

Beware Fake News

How Influence Operations Challenge Liberal Democratic Governments

Author: Eric Jardine

Cet essai est disponible en français.

The 2016 US presidential election was a tumultuous time. In the weeks and months leading up to that Tuesday, November 8, social media sites, such as Twitter and Facebook, were flooded with “fake news” (Howard et al. 2017). Investigations following the election of Donald Trump as the forty-fifth president of the United States revealed that extensive foreign influence had played a role during the campaign, its efforts aimed largely at affecting the course of the election. Most fingers pointed directly to the Russian Federation and the regime of President Vladimir Putin as the most likely culprits (National Intelligence Council 2017).

This was not by any means the first use of social media in influence operations. A few years earlier, for example, the Islamic State terrorist organization (ISIS) used extensive Twitter campaigns to spread propaganda, encourage radicalization and recruit foreign soldiers for its war in Iraq and Syria (Klausen 2015).

Influence operations, whether launched by governments or non-state actors, existed long before social media, but what is new about contemporary influence operations is their scale, severity and impact, all of which are likely to grow more pronounced as digital platforms extend their reach via the internet and become ever more central to our social, economic and political lives. Such efforts represent a clear cyber security challenge. Yet, democracies, which depend on the open and free sharing of information, are particularly susceptible to the poison of influence operations that spread fake news, disinformation and propaganda. The whole edifice of democratic governance is based on the assumption of an informed citizenry with a common sense of facts, shared public narratives and a solid trust in the information provided by institutions. This entire assemblage is threatened by carefully crafted influence operations and will only grow worse as new “deep fake” technologies come into play.

The Scope of the Problem

By one false account, in 2016, Democratic Party nominee Hillary Clinton and her chief of staff, John Podesta, were operating a child sex ring out of a pizza parlour’s basement in Washington, DC. What started as a malicious internet rumour quickly morphed into a social media trend. The hashtag #pizzagate went viral as thousands of accounts tweeted “evidence” both for and against the story. Many of these tweets originated outside of the United States, with disproportionately large clusters coming from the Czech Republic, Cyprus and Vietnam. Shortly after the election, this fictitious online tale made a sinister cross-over into the physical world, as one of the story’s followers, Edgar Welch, drove to

Washington with an assault rifle. He entered the pizzeria, demanding to see the basement (the building does not have one) and fired off three shots. What began as online disinformation had taken a terrible turn (Fisher, Cox and Hermann 2016).

The pizzagate story is just one illustration of an increasingly prevalent problem of online influence operations by foreign governments and non-state actors. While a healthy ecosystem involves the free flow of information and interpretation of facts, large swaths of online influence operations to date, particularly as they are directed toward the West, can be colloquially called “fake news,” meaning content that is “intentionally and verifiably false, and [that] could mislead readers” (Allcott and Gentzkow 2017, 213). Beyond subverting the facts, fake news also plays another role. It is crafted to resonate with its readers. Such resonance does not arise purely through information. Resonance can also be based on sentiment or a reader’s sense of its truth, creating what could be called a folkloric element (Frank 2015).

If fake news was only about spreading incorrect information, then those who believed such stories would have to be either ignorant or undiscerning about news in general, or willfully ingesting false content. Viewing fake news as a genre of folklore, as Russell Frank has proposed (*ibid.*), raises a third possibility, that fake news is appealing because it delivers a moral narrative or confirms sentiments that people already hold. From this perspective, the ISIS social media propaganda about the corruption of the West (Klausen 2015) or the fake news stories about the health of Hillary Clinton during the 2016 election (Milligan 2016) share a common foundation: they propagate “alternative” information *and* present a moral narrative that people holding similar views can latch on to.

Influence operations have a long history, but their potential reach today is scaled up by the enormous user populations of digital platforms and applications. (Photo: Unsplash.com)

Influence operations using messages combining these informational and political-parable-like qualities can be launched by state actors, non-state actors or some combination of both. Efforts at influencing information environments have a long history, but today, the potential scale of influence operations is decisively affected by new digital platforms with vast numbers of users. Facebook alone has roughly 2.25 billion users. Twitter has 336 million. Mobile messaging applications that allow users to share threads and stories likewise capture huge proportions of the internet-using population, with 100 million Telegram users, 1.5 billion WhatsApp users and 1.0 billion Viber users, not to mention the numerous smaller messaging applications that exist online.

The scaling effect of social media gives a simple boost to terrorist organizations that seek to radicalize individuals or recruit foreign fighters. For example, ISIS ran a highly advanced online influence operation. On Twitter, this process spanned geography, with carefully selected fighters in Syria and Iraq tweeting photos that were then vetted and shared by third parties and individuals linked to ISIS but living in the West (Klausen 2015). Through this simple gatekeeper methodology, ISIS was able to put forward a coordinated influence campaign, designed to showcase a skewed image of the glories of war and life under ISIS.

