

The Human Consequences of Computational Propaganda

Eight Case Studies from the 2018 US Midterm Elections

EXECUTIVE SUMMARY

Katie Joseff and Samuel Woolley
Editors
Institute for the Future



 INSTITUTE FOR THE FUTURE

DIGITAL INTELLIGENCE
FUTURES LAB | iftf.org

© 2019 Institute for the Future. All rights reserved.

introduction

Computational propaganda continues to be used to influence elections worldwide despite the efforts of technology firms to combat the flow of misinformation, disinformation, and automated politicking. A wide range of political groups now combines the use of social media platforms, bots, and big data analytics with communication tactics like disinformation and politically motivated trolling in order to manipulate public opinion (Woolley & Howard, 2018). Internally, there is evidence that US political campaigns, activists, and individuals have leveraged the amplification power of bots to boost particular messages and drown out others during major political events (Metaxas & Mustafaraj, 2012; Woolley & Guilbeault, 2017). Externally, foreign governments have used the same tools to sow polarization and sway US elections (Allcott & Gentzkow, 2017).

In many instances, groups have deployed computational propaganda broad-scale—simply using bots or organized groups of people to flood platforms like Twitter, Reddit, YouTube, and Facebook with false, defamatory, or misleading content. In other circumstances, however, shrewder actors (both foreign and domestic) have leveraged these same tools and tactics against particular social groups and issue-focused voting blocs. The results for democratic communication can be devastating. The 2018 US midterm elections saw a spike in computational propaganda attacks against these often-vulnerable populations. While the type of individual “psychographic” political targeting imagined by groups like Cambridge Analytica remains in development, larger-scale efforts to use social media to divide and conquer already embattled US social groups, including ethnic and religious minority populations, are on the rise. There are very real human consequences to the usage of computational propaganda. When computational propaganda flows to groups with already marginalized voices or to those on the precipice of political alienation, the results range from chilling effects and disenfranchisement to psychological and physical harm.

This series of research papers works to track the use of computational propaganda during the 2018 US midterms in order to set the stage for similar research and inquiry during future events. We conducted case studies of eight different social and issue-focused groups. Where possible, we worked with authors, researchers, and experts who represented the groups they covered. This was a deliberate choice aimed toward expanding the cadre of voices speaking about issues of computational propaganda. The eight case studies on Computational Propaganda in the US are:

- Anti-Climate Science
- Anti-Immigrant
- Anti-Latinx
- Black Women Gun Enthusiasts
- Jewish Americans
- Moderate Republicans
- Muslim Americans
- Women’s Reproductive Rights

literature review

We found, among all groups studied, that:

- Computational propaganda was used during the midterm elections as a tactic to manipulate, threaten, and divide groups.
- This propaganda took various forms, ranging from issue-specific false news or disinformation to racially and religiously specific defamatory content.
- Vulnerable demographics within already marginalized social groups—including young people, the elderly, military veterans, and those for whom English is a second language—were often particular targets of disinformation and harassment.
- In already embattled minority and social groups, both human groups and bot armies spread harassing content and misleading information about voting and worked to foment in-group fighting both online and off-line.
- In issue-focused groups and other social groups—gun rights/control advocates and moderate Republicans, for instance—propagandists worked to polarize, inflame, and misinform.

Beyond what can be said generally, computational propaganda took a variety of forms specific to each case. The details are provided in the “Findings” section (page 6).

Analyses during the 2016 US presidential election produced significant evidence of both automated and organized human astroturf group efforts to control political conversation using social media. Armies of bots and people leveraged social media to spread false information during the contest with varying degrees of success. Researchers found that automated bot accounts generated one-fifth of the entire conversation about the election on Twitter (Bessi & Ferrara, 2016). Of these human-mimicking accounts, pro-Trump bot accounts out-tweeted pro-Clinton bot accounts by a ratio of 5-to-1 (Kollanyi, Howard & Woolley, 2016). Three months prior to the election that year, pro-Trump fake news stories were shared 30 million times on Facebook, in comparison to just 8 million shares of pro-Clinton fake stories (Allcott & Gentzkow, 2017).

