again, the fact that these hashtags were written in Russian indicates that they were aimed at a Russian audience, or at a Russia-friendly audience in the Donbass, not at a mainstream Ukrainian or other foreign audience.

Finally, in the second half of 2015, IRA's Russian-oriented activity switched to the Belarus elections (again in Russian, not in Belorussian) and on support of Leonid Slutsky, Chairman of the Duma Committee on CIS Affairs, Eurasian Integration, and Relations with Compatriots. Mr. Slutsky was one of the first Russian individuals to be targeted by US sanctions in the wake of Russian annexation of Crimea. Both topics advance an "us vs. them" narrative, supporting Russia's allies (Belarus' president Alexander Lukashenko, Slutsky) and disparaging its enemies (US, European, and other powers that imposed sanctions on Russia). The us vs. them narrative is complementary to the story that Russia's government bears no responsibility for sanctions or other sources of economic hardship, and deserves broad support from the Russian populace.

Appendix F: A Cross-Platform Case Study

Throughout this dataset, Black Matters US is an interesting case of broad, cross-platform online presence, with clear intent to translate online activity into in-person events.

The IRA's Black Matters US persona described itself as a "fast-growing group of online activists, true believers in the course and ideals of the Black rights movement." That persona created and maintained accounts across many social media platforms, including Facebook, Twitter, Instagram, YouTube, Google+, and Tumblr. The IRA also purchased Google ads to promote the associated Black Matters website and used several email addresses associated with the accounts, along with a PayPal account to receive and encourage donations to the false organization.

Black Matter US's most active account in this dataset was its Twitter presence, @blackmattersus. The account started its activities in October 2015 and grew to 20,019 followers before being suspended in July 2017. The Twitter account was used consistently to promote the Black Matters US website. Administrators also created Twitter profiles for individual "contributors" to the website. One of them, @Crystal1Johnson, was the 6th most retweeted IRA account in the dataset.

Much of the site's content, like the tweets themselves, were intended to exacerbate racial tensions in the United States. At least three Facebook accounts were also associated with the Black Matters US campaign, but we only have evidence that one of these, simply called "Black Matters," was ever linked from the @blackmattersus Twitter page. With a follower count of 222,846, the primary purpose of this Facebook page was to encourage participation in offline rallies around police violence against minorities. The much smaller Facebook dataset contains just one link to blackmattersus.com but references public rallies in at least six different cities, including New York City, Jackson (Mississippi), Atlanta, Detroit, Charleston, and Baltimore.

Figure 3: A Facebook Event Created by the Black Matters Account



Black Matters added an event.
July 6 at 9:55am · 🌐

**NEVER FORGET
ERIC GARNER**

JUL 17 **Rally In Memory Of Eric Garner**
Jul 17 - Jul 18 · 202 Bay St, Staten Island, NY 10...
484 people interested · 134 people going

★ Interested

👍 Like 💬 Comment ➦ Share

The first cross-platform link from Twitter to the Black Matters Facebook page goes as far back as December 12, 2015 and reads “Everywhere together! Join us on Facebook!” Archived pages also show that by May 2016, the @blackmattersus Twitter bio was updated to promote the same Facebook site.

The IRA administrator’s intention to direct web traffic from Twitter to Facebook also seems evident from the outbound and inbound links between the @blackmattersus Twitter handle and the Black Matters Facebook page. We do not have any evidence that Black Matters Facebook cited any Twitter accounts, but 55 cross-platform links land directly on the Black Matters Facebook site from the Twitter account. There are a further 14 links from @blackmattersus Twitter to various Facebook events - protests and rallies hosted again in some cases also by the Black Matters Facebook page.

We also observed multiple instances of activity on both Facebook and Twitter accounts at the same time. The table below shows a selected timeline for activity on March 26, 2016 for the BlackMattersUS campaign across both its Facebook and Twitter accounts.