On other digital platforms, such as YouTube, ISIS used the huge user population (1.8 billion users) and hours of consumed videos (up to one billion hours daily) to spread propaganda videos to glorify its terrorist agenda (Gillespie 2018). As Tarleton Gillespie put it, “ISIS has proven particularly skilled at using social media in this way, circulating glossy recruitment magazines and videos documenting the beheading of political prisoners and journalists” (*ibid.*, 55). The goal was to reach those individuals who might be swayed by ISIS’s messages and encouraged to undertake homegrown operations or become a foreign fighter.

The increased scale of information operations also plays out via new socio-technical algorithmic assemblages. Algorithmic bots, specially designed programs that use computer processing power to spread content via fake user accounts, have helped to generate and pollute the online information ecosystem. Such bots are particularly active during political events. The 2016 US election swung partially due to changed, and somewhat unexpected, shifts in voter preferences in Michigan. Within this key battleground state, as research from the Computational Propaganda program at Oxford indicates, non-professional news (fake news) was shared more frequently via social media than professional, mainstream news (Howard et al. 2017). More troubling still, news produced by reputable media outlets (*The New York Times*, for example) hit its lowest point as a proportion of content the day before the election (ibid.). These trends were exacerbated by bots' activity.

The growing sophistication of artificial intelligence (AI) and machine-learning algorithms also points to a potential new qualitative change in influence operations. Generally, people tend to trust the written word somewhat less than they do audio and, in particular, video media. A news story might say that Hillary Clinton is ill, but the story would appear more believable if Clinton were to say so herself — or at least if she were to *seem* to say so. AI can now be leveraged to generate so-called “deep fake” videos, which actually involve faked video of a person saying fake news (Giles 2019). Deep fakes are hard to spot and will greatly increase the qualitative impact of fake news and foreign influence operations.

With enhanced scale, increasing automation and the capacity for pernicious deep fakes, influence operations by foreign governments and non-state actors have gained a new edge. Operations that would have been manageable in a predigital age are now a very real challenge to liberal democratic regimes.

The Challenge

Democracy is fundamentally based on trust — trust of each other, trust in institutions and trust in the credibility of information. Influence operations, in particular those run by foreign governments or malicious non-state actors, can pollute an information environment, eroding trust and muddying the waters of public debate.

The discourse surrounding the 2016 US presidential election is a case in point. Debate during the campaign was marked by a high level of rancour. Since the election, survey respondents have indicated that they feel civility and trust in major institutions within the United States have declined as the opposing ideological camps have hardened their positions. For example, one survey found that fewer than 30 percent of people trusted media institutions and, more broadly, fully 70 percent of respondents thought that there was less civility (Santhanam 2017).

Fake news and other influence operations are made more powerful by “filter bubbles” (Pariser 2012). The term describes the result of the algorithmic machinations that lead people into relatively contained online information ecosystems of their own making. Once within such a bubble, people tend to get more of what they like, based on their earlier online choices, whether those are funny YouTube videos of cats or ideologically infused podcasts and posts. The troubling part is that the commercial aim of platform filters — namely, to give people what they want to encourage consumption of content — tends to play out badly in the political space. They lead people to hear their own message rather than others' points of view, in an echo chamber reinforced by algorithms. While democracy requires the free exchange of information and ideas, filter bubbles tend to isolate users. In a filtered environment, information does not circulate widely and freely.

News produced by reputable traditional outlets hit its lowest point as a proportion of shared media content the day before the 2016 US presidential election. (Photo: Osugi / Shutterstock.com)

Solutions and Ways Forward

Malicious influence operations are a growing problem, exacerbated by social media platforms that enable the scale-up of misinformation, disinformation, propaganda and information disruption operations and new algorithmic technologies that might potentially cause us to even distrust our own eyes.

Modest, but meaningful, changes are possible and necessary. Broadly, countering the problem means addressing three aspects: exposure, receptivity and counter narrative.

Exposure is at the core of the problem of fake news and other forms of influence operations. A person might be psychologically ripe for radicalization, but, without exposure to ISIS's message, may never tip over the edge. Likewise, an electorate's exposure to fake news during an election cycle may affect political discourse and even electoral outcomes. Simply put, reducing exposure to influence operations reduces their effects.