The data of American citizens was also compromised for the purposes of digital political manipulation in 2016. The personal data of an estimated 87 million Facebook users was exploited by Cambridge Analytica, a data consulting firm that specialized in political behavioral microtargeting and was employed by the Ted Cruz and Donald Trump campaigns (Kang & Frenkel, 2018). Firms like Cambridge Analytica, as well as foreign organizations, leveraged the age-old strategy of “divide and conquer” online—exploiting US diversity to encourage disunion (Howard et al., 2017). For instance, more than half of the Facebook advertisements bought by the Russian Internet Research Agency during the 2016 US elections focused on race (Penzenstadler, Heath & Guynn, 2018). These roughly 2,000 advertisements, contained statements like “Colorado Springs cops apprehend two men for ‘driving while black’” and “Don’t mess with TX border patrol, always guided by God” (Keating, Schaul & Shapiro, 2017; Danner, 2018). They led to 25 million impressions by Facebook users (Penzenstadler, Heath & Guynn, 2018). Divisive issue advertisement campaigns, which were purchased by a variety of foreign and domestic actors,

targeted battleground states and specific individuals—low-income individuals saw more advertisements about immigration and race, and whites saw more advertisements about nationalism and immigration than any other racial group (Kim et al., 2018).

While these statistics provide a sense of the scale and pervasiveness of computational propaganda, they do little to convey the magnitude of damage that disinformation attacks continue to inflict upon targeted social groups and issue publics. The online sphere is already inhospitable for marginalized individuals in the United States—27% of African Americans, 30% of Latinx, 35% of Muslims, and 63% of LGBTQ+ community members have experienced harassment online due to their identity (ADL, 2019). Correspondingly, women are twice as likely as men to have experienced online harassment due to their gender (Duggan, 2017). To what extent do disinformation and broader computational propaganda campaigns work to foment online and offline harassment and violence? Do disinformation and harassment campaigns discourage or galvanize political participation by marginalized and minority groups?

In short, this series works to answer the question: What are the human consequences of computational propaganda?

“ Women are twice as likely as men to have experienced online harassment due to their gender.

methods

These case studies involved a mixed methods approach of qualitative and quantitative social science research. Authors conducted interviews, fieldwork, and focus groups with the social and issue-focused groups they studied to garner an in-depth understanding of how particular communities experience political communication and computational propaganda over social media. Where possible, they used these findings to inform large dataset analyses that focused on the kinds of social media content from various platforms being targeted at particular communities during the dates surrounding the 2018 US midterms. Teams worked collaboratively with Digital Intelligence Lab data scientists to both gather and analyze data pertaining to their groups.

Papers went through three drafts of revision after feedback from the Digital Intelligence Lab leadership team and affiliate researchers. All papers went through a fourth critique process of peer-review from external experts. Each paper was reviewed a minimum of two times by leaders in the field of computational propaganda from outside academic institutions and think tanks. The review process was non-blind and was specifically focused on feedback regarding the papers' mixed methodologies and particular social focuses.

We sought to collaborate with case study authors who self-identify as members of the groups they studied. Prioritizing the work of authors who represent the communities studied is not only important with regard

to narrative agency, but it also arguably allows for a more nuanced consideration of social context that enables researchers to more effectively uncover the systems of power central to computational propaganda. Regrettably, we were unable to find authors from the African American or LGBTQ communities who were available to work with us on this series. As both communities were highly targeted in the 2016 election, it is imperative that they are included in future iterations of this research (US House of Representatives Permanent Select Committee on Intelligence, 2017; Kim et al., 2018).

We also worked with representatives from civil society groups (including ADL and Witness) focused on the study of defamation, harassment, political speech, human rights, internet policy, cyber-security, and social media. Experts affiliated with the following academic entities authored or co-authored papers: University of Washington, University of California at Berkeley, Stanford University, University of Pennsylvania, West Virginia University, University of Texas at Austin, University of Minnesota, and University of Illinois at Chicago.