Table 7: Vignette of BlackMattersUS Activity Timeline

Account	Time	Post
Facebook	11:18:55	"Guys, the rally is starting in less than one hour! We're meeting on Mississippi str., Mississippi State Capitol; Main entrance, near Liberty bell monument. Don't forget to take pictures! See ya'll!"
Twitter	13:45:00	Convicted Rapist Still A Threat\nhttps://t.co/OcAe5c8O74\n#rape #BlackMattersUs https://t.co/Yq3jxvuKJI
Twitter	13:53:36	3 White Students Charged With Hanger Assault On Black Classmate\nhttps://t.co/NEAAmTH91r\n#USviolence #GodBlessAmerica https://t.co/2U9OZ4uyWb
Twitter	14:24:44	Bill O'Reilly blamed Black men and boys for their own deaths\nhttps://t.co/dSJCIchmZ6\n#Foxnews #Racism https://t.co/Dz2Lgu6XrD
Twitter	14:36:55	Cops Roughed Up Black Student. Again\nhttps://t.co/4VIRv6KcAi\n#racism #BlackMatters https://t.co/9iyTTXI6dZ
Facebook	15:15:29	"Thanks everyone for attending the rally! It was really a great experience for all of us. Change is coming! We are the change! Racism is not our heritage. We refuse to honor racism and white supremacy. We're not going to give up until Confederate emblems are banned. The fight is not over."

Source: Authors' highlights based on data provided by the SSCI.

The Black Matters Facebook account was eventually suspended on August 2016. The Twitter account publicized this shutdown in a tweet on 8/22/16 and incorporated it into its messaging by accusing Facebook of “supporting white supremacy.” The next day, on 8/23/16, is the earliest activity we have for a new Facebook page called simply “BM”. There are no online traces of the new BM Facebook page before this date: we can hypothesize that the IRA simply created a the new “BM” Facebook page after the suspension of the Black Matters page.

Content from the newly created BM Facebook page (which grew to 103,013 followers) differs significantly from the suspended Black Matters page. A new logo was created, and while there are still some links to various rallies, the focus of the new Facebook page shifts in two ways.

As with the @blackmattersus Twitter page, links promoting content from the dedicated website are prominent. Whereas there are no links to the website from the Black Matters account, BM contains over 1,200 such links. Further, these links tends to be posted in “batches”: closely in time, one after the other, suggesting the consistency of quota-driven messaging. And second,

around this time, administrators began the positive black empowerment, black beauty messaging that would later feature prominently in the BlackMattersUS Instagram account.