In liberal democracies, where freedom of expression is enshrined as a fundamental right, governments often cannot directly censor the information being shared online. Furthermore, the primary infrastructure for disseminating information during an influence operation (such as social media platforms) is owned and operated by private companies. So, although governments are limited in their ability to constrain exposure, the companies that own the platforms are not. Facebook, Twitter and YouTube (run by Alphabet) can all directly control what sort of information flows across their networks.

While platforms historically avoided explicit content moderation, and to some extent still do, arguing that they are not publishers, consumers have begun to express a desire for some moderation of more extreme and polarizing content, such as white supremacist content or fake news stories. The platforms are able to oblige (Gillespie 2018). These systems can moderate, and so control, exposure to information through two complementary methods. First, platforms now leverage their vast user bases, encouraging users to flag and report content that is potentially objectionable. The platforms then evaluate the flagged content. If it is found to be in violation of a platform's terms of service or community guidelines, it can be removed and the account that posted it can be banned (ibid.). Besides these human-driven methods, many firms are using automated detection systems to flag and pull down content. With more data, these approaches will improve further still. Through both measures, the platforms are working to limit the worst effects of malicious influence operations by reducing exposure to such content as ISIS beheading videos, "conspiracy videos" and hate-infused tweets.

Another method for countering influence operations is to build up people's online "immunity" so that they have less receptivity to misleading, false and polarizing information. Broad-based educational initiatives that aim to increase user awareness of fake content might be helpful, if hugely costly. Inoculating key points (people) within a network is likely more effective and cheaper (Christakis and Fowler 2011). Targeted engagement with individuals at the centre of networks (high network centrality scores, in social network analysis terms) could help promote immunity of the herd and reduce receptivity to fake content (Halloran et al. 2002).

Finally, governments and traditional media institutions can work to create their own narratives of events that can counter the influence operations of others. The effectiveness of such counter narratives is conditional upon the trust that users place in their sources, so initiating these efforts swiftly to stem the tide of disruptive influence operations aimed at diminishing user trust is key. Their effectiveness is also likely a function of how well traditional producers adapt to changing media. The current social networking ecosystem is driven by clickbait content. Sending out boring titles into this sort of maelstrom will likely fall flat.

If done right, meeting the messages of foreign influence operations with a counter narrative can have a positive effect on the perceptions of internet users. The public's willingness to believe climate change denial stories, for example, is reduced if exposure to that disinformation is quickly paired with countering narratives that highlight the flaws in anti-climate change science and point to the climate change consensus that exists within the scientific community (Cook, Lewandowsky and Ecker 2017). In short, while refuting disinformation is an ongoing struggle rather than a quick win, governments — helped by platforms — can counter one influence operation with another. Doing so can help preserve trust, while also retaining the free flow of information that is at the core of liberal democratic governance.

Conclusion

Influence operations targeting liberal democratic regimes are deeply troubling. They disrupt the twin bedrocks of effective democratic governance: the free flow of information and trust. These campaigns can be undertaken by malicious foreign governments who aim to sow chaos, or by non-state actors, such as ISIS, who seek to radicalize disaffected individuals in the West. Countering these operations is both necessary and possible. Such efforts require the engagement of not only governments but also the platforms. Working together, these actors can preserve liberal democratic governance by minimizing exposure to fake news and other influence operations, promoting user immunity and promulgating counter narratives to misinformation.

Works Cited

Allcott, Hunt and Matthew Gentzkow. 2017. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives* 31 (2): 211–36. <https://web.stanford.edu/~gentzkow/research/fakenews.pdf>.

Christakis, Nicholas A. and James H. Fowler. 2011. *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives — How Your Friends' Friends' Friends Affect Everything You Feel, Think, and Do*. New York, NY: Back Bay Books.

Cook, John, Stephan Lewandowsky and Ullrich K. H. Ecker. 2017. "Neutralizing misinformation through inoculation: Exposing misleading argumentation techniques reduces their influence." *PLoS ONE* 12 (5): e0175799. <https://journals.plos.org/plosone/article/file?id=10.1371/journal.pone.0175799&type=printable>.

Fisher, Marc, John Woodrow Cox and Peter Hermann. 2016. "Pizzagate: From rumor, to hashtag, to gunfire in D.C." *The Washington Post*, December 6.

Frank, Russell. 2015. "Caveat Lector: Fake News as Folklore." *The Journal of American Folklore* 128 (509): 315–32. doi:10.5406/jamerfolk.128.509.0315. www.researchgate.net/publication/281601869_Caveat_Lector_Fake_News_as_Folklore.

Giles, Martin. 2019. "Five emerging cyber-threats to worry about in 2019." *MIT Technology Review*. www.technologyreview.com/s/612713/five-emerging-cyber-threats-2019/.

Gillespie, Tarleton. 2018. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. New Haven, CT: Yale University Press.