“ African American and LGBTQ communities were highly targeted in the 2016 election and it is imperative that they are included in future iterations of this research.

findings

Our collective research uncovered a variety of tactics used to attack and disenfranchise particular groups, including disinformation, targeted political trolling campaigns, bot-driven censorship, and even intra-group political harassment. There are six overarching themes: (1) Human social media users, not bots, produced the majority of harassment—but bots continue to be used to seed and promote coordinated disinformation narratives; (2) Adversarial groups are co-opting images, videos, hashtags, and information previously used or generated by social and issue-focused groups—and then repurposing this content in order to camouflage disinformation and harassment campaigns; (3) Disinformation campaigns utilize age-old stereotypes and conspiracies—often attempting to foment both intra-group polarization and external arguments with other groups; and (4) Social media companies’ responses to curb targeted harassment and disinformation campaigns have not served to effectively protect the groups studied here from such content.

Humans are the primary harassers, while bots seed and amplify disinformation narratives.

Across the series, researchers found that the majority of harassment perpetrated on social media was enacted by real humans without automation. Of the social media accounts spreading harassment, automated accounts (bots) made up the minority in the studies of Jewish Americans, Muslim Americans, Latinxs, and moderate Republicans. When bots were used, however, they worked to significantly amplify highly partisan or harassing content. Bots are still being used to suppress communication between particular groups in roadblock efforts to generate noise around activist or social group hashtags (Woolley, 2018). Individuals interviewed for the Jewish American study consistently reported that, out of all the social media platforms, they experienced the worst harassment on Facebook—where political bots are less of a documented problem than on Twitter, Reddit, YouTube, and other more open platforms. Moreover, respondents said that—due to Facebook’s real-name policy—their harassers on that platform

generally posted harassing, hateful, and threatening content using their full legal names without reprisal from moderation teams or authorities (Woolley & Joseff, 2019).

Though humans tend to perpetuate the lion’s share of disinformation and harassment levied against social groups and issue-focused groups, politically motivated bots still play a serious role. These political bots are often used to seed and fertilize content among human communities and give that content the illusion of widespread popularity. The authors of our Moderate Republicans study found that roughly 10 percent of the users discussing the 2018 Arizona Primary on Twitter were bots. Despite this small number, they had an outsized effect on communication on that platform due to their automated-speech capability. Bots spread not only more disinformation news articles than human users but “more divisive and extreme content than human users” (Kumleben, 2019). The group studying pro-choice and pro-life campaigners found that bots were highly central to the networks discussing those topics (Nonnecke et al., 2019). Interestingly, pro-life bots were more likely to send and retweet harassing language, while pro-choice bots were more likely to perpetuate political divisiveness (Nonnecke et al., 2019).

Adversarial groups are co-opting social groups’ and issue publics’ images and messages—and reappropriating them as harassment tools.

The author of the study on climate change found that a key rhetorical strategy of fossil fuel companies is “cultural co-opting,” which entails digitally rebranding coal as “green” and carbon dioxide as “a benign gas that is essential for all life” (Goldenberg & Bengtsson, 2016 qtd. in Guilbeault, 2019). Similarly, the studies on immigration groups, Muslim Americans, Jewish Americans, and the Latinx community discovered co-option of in-group narratives and content by adversarial groups.

Pro-immigration groups rely heavily on images and video to invigorate audiences and communicate important information online (McAweeney & Makam, 2019). Trolls and anti-immigration groups are corrupting this important mode of communication, photoshopping “know your rights” flyers from pro-immigration groups in order to spread disinformation about what to do when stopped by ICE. These same adversaries also leverage tactics of targeted harassment by altering photos of pro-immigration officials to make them look like ICE officers or white supremacists (McAweeney & Makam, 2019). The immigration group found that 100 percent of tweets co-opted by adversaries for the purposes of disinformation and harassment had larger audiences than the original, non-manipulated information.

“Many Muslim American hashtags have been ‘taken over by haters’ and are being used for targeted harassment and abuse.”

