DISINFORMATION AND DEMOCRACY: THE INTERNET TRANSFORMED PROTEST BUT DID NOT IMPROVE DEMOCRACY

Anya Schiffrin[1]

Professor at Columbia University School of International and Public Affairs

Abstract: Recent years have seen a marked shift in global attitudes toward social media platforms. In 2011, Facebook was hailed as a platform that would bring democracy to the world, Google was breaking new ground in convenience and access to information, and the protests taking place in Iran, Egypt, Tunisia, Bahrain, and many other countries were spurred in part by bloggers and social media commentators who used the platforms to galvanize people and encourage them to take to the streets. But, by 2017, we had learned that although the Internet transformed protest, it has not much improved democracy. Moreover, we learned again the lesson that the post-Cold War democracies had apparently forgotten: that misinformation and propaganda are powerful, and that repeating "big lies" can persuade susceptible people of nonsensical and dangerous ideas. This essay will examine the various sources and forms of disinformation that are most prevalent in today's political and media environment, the implications of this new reality on democracy, and the ways in which government can and must respond.

2016 was the year that public opinion turned against social media and big tech companies. In 2011, Facebook was hailed as a platform that would bring democracy to the world. We were grateful to Google. The protests in Iran, Egypt, Tunisia, Bahrain, and many other countries were spurred in part by bloggers and social media commentators who used social media to galvanize people and encourage them to take to the streets.

By 2017, we had learned that although the Internet had transformed protest, it has not much improved democracy. Moreover, we learned again a lesson that the post-Cold War democracies had apparently forgotten: that misinformation and propaganda are powerful and that repeating "big lies" can persuade susceptible people of all kinds of nonsensical and dangerous ideas. This should not have been a surprise, but critic Norah Ephron once said that "people have a shocking capacity

to be surprised by the same things over and over again." The question now is what to do. Regulation of social media platforms comes up repeatedly, but of what kind is less clear.

Of course, it was not all boundless optimism in 2011. Even before the Arab Spring, critics like Evgeny Morozov had warned that the Internet could be used as a tool of surveillance, and Cass Sunstein and Markus Prior had warned that giving everyone the right to select the news they wanted to read would compromise the marketplace of ideas.[2] What wasn't clear at the time was the scale of the disinformation that would flood the Internet and the effect this could have on voting. It didn't seem plausible that people would be so susceptible to lies on the Internet and that they would resist reasoned attempts to explain facts, that truth would seem not to matter. By 2017, it became clear that anger over social inequality had turned into the conflation of privilege with expertise, and that many hated experts. Global demagogues stoked the fires of this hatred with constant attacks on the judiciary, the media, science, climate change scientists, and any institution that could undermine their agendas.[3] At the time of this writing, it does not seem an exaggeration to say that disinformation spread by social media has undermined the functioning of democracy globally. But if social media is undermining our ideas of democracy, how can we solve the problem without also undermining the processes of democracy?

Looking Back at the Optimistic Debates of 2010 and 2011

A few months before the Arab Spring, two books were published that discussed the role of digital technology on society and democracy. One, *The Net Delusion*, by Evgeny Morozov, got widespread attention for its robust attack on the "techno optimists" who were foolish enough to believe that the likes of Facebook could bring about social change and force governments to become more accountable and democratic. "A dictator who answers his cell phone is still a dictator," Morozov wrote. Further, he argued, sophisticated authoritarian regimes would be able to use the web not just for propaganda purposes but to track their opposition; so that digital technology was actually helping authoritarian regimes survive—a point that, with the passage of time, no longer seems novel.