Halloran, M. Elizabeth, Ira M. Longini Jr., Azhar Nizam and Yang Yang. 2002. "Containing Bioterrorist Smallpox." *Science* 298 (5597): 1428–32. <http://science.sciencemag.org/content/298/5597/1428>.

Howard, Philip N., Gillian Bolsover, Bence Kollanyi, Samantha Bradshaw and Lisa-Maria Neudert. 2017. "Junk News and Bots during the U.S. Election: What Were Michigan Voters Sharing Over Twitter?" *Data Memo 2017.1*, March 26. Oxford, UK: Project on Computational Propaganda. <https://comprop.oii.ox.ac.uk/wp-content/uploads/sites/89/2017/03/What-Were-Michigan-Voters-Sharing-Over-Twitter-v2.pdf>.

Klausen, Jytte. 2015. "Tweeting the *Jihad*: Social Media Networks of Western Foreign Fighters in Syria and Iraq." *Studies in Conflict & Terrorism* 38 (1): 1–22. www.tandfonline.com/doi/abs/10.1080/1057610X.2014.974948.

Milligan, Susan. 2016. "Hillary's Health: Conspiracy or Concern?" *U.S. News*, August 15. www.usnews.com/news/articles/2016-08-15/hillarys-health-conspiracy-or-concern.

National Intelligence Council. 2017. *Assessing Russian Activities and Intentions in Recent US Elections*. Office of the Director of National Intelligence, Intelligence Community Assessment 2017-01D, January 6. www.dni.gov/files/documents/ICA_2017_01.pdf.

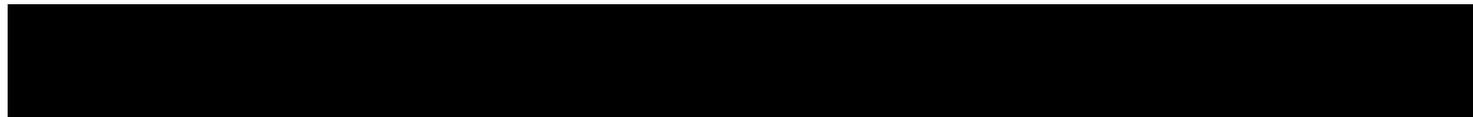
Pariser, Eli. 2012. *The Filter Bubble: What the Internet Is Hiding from You*. London, UK: Penguin Books.

Santhanam, Laura. 2017. "New poll: 70% of Americans think civility has gotten worse since Trump took office." *PBS News Hour*, July 3. www.pbs.org/newshour/politics/new-poll-70-americans-think-civility-gotten-worse-since-trump-took-office.

The opinions expressed in this article/multimedia are those of the author(s) and do not necessarily reflect the views of CIGI or its Board of Directors.

About the Author
[Eric Jardine](#)

Eric Jardine is an assistant professor of political science at Virginia Tech and a fellow at CIGI.



CIGI

[Africa \(292\)](#)
[Arctic \(309\)](#)
[Artificial Intelligence \(325\)](#)
[Big Data \(884\)](#)
[Brexit \(897\)](#)
[Central Banking \(457\)](#)
[China \(1348\)](#)
[Clean Technology \(232\)](#)
[Climate Change \(2614\)](#)
[Conflict Management \(1342\)](#)
[Cyber Security \(1001\)](#)
[Democracy \(801\)](#)
[Digital Currency \(168\)](#)
[Platform Governance \(1203\)](#)
[Diplomacy \(1213\)](#)
[Diversity & Inclusion \(227\)](#)
[Emerging Technology \(432\)](#)
[Financial Systems \(1593\)](#)
[Future of Work \(121\)](#)
[G20/G7 \(1808\)](#)
[Gender \(463\)](#)
[Human Rights \(480\)](#)
[IMF \(1164\)](#)
[India \(230\)](#)
[Indigenous Peoples \(522\)](#)
[Innovation \(577\)](#)
[Innovation Economy \(944\)](#)
[Intellectual Property \(960\)](#)
[Internet Governance \(1150\)](#)
[Investor State Arbitration \(403\)](#)
[Monetary Policy \(608\)](#)
[NAFTA/CUSMA \(1415\)](#)
[Oceans \(142\)](#)
[Patents \(135\)](#)
[Productivity \(52\)](#)
[Refugees \(718\)](#)
[Sovereign Debt \(593\)](#)
[Standards \(193\)](#)
[Surveillance & Privacy \(623\)](#)
[Sustainable Development \(690\)](#)
[Systemic Risk \(67\)](#)
[Trade \(2994\)](#)
[Traditional Knowledge \(231\)](#)
[WTO \(620\)](#)

-
-
-