Many Muslim American hashtags (including the primary hashtag #CAIR of the Council on American-Islamic Relations, one of the largest Muslim American advocacy groups in the US) have been “taken over by haters” and used for targeted harassment and abuse (Pakzad & Salehi, 2019). CAIR members have attempted to “reclaim the hashtag” but feel that they are “constantly dealing with the possibility that there’s just so much more keyboard power on the other side” (Pakzad & Salehi, 2019). The authors of the Jewish case also discovered co-option of terms and ideas. Anti-Semitic trolls online have co-opted Hebrew words, including shoah and goyim, and turned them into derogatory terms (Woolley & Joseff, 2019). The group studying Latinx Americans found that there is a “data void”

regarding Latinx and political involvement on websites like Reddit (Flores-Saviaga & Savage, 2019). Trolls take advantage of this data void to create coordinated disinformation campaigns about the US Latinx community, often impersonating experts in attempts to sow discord and false narratives.

Disinformation campaigns utilize age-old stereotypes and conspiracies—often in attempts to pit marginalized groups against one another.

Researchers studying pro-choice and pro-life campaigners found that one of the most common false social media narratives was the claim that Margaret Sanger founded Planned Parenthood with the intention of carrying out race-based abortions (Nonnecke et al., 2019). One of the most widely disseminated pro-life tweets uncovered during this research claimed that Hillary and Bill Clinton “support the extermination of minority babies” and attend KKK rallies (Nonnecke et al., 2019). In the study of disinformation targeting moderate Republicans, the most commonly bot-shared disinformation website was America-hijacked.com, an anti-Semitic conspiracy website dedicated to attacking the ‘Israel Lobby’ (Kumleben, 2019). Throughout debate about climate science, the conspiracy persisted that climate scientists cannot be trusted because they falsely represent climate science to further their own careers and because they are funded by liberal donors with partisan agendas (Guilbeault, 2019). Divisive and misinformative stories featuring racist stereotypes—particularly regarding the supposed “migrant caravan” of terrorists and criminals—were prominent in the studies on immigration activists, Latinx Americans, Muslim Americans, and Jewish Americans. Conspiracies about governmental manipulation of the public and secret cabals running the world—a “new world order”—were featured in both the study of gun rights/gun control activists and Jewish Americans.

Social media platforms' responses to targeted harassment and disinformation campaigns continue to be inadequate.

Interviewees (including leaders from a variety of social and religious communities) across multiple studies expressed serious concerns with social media companies' current efforts to deal with disinformation and harassment. Dozens of interviewees representing different social groups described incidents in which they or their loved ones had reported racial or religious harassment and subsequently received no feedback or support from the firms. Other interviewees stated that they faced so much harassment online that it would be a "full-time job" to block and report all malicious users" (Pakzad & Salehi, 2019).

“Several Jewish American interviewees who had been doxxed reported that they received little help from social media companies.

A repeated critique from interviewees across multiple groups was that they felt the platforms lacked general diversity and that they also failed to present nuanced understandings of minority user groups. Interviewees said they also felt platforms failed to adequately protect minority populations from the disproportionate amount of computational propaganda they face. As an example, one interviewee noted that a Twitter account audit in June erroneously led to the removal of a key

organizer of the Black Lives Matter marches, seriously disrupting a planned protest. The interviewee felt that it could have been avoided if Twitter had clearer lines of communication with social rights activists (McAweeney & Makam, 2019).

Interview subjects across cases were also concerned about the radicalization of white supremacists and other alt-right groups on social media. In discussing closed and insular messaging and blogging platforms (e.g., Gab), a Muslim American interviewee stated: "They're just sharing a lot of this hate speech within and among themselves which is very scary because we don't know where that's going to lead them" (Pakzad & Salehi, 2019).

A particularly concerning abuse of social media is the malicious phenomenon of doxxing—the release of an individual's private information (and sometimes the information of their family and friends) on social media without their consent. This release of personal information enables harassers to have an outsized impact on the security and comfort of activists on social media. Doxxing can lead to torrents of hate on all social media platforms, private phone numbers, and email addresses. It can also facilitate in-person stalking and offline violence. Several Jewish American interviewees who had been doxxed reported that they received little help from social media companies, some even had to hire personal security and cybersecurity firms (Woolley & Joseff, 2019). They and several interviewees across other groups expressed frustration that the companies had no clear point of contact or protocol for helping them stem the onslaught of harassment in such cases.