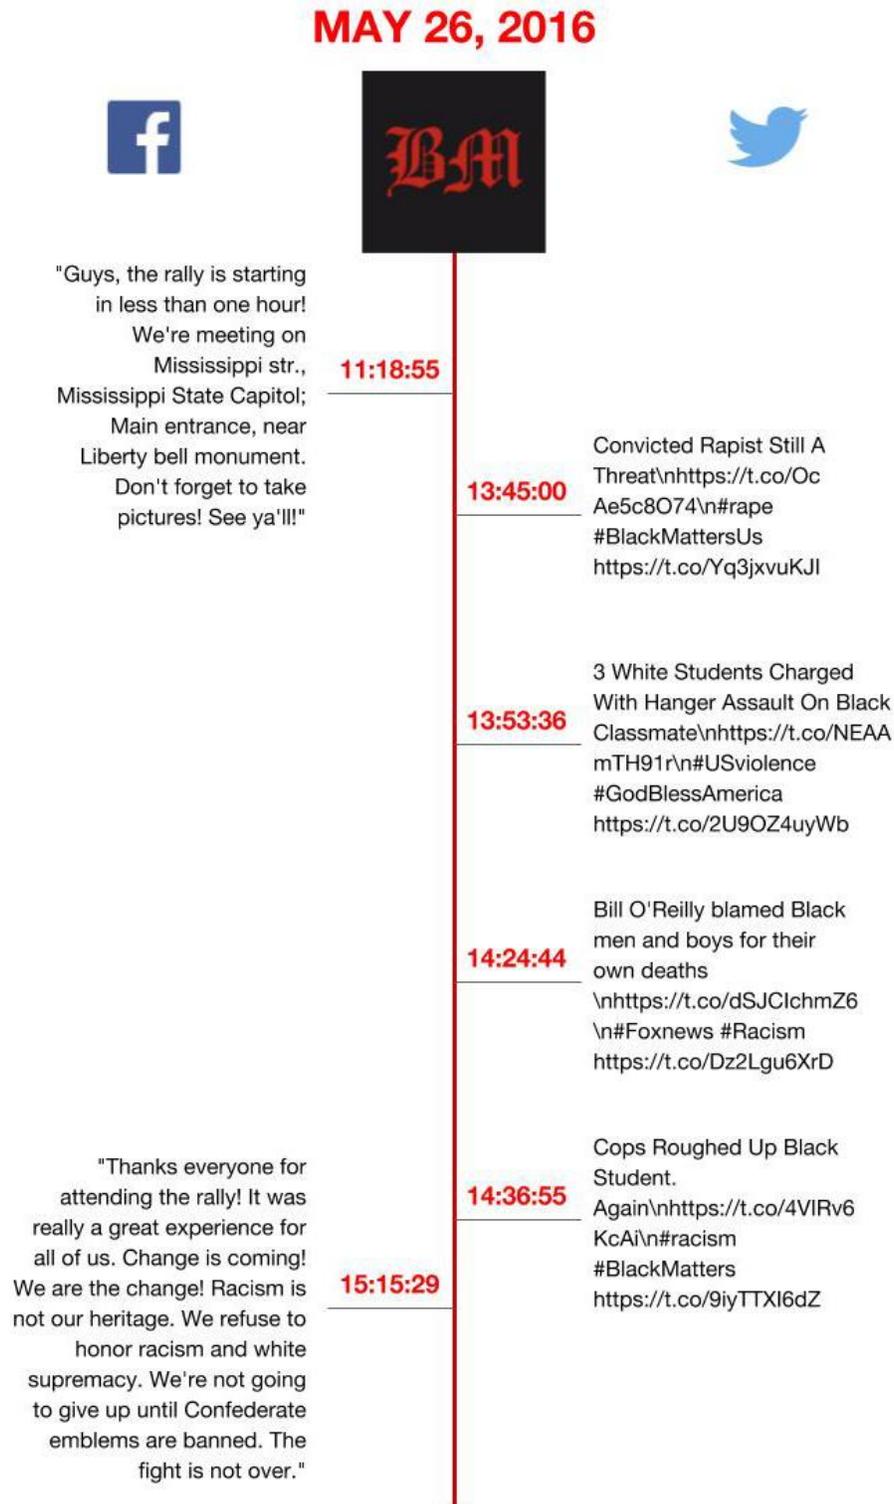
Figure 4: An Image from Black Matters US



Around the time of the initial Facebook page’s suspension, a new “meetup” feature appeared on the Black Matters website. This feature was used to coordinate rallies and protests, and was often promoted on the @blackmattersus Twitter account. It is possible that this feature was intended to fulfill the organizing role of the then-suspended Black Matters Facebook account. Unlike the Black Matters page, which appeared to the ending point for web traffic, the creation of the BM Facebook page was designed to direct viewers to the blackmattersus.com news site. The site is indeed less susceptible to the platform’s detection efforts: at the time of this report’s writing, the page is still online.

This shift in messaging toward building a positive black identity was applied most of all to the BlackMattersUS Instagram account. While online traces suggest the account was active as early as March 2016, the earliest data provided on this account’s activity to the Committee began in March 2017. This account ultimately gathered 28,466 followers and was most active in disseminating positive black empowerment messages similar to those found on the BM Facebook page. The Instagram account appears more subtle in its affiliation with the rest of the Black Matters US accounts: for instance, it did not share links to blackmattersus.com or to rallies/protest. Instead, a combination of #BlackExcellence, #blacklove, and #BlackIsBeautiful are used in the vast majority of posts.

Figure 5: Black Matters Cross Platform Activity



Source: Authors' calculations based on data provided by the SSCI.

Appendix G: Accounts Operated from the Same IP Address, by Category

We examined whether individual IP addresses created and operated accounts targeting different communities. The dataset from Twitter contained 3,841 unique Twitter screen names but only 3,138 unique "create IP" records, with 89 IP addresses matching two or more Twitter screen names.

We focused on these 89 IP addresses and the screen names matching them. For every IP address, we could identify the set of categories that the IP targeted. We were particularly interested in IP addresses that targeted multiple distinct categories. Thus, for every distinct pair of categories, we counted the number of unique IP addresses that matched accounts in both categories. The results are in Table 8, and the same analysis for Last IP (the most recent IP address to operate the account) is in Table 9.