But it was a book that got far less attention that turned out to be more immediately prescient. Using a data set of Islamic countries from around the world, political science professor Philip Howard argued that digital technology was bringing communities together, providing vast amounts of information to closed societies; and forcing governments to become more accountable. This in itself was making the world more democratic. A few months after these two books appeared, the Arab Spring revolutions cemented the idea that digital technology was a force

for political change. The new conventional wisdom became that the Internet had dispersed the power of international organizations and governments and emerging communities online have undermined traditional state authority. From mobile money to crowd-mapping crises and bringing citizens together to report on news, distribute information, and organize politically, digital technology had the potential to leave obsolete power structures behind. Scholars such as Zeynep Tufekci and Jennifer Earl argued that the "affordances" of the web had transformed protest in part by lowering the amount of time, effort, and money it required, and by making it easier to gather large numbers of disparate people from around the world into new communities.[4] The recent scholarship makes it clear that the nature of activism and protest had changed and that the web was not just recreating earlier forms of protest.

2016, and the Values of Big Tech: Make Billions by Spreading Millions of Dangerous Lies

By 2016, it was apparent that something had gone very wrong; many of the optimists of 2010 and 2011 had changed their thinking, warning of the dangers of digital technology. Wael Ghonim, whose Facebook pages are credited with galvanizing the protests in Egypt, declared that the web had become a "mobocracy." Along with Emily Parker, Ghonim launched a site called Parlio that was meant to encourage civilized and expert discourse online about vital topics of the day.[5] The site never garnered a large following but was bought by Quora and eventually closed down. Philip Howard began studying bot activities and disinformation during the 2016 elections in Europe and the US and came up with some startling numbers about the amount of disinformation shared over Twitter.[6] Howard and his colleagues at the Computational Propaganda Research Project at the Oxford Internet Institute looked at

> "How many Facebook users saw what kinds of disinformation, when they saw it, and how often this took place is unclear, in part because Facebook consistently refused to provide information to researchers about what political advertisements it displayed and who saw them."

seven million tweets that used hashtags related to the 2016 election between November 1 and November 11 in 16 swing states. After developing a typology based on the URLs included in these tweets, which sorted all tweets into six categories including professional political content such as government and campaign sources, professional news outlets, and polarizing and conspiracy content;

Howard and his colleagues found overwhelming levels of news from Russian outlets, Wikileaks, and "junk news" sources flooding Twitter just before the 2016 US presidential elections.[7] Howard and his colleagues also noted that in these 16 swing states, levels of "junk" and polarizing news exceeded those of the United States as a whole.

How many Facebook users saw what kinds of disinformation, when they saw it, and how often this took place is unclear, in part because Facebook consistently refused to provide information to researchers about the political advertisements it displayed and who saw them. According to Howard:

> At this point Facebook is the single most important platform for public life in the vast majority of countries. Its advertising algorithms allow politically moti-vated advertisers to reach a purposefully selected audience. Unfortunately, the company provides no public record of the political advertisements it serves to users, and there is no systematic way for analysts to measure the spread of junk news. For other kinds of media, political candidates must declare their spon-sorship and file copies with the FEC. In the US election, for example, Trump spent 70 million on Facebook ads we'll never see.[8]

Without knowing what people saw, how many times, and for how long, it is difficult to know whether or how much of an effect disinformation had on voting patterns. An early study was released in early 2017 by economists Hunt Allcott and Matthew Gentzkow and concluded that "fake news" had no effect on the US elec-tions.[9] The earlier draft of the study, however, was based on some assumptions that seemed shaky at best, including the assumption that one piece of fake news was comparable to 36 negative campaign advertisements.[10] By the time the paper was published, it had already been circulated widely and read closely by senior people at Facebook. Another questionable part of the study includes the authors giving their subjects "placebo" headlines so as to compare their judgement of real news with their judgement of fake news. However, understanding media effects is far more complicated than doing a scientific experiment in which randomized control groups are necessary. It is possible that there was a backfire effect resulting from the placebo headlines, or even what scholars call "misinformation persistence."[11]

For all of these reasons, we don't actually know whether or how disinformation affected the 2016 elections. The use of social media to move public opinion is a relatively new phenomenon, and the speed and volume of the incorrect messages transmitted by social media may be unprecedented.