a path forward

Our case studies reveal serious gaps in understanding the problems, as well as glaring shortcomings in policies (both internal to the companies and governmental) for mitigating effects. Taken together, our analyses suggest an imminent need for further investigations into the human impact of computational propaganda. It is important to quantify the magnitude of disinformation and bot-driven trolling, and big data analyses are helpful in this regard. However, to fully understand the repercussions of computational propaganda upon democratic processes including elections, researchers must conduct field work with the social and issue-focused voting blocs that are disproportionately targeted during these events. Researchers must speak to the people who are affected, including social, cultural, and religious leaders. Such qualitative research with groups provides sociological, psychological, and cultural depth to quantitative analysis of large datasets. Because of this, we advocate for a mixed methods approach to studying computational propaganda and for a centering of human experience in this work.

The 2020 US elections

The computational propaganda attacks upon social and issue-focused groups during the 2018 US midterms have serious implications for the 2020 presidential election. The particulars of these disinformation and harassment campaigns provide insight into the social media political communication tactics voters are likely to experience in 2020.

Those concerned with computational propaganda in the forthcoming presidential elections should look out for the following:

- Attacks upon social and issue-focused groups.
- Disinformation campaigns against these groups aimed at disenfranchising voters.
- Coordinated harassment campaigns aimed at silencing minority populations.
- Threats of violence against minority social and religious groups.
- The use of bots on Twitter, Reddit, and other sites to kindle and inflame disinformation and political attack campaigns.
- The use of Facebook group pages as well as community- and location-targeted advertisements as a mode for spreading disinformation and inflaming political tensions there.
- The use of Islamophobia as a digital rallying cry for white supremacist and other extremist groups attempting to recruit more moderate voters.
- Large-scale disinformation and harassment campaigns targeting Latinxs and African Americans.
- Coordinated vilification campaigns aimed at Jewish Americans, women's reproductive rights groups, and immigration activists.
- Coordinated campaigns aimed at polarizing issue-focused groups, including gun rights and gun control groups and climate change supporters and deniers.
- Growth of "divide and conquer" and radicalization campaigns upon moderate Republicans and Democrats.
- Foreign use of computational propaganda during the US contest by countries other than Russia.

Responding with policy

There are clear channels for fighting back against group-focused computational propaganda. This paper series provides evidence that lays out the groundwork for technological, legal, and civil actions focused upon creating a safer and more accessible online public sphere. Our mapping of the issues at hand not only displays the human cost of computational propaganda but also highlights specific avenues for intervention and prevention.

Possible broad-scale policy interventions include:

- Adapting current First Amendment prohibitions on hate speech to online platforms and technologies—e.g., building upon the federal cyberstalking statute (18 USC § 2261A).
- Modifying Section 230 of the Communications Decency Act of 1996 that protects social media platforms from legal liability.
- Identifying state-sponsored trolling and hate-based computational propaganda as new forms of human rights abuse under US and international law.
- Establishing a specific federal commission aimed at monitoring digital political communications and expanding FEC, FTC, and FCC purview to include this work.
- Passing of acts aimed at establishing public and democratic accountability for social media firms.

Additionally, in following and cataloging recent policy responses to the rise of computational propaganda, we have noticed several themes and have developed corresponding recommendations:

- Federal governments, including the US government, are failing to pass laws that meaningfully work to combat online disinformation, political automation, and politically motivated digital harassment. The Honest Ads Act (a low-hanging piece of legislation that would force companies to disclose who buys political ads during elections) is a prime example of a bill that has stalled out due to the current political climate and a lack of true commitment from technology firms (Ravel, Woolley & Sridharan, 2019).

- Where governments have succeeded in passing legislation (e.g., Germany), the laws have been heavy-handed and lacking in technical expertise. Leveraging multimillion-dollar fines against firms that fail to take down disinformation or propaganda might force them to hire more people to police content, but such fines do not address flaws at the heart of technology design caused by a lack of ethics and the Silicon Valley's imperative for growth. Policies to combat computational propaganda must start somewhere, but legislators should be wary of regulation that doesn't pay heed to empirical research or the mechanics of computing. There will be few easy fixes and quick solutions to this issue.
- Policymakers must work closely with technology designers, social scientists, security professionals, and information experts to craft laws that challenge computational propaganda holistically rather than mitigating perceived symptoms. For instance, it is not feasible to simply ban bots on a site like Twitter or across the web. These tools are infrastructural and play myriad roles online beyond those that are political. A reasonable law might advocate for clear tagging of potentially automated social media profiles in order to mitigate attempts to manipulate, but it should not aim to eradicate the use of automated software on these platforms writ large.
- It is time for the US FEC, FTC, and FCC to more effectively regulate digital political communications during elections. More political advertising—and, consequently, more harmful propaganda—is being spread online than ever before, and the money spent in this domain continues to grow.