The tables show that categories on different sides of the US political spectrum were often targeted by the same IP address. There is also overlap between US and Russian categories, as well as IPs shared between other categories, but the overlaps between US categories on opposite ends of the political spectrum are more numerous.

Table 8: Category Overlap by Shared IP Addresses (Create IP)

Category	Arabic	French	German	Italian	Food Hoax	Liberal - Mix	Liberal - Black	Liberal - LGBT	Liberal - Muslim	Local/news	Conservative 1	Conservative 2	Conservative 3	Conservative 4	Russian	SMM	Spanish	Mixed
Arabic	0	2	2	1	1	1	1	0	0	1	2	0	0	0	2	0	0	6
French		0	1	2	0	1	1	0	0	1	1	0	0	0	1	1	0	2
German			0	1	0	2	1	0	0	1	2	0	1	1	2	1	0	5
Italian				0	0	1	1	0	0	1	1	0	0	0	1	1	0	2
Food Hoax					0	0	0	0	0	0	0	0	0	0	0	0	0	1
Liberal - Mix						0	8	1	3	3	7	1	2	4	2	1	1	7
Liberal - Black							0	1	2	3	9	0	1	2	3	0	0	4
Liberal -LGBT								0	1	0	2	0	0	1	0	0	0	1
Liberal - Muslim									0	0	2	0	0	1	0	0	0	1
Local/news										0	3	0	1	0	2	0	0	1
Conservative 1											0	1	5	5	3	0	1	5
Conservative 2												0	3	2	0	0	1	1
Conservative 3													0	7	1	1	1	2
Conservative 4														0	2	1	1	9
Russian															0	0	0	10
SMM																0	0	4
Spanish																	0	1
Mixed																		0

Source: Authors' calculations based on data provided by the SSCI.

Table 9: Category Overlap by Shared IP Addresses (Last IP)

Category	Arabic	French	German	Italian	Food Hoax	Liberal - Mix	Liberal - Black	Liberal - LGBT	Liberal - Muslim	Local/news	Conservative 1	Conservative 2	Conservative 3	Conservative 4	Russian	SMM	Spanish	Mixed
Arabic	0	3	6	2	3	5	2	1	1	2	7	1	2	5	7	2	1	11
French		0	2	3	1	2	2	1	1	2	2	1	1	1	2	2	1	3
German			0	2	2	5	2	1	1	2	6	1	3	5	6	2	1	9
Italian				0	1	2	2	1	1	2	2	1	1	1	2	2	1	3
Food Hoax					0	2	1	1	1	1	2	1	1	2	2	1	1	3
Liberal - Mix						0	10	2	4	4	11	2	4	8	7	2	2	11
Liberal - Black							0	3	4	4	11	2	3	3	7	2	1	5
Liberal - LGBT								0	3	1	4	1	1	2	2	1	1	2
Liberal - Muslim									0	1	4	1	1	2	2	1	1	2
Local/news										0	4	1	2	1	3	1	1	2
Conservative 1											0	3	11	10	12	1	2	10
Conservative 2												0	43	12	2	1	2	5
Conservative 3													0	17	8	2	2	9
Conservative 4														0	9	2	2	15
Russian															0	2	1	19
SMM																0	1	5
Spanish																	0	2
Mixed																		0

Source: Authors' calculations based on data provided by the SSCI.

Appendix H: Heterophily and Homophily of IRA Accounts on Twitter

IRA accounts created social ties to each other on Twitter by mentioning each other in their Tweet posts. These social ties serve two functions: first, they make the IRA Twitter accounts appear to belong to real social communities, integrated into the same cultural spaces as their respective audiences. The objective is to make IRA accounts appear more trustworthy and real to their respective audiences. The second function is to reinforce their messaging: an IRA community can spread a particular message through Twitter conversations with each other, and thus expose their audience to aligned propaganda from multiple accounts all in the same community.

We quantified these observations by calculating the homophily and heterophily values for IRA Twitter communities. Heterophily is the ratio of observed number of social ties between two different communities to the expected number of social ties between the same communities if network connections were created at random, a common baseline in social science. Homophily is the same quantity but between a community and itself.