Designing for democracy

Technology firms along with governments bear significant responsibility for the current state of global political communication. Facebook, Twitter, Google, and others might have originally envisioned their products as technologies and not media, but they can no longer avoid the fact that the algorithms they have constructed have curated the news and information people use to form understandings of the world. These companies may argue that they are not the arbiters of truth, but they are certainly mediators and curators of information—and they should be held accountable as such. They are not service providers, à la telecommunications companies; they are creators of *social* media. Their algorithms are gameable and can result in biased or hateful trends. What spreads on their platforms can give rise to very real offline political violence, to say nothing of psychological damage from abusive online behaviors such as trolling and harassment.

In short, the content that the global public views over social media is being manipulated by powerful entities, and these companies must act. The same laws that govern one-to-many tools like television and radio may not be applicable to these many-to-many platforms. There is a serious need for these enormously wealthy technology companies to make some clear self-regulatory and user-based moves, including:

- Creating new user tools on social media that allow for greater control and better company-user communication with regard to harassing and derogatory messaging—beyond muting, blocking, reporting, and unfollowing.
- Standardizing platform responses to onslaughts of trolling due to doxxing and other targeted attacks.
- Deplatforming white supremacy and hate speech.
- Doing more about both politically motivated and general identity theft.
- Addressing accounts that use automation to artificially boost political content and political bots; this should include labels for highly automated accounts. The effect upon trending algorithms—and broader online political discussion—needs to be mitigated.

“ The institutional lack of focus on the ethical consequences of a technology sector that has experienced unbridled growth and unprecedented innovation has brought us to where we are today.

To be clear, our analyses suggest that computational propaganda is not a problem that can be dealt with by platforms alone. Our research shows that self-regulation will not be enough to truly check this issue. Nor do our findings suggest that these corporations, which have the prerogative to protect trade secrets (including algorithms) as well as shareholders, should be expected to consistently make decisions in the interest of the general and political health of society. Computational propaganda should not be an issue that is foisted upon already embattled and underfunded civil society entities and news outlets. While it is important for companies like Facebook to partner with groups like the Associated Press and Snopes, social media firms must not simply pass potential disinformation—and the role of verification—to these entities without access to the data and processes that informed content flagging in the first place. There is room for sensible third-party oversight of internal processes for both prioritizing certain kinds of information for users and flagging other data as potentially harmful.

Finally, this isn't simply a user problem. It's not solvable through a new feature or app that helps people improve their media literacy—although such tools can certainly be useful. Our work suggests that larger systems—technological, governmental, and otherwise—bear the bulk of responsibility for computational propaganda. The institutional lack of focus on the ethical consequences of a technology sector that has experienced unbridled growth and unprecedented innovation has brought us to where we are today. We need an ethical operating system. It's time for technologists to design with both democracy and human rights in mind. This means not just building more democratic, less-biased algorithms. It means social media companies and other technology firms must work with and invest in a wide variety of US and international social and issue-focused groups—especially marginalized, and often unheard, minority populations—to build a more equitable public sphere offline as well as online.

references

- Allcott, H., & Gentzkow, M. (2017). Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, 31(2), 211–236.
- Bessi, A., & Ferrara, E. (2016). Social bots distort the 2016 US Presidential election online discussion.
- Burns, J. (2018, May 19). Whistleblower: Bannon Sought To Suppress Black Voters With Cambridge Analytica. *Forbes*. Retrieved from <https://www.forbes.com/sites/janetwburns/2018/05/19/cambridge-analytica-whistleblower-bannon-sought-to-suppress-black-voters/#36b8c4f87a95>
- Danner, C. (2018, May 12). More Than Half of Russian Facebook Ads Focused on Race. Retrieved April 19, 2019, from *New York Magazine* website: <http://nymag.com/intelligencer/2018/05/more-than-half-of-russian-facebook-ads-focused-on-race.html>
- Duggan, M. (2017, July 11). Online Harassment 2017 | Pew Research Center. Retrieved April 19, 2019, from <https://www.pewinternet.org/2017/07/11/online-harassment-2017/>
- Howard, P. N., Kollanyi, B., Bradshaw, S., & Neudert, L. M. (2018). Social media, news and political information during the US election: Was polarizing content concentrated in swing states?. *arXiv preprint arXiv:1802.03573*.
- Kang, C., & Frenkel, S. (2019, January 29). Facebook Says Cambridge Analytica Harvested Data of Up to 87 Million Users. *The New York Times*. Retrieved from <https://www.nytimes.com/2018/04/04/technology/mark-zuckerberg-testify-congress.html>
- Keating, D., Schaul, K., & Shapiro, L. (2017, November 1). Analysis | The Facebook ads Russians showed to different groups. Retrieved April 19, 2019, from *Washington Post* website: <https://www.washingtonpost.com/graphics/2017/business/russian-ads-facebook-targeting/>
- Kim, Y. M., Hsu, J., Neiman, D., Kou, C., Bankston, L., Kim, S. Y., ... & Raskutti, G. (2018). The stealth media? Groups and targets behind divisive issue campaigns on Facebook. *Political Communication*, 35(4), 515–541.
- Kollanyi, B., Howard, P., & Woolley, S. C. (2016). Bots and automation over Twitter during the second US presidential debate.
- Metaxas, P. T., & Mustafaraj, E. (2012). Social media and the elections. *Science*, 338(6106), 472–473.
- Online Hate and Harassment: The American Experience. (2019, February 13). Retrieved April 19, 2019, from *Anti-Defamation League* website: <https://www.adl.org/onlineharassment>
- Penzenstadler, N., Heath, B., & Guynn, J. (2018, May 11). We read every one of the 3,517 Facebook ads bought by Russians. Here's what we found. Retrieved April 19, 2019, from *USA TODAY* website: <https://www.usatoday.com/story/news/2018/05/11/what-we-found-facebook-ads-russians-accused-election-meddling/602319002/>
- Ravel, A. M., Woolley, S. C., & Sridharan, H. (2019). Principles and Policies to Counter Deceptive Digital Politics. *Institute for the Future & MapLight*, 26.
- Social Media Advertisements [US House of Representatives Permanent Select Committee on Intelligence]. (2017). Retrieved April 19, 2019, from <https://intelligence.house.gov/social-media-content/social-media-advertisements.htm>
- Woolley, S. (2018). *Manufacturing Consensus: Computational Propaganda and the 2016 United States Presidential Election* (Ph.D. Dissertation). University of Washington, Seattle, WA.
- Woolley, S. & Howard, P. (2018). *Computational Propaganda: Politicians, Political Parties and Political Manipulation over Social Media*. Oxford University Press, Oxford: UK.
- Woolley, S., & Guilbeault, D. (2017). *Computational Propaganda in the United States of America: Manufacturing Consensus Online* (No. 2017.5; p. 37). Oxford, United Kingdom: Oxford Internet Institute, University of Oxford.

INSTITUTE FOR THE FUTURE

Institute for the Future is the world's leading futures thinking organization. For over 50 years, businesses, governments, and social impact organizations have depended upon IFTF global forecasts, custom research, and foresight training to navigate complex change and develop world-ready strategies. IFTF methodologies and toolsets yield coherent views of transformative possibilities across all sectors that together support a more sustainable future. Institute for the Future is a registered 501(c)(3) nonprofit organization based in Palo Alto, California.

DIGITAL INTELLIGENCE LAB

The **Digital Intelligence Lab** at Institute for the Future is a social scientific research entity conducting work on the most pressing issues at the intersection of technology and society. We examine how new technologies and media can be used to both benefit and challenge democratic communication.



INSTITUTE FOR THE FUTURE
201 Hamilton Avenue
Palo Alto, CA 94301
www.iff.org