If the IRA accounts were operated to mimic cohesive communities, their homophily values would be higher than their heterophily values. Indeed, Table 10 shows this to be generally the case. Almost every category in the table has homophily (on the diagonal) higher than heterophily (off the diagonal). The first exception, Local/News, consists of accounts pretending to be media outlets, which do not rely on social cohesion to spread messages and thus have no need to form a strong community. The second exception, Conservative 3, appears to function in part as an auxiliary amplification community boosting the influence of Conservative 1, the primary team of conservative sock puppet accounts.

Table 10: Heterophily and Homophily by Source and Target

Category	Arabic	German	Food Hoax	Liberal - Mix	Liberal - Black	Liberal - LGBT	Liberal - Muslim	Local/news	Conservative 1	Conservative 2	Conservative 3	Conservative 4	Russian	SMM	Mixed
Arabic	58.0	0.0	0.0	0.0	0.0	0.0	0.0	7.8	0.0	0.0	0.0	0.0	0.1	0.0	0.1
German	0.3	31.2	0.0	0.0	0.0	0.0	0.0	2.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Food Hoax	0.0	0.0	29.5	0.5	0.5	0.0	0.0	0.0	1.4	0.0	0.0	0.6	0.0	0.9	0.0
Liberal - Mix	0.0	0.0	1.1	8.2	18.5	0.0	0.0	0.0	0.4	0.0	0.0	0.8	0.0	3.1	0.1
Liberal - Black	0.0	0.0	0.0	8.6	22.0	0.0	0.0	0.0	0.4	0.0	0.0	0.2	0.0	2.3	0.0
Liberal - LGBT	0.0	0.0	0.0	0.0	0.0	1280.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Liberal - Muslim	0.0	0.0	0.0	0.0	0.0	0.0	768.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Local/news	0.4	0.0	0.0	6.3	15.4	0.0	0.0	0.8	3.4	0.0	0.7	0.6	0.0	2.1	0.1
Conservative 1	0.1	0.1	1.0	1.0	1.2	0.3	0.0	0.3	9.4	0.1	0.8	0.7	0.1	6.2	0.1
Conservative 2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.8	45.7	0.0	0.0	0.0	0.0	0.0
Conservative 3	0.0	0.0	1.2	0.6	0.5	0.0	0.0	0.0	12.0	0.4	2.1	1.0	0.0	0.0	0.2
Conservative 4	0.0	0.0	1.3	3.3	1.7	0.0	0.0	0.0	1.8	0.0	0.3	8.3	0.0	1.2	0.2
Russian	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	2.8	0.0	0.0
SMM	0.0	0.0	0.2	1.5	2.3	0.0	0.0	0.0	3.2	0.0	0.1	0.3	0.0	22.3	0.2
Mixed	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.0	0.0	1.9	0.3	1.1

Note: The rows are the sources and the columns are the targets.

Source: Authors' calculations based on data provided by the SSCI.

Series Acknowledgements

The authors gratefully acknowledge the support of the European Research Council for the research project, “Computational Propaganda: Investigating the Impact of Algorithms and Bots on Political Discourse in Europe,” Proposal 648311, 2015-2020, Philip N. Howard, Principal Investigator. Additional support for this study has been provided by the Ford Foundation and the Hewlett Foundation. This particular Working Paper did not involve the analysis of human subject data. However, overall project activities were approved by the University of Oxford’s Research Ethics Committee (CUREC OII C1A15-044), with additional approvals to study the flow of misinformation over social media in the US (CUREC OII C1A17-054) and the activities of fake and trolling accounts (CUREC OII C1A17-076). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funders, the Oxford Internet Institute, or Oxford University.

This publication and its conclusions are in part based on the analysis of social media content that the authors were provided by the Senate Select Committee on Intelligence under the auspices of the Committee’s Technical Advisory Group, whose members serve to provide substantive technical and expert advice on topics of importance to ongoing Committee activity and oversight. The findings, interpretations, and conclusions presented herein are those of the authors, and do not necessarily represent the views of the Senate Select Committee on Intelligence or its Membership. All of the data provided by the US Senate was for fake accounts operated by the IRA, not human subjects, and no user data on real human subjects was used in this analysis.

Author Biographies

Philip N. Howard is Director of the Oxford Internet Institute, and a statutory Professor at Balliol College, Oxford. He writes about information politics and international affairs, and is the author of eight books, including *The Managed Citizen*, *the Digital Origins of Dictatorship and Democracy*, and *Pax Technica: How the Internet of Things May Set Us Free or Lock Us Up*. He has won multiple “best book” awards, and his research and commentary writing has been featured in the *New York Times*, *Washington Post*, and many international media outlets. Foreign Policy magazine named him a “Global Thinker” for 2017 and the National Democratic Institute awarded him their “Democracy Prize” for pioneering the social science of fake news.

Dr. Bharath Ganesh is a researcher at the Oxford Internet Institute focusing on hate speech and far right digital cultures in Europe and North America. He uses qualitative and quantitative methods to study the growth, scale, and impact of the networks and cultures that underwrite hate and how the far right exploits social media platforms. Dr. Ganesh contributes to the OII’s Computational Propaganda Research Project and the VOX-Pol Network of Excellence that focuses on violent online extremism across the world. Dr. Ganesh’s research and findings have contributed to various programs and newspapers, including BBC broadcasts, *Deutsche Welle*, and the *New Statesman*. He has shared his expertise in meetings of the United Nations Counter-Terrorism Executive Directorate and the UK Houses of Parliament. Dr. Ganesh completed his PhD in Human Geography at University College London in 2017, funded by the Bonnart Trust for its dedication to combating hate and intolerance.

Dr. Dimitra (Mimie) Liotsiou is a researcher at the Oxford Internet Institute working on computational propaganda in Philip Howard’s team. Her research focus is on developing computational models and methods for understanding and analyzing patterns of behavior in online interactions, informed by the social sciences. Her research at Oxford was covered by *Newsweek* and *TechCrunch*, among other media outlets. Her research interests relate to the areas of online social influence, causal inference, social network analysis, computational social science, and data science. She holds a PhD in Computer Science from the University of Southampton. Her PhD focused on causal inference for estimating the social influence of online communications on real-world outcomes, at the individual and collective level. This PhD research was honored with the Best Poster award at the 2016 International Conference on Social Informatics in Seattle, Washington, for the poster accompanying her full-length paper (in proceedings). She holds an MSc in Operational Research from the University of Southampton, and a BA (Hons) in Computer Science from the University of Cambridge. She is a member of the ACM, and of the London Computational Social Science Initiative.

Dr. John W. Kelly is the founder and CEO of Graphika, a social media analysis firm founded on technology he invented that blends social network analysis, content analysis, and statistics to make complex online networks understandable. Dr. Kelly is also a recognized expert on advanced computational techniques for measuring online behavior. Most recently, he was invited to provide his expert testimony on foreign interference in the U.S. presidential election before the Senate Select Committee on Intelligence. He is also an affiliate at the Berkman-Klein Center for Internet and Society at Harvard University, where he works with leading academics to design and implement empirical studies of the internet’s role in business, culture, and politics around the world. A quantitative social scientist by training, Dr. Kelly earned his Ph.D. in Communications

from Columbia University, and has also studied at Stanford and at Oxford's Internet Institute. Today, he leads a team of business professionals and data scientists to bring new analytic technologies to market via the Graphika platform.

Camille François works on cyber conflict and digital rights online. She is the Research and Analysis Director at Graphika, where she leads the company's work to detect and mitigate disinformation, media manipulation and harassment. François was previously the Principal Researcher at Jigsaw, an innovation unit at Google that builds technology to address global security challenges and protect vulnerable users. François has advised governments and parliamentary committees on both sides of the Atlantic on policy issues related to cybersecurity and digital rights. She served as a special advisor to the Chief Technology Officer of France in the Prime Minister's office, working on France's first Open Government roadmap. François is a Mozilla Fellow, a Berkman-Klein Center affiliate, and a Fulbright scholar. She holds a Master's degree in human rights from the French Institute of Political Sciences (Sciences-Po) and a Master's degree in international security from the School of International and Public Affairs (SIPA) at Columbia University. François' work has been featured in various publications, including the *New York Times*, *WIRED*, *Washington Post*, *Bloomberg Businessweek*, *Globo* and *Le Monde*.



UNIVERSITY OF
OXFORD



This work is licensed under a Creative Commons Attribution – Non
Commercial – Share Alike 4.0 International License