However, the example of Fox News is instructive. When Fox News began, it was assumed that people watching it would be influenced to vote for the Republican party, and early research suggested that this was indeed the case.[12] In the fall of 2017, a more definitive paper was published in the American Economic Review that solved the question of causation.[13] Consumption of Fox News pushed

voters to vote for Republican candidates. The prevalence of junk news also suggests that voters with low exposure to information participated in elections at high rates, perhaps a different trendline from past elections, when the assumption was that those who were uninformed and didn't follow the news were also the people who did not vote.

Propaganda, lies, and "truthiness" have been around for hundreds of years and used by many political candidates, corporations, and religions to persuade and mislead. What is different today is the speed and volume of disinformation. We simply do not know what it means for the electorate when millions of Russian propaganda messages are targeted at swing states. We can guess, but the research has not yet been done, and the information is not available for us to know with absolute certainty.

Even so, it is not too early to take action. When there is a strong possibility of danger, society must act. Governments don't wait for everyone to be in an automobile accident before mandating that air bags should be put in every car. Now is the time to consider low-hanging policy measures that may help the situation. All policy involves tradeoffs, but we need to consider measures are acceptable to a democratic society.

MIT researcher Yochai Benkler said at an October 2016 talk at Columbia University's School of International and Public Affairs, that there are five parties circulating fake news:

1. Bodies thought to be close to the Russian government that circulate propaganda and disinformation with the intent of sowing confusion and distrust;

2. Right-wing US groups such as Breitbart;

3. Groups that make money from circulating disinformation such as the notorious Macedonians profiled by Buzzfeed in the fall of 2016;[14]

4. Formal campaigns using behavioral marketing tools, such as Cambridge Analytica;[15]

5. Peer-to-peer distribution networks, including the far-right activists of 4Chan.

Benkler says:

> The problem is potentially sufficiently serious enough that we should spend a lot of money quickly to figure out what is happening so we know what measures to take. At a minimum we should support transparency in political advertising and that should include anyone paid to comment online on (or) spread political information even if it's by marketing companies as well as the commercial equivalent of the 50-cent army.

While millions of dollars are being spent on research, there should be a focus on policy prescriptions that can be put in place quickly. One example of policy-making which involves an acceptable tradeoff was the bi-partisan bill introduced in the US Senate in October 2016 that would have required social media companies

to tell the Federal Elections Commission the source of funding for online political advertisements.[16] Just as we require disclosure for political advertising on television, so too should we disclose the sources of paid information online. Technology expert Julia Angwin notes that such laws would not cover paid commenters, but Benkler says that the law could be expanded to cover people paid to comment online.

In many cases, countries with laws against hate speech and incitement will need to find democratic ways to enforce them online so that the fight against disinformation does not become an excuse for corporate and government censorship. Asking big tech companies to deal with the problem on their own opens the way to corporate censorship, free expression advocates have consistently warned.[17] On the other hand, it is important not to let technology companies use free speech as an excuse not to take action.

In their comprehensive report on "Information Disorder" for the Council of Europe, Claire Wardle and Hossein Derakhshan discuss the need to create cultures of truth and provide recommendations for governments, journalists, technology companies, and other parties.[18] While many of the fixes being proposed, including media literacy education and changes in ownership models, are long term, changing norms and culture will be part of getting back to a culture of truth and evidence. The topic is too important to leave to tech companies to handle alone and without disclosure. Government, academia, and civil society need to lead the conversation on how to address the problem of the millions of lies and propaganda mentions that can spread so quickly on social media. E-Bay creator and philanthropist Pierre Omidyar wrote in an October 2017 op-ed:

> Just as new regulations and policies had to be established for the evolving online commerce sector, social media companies must now help navigate the serious threats posed by their platforms and help lead the development and enforcement of clear industry safeguards. Change won't happen overnight, and these issues will require ongoing examination, collaboration and vigilance to effectively turn the tide.[19]

In fact, more needs to be done. Columbia law professor Tim Wu believes that Facebook should become a nonprofit or public benefit corporation, and Columbia University professor Joseph E.Stiglitz argues that Facebook is similar to a public utility and should be strongly regulated. Privacy, taxation and distribution of dis/misinformation are all areas where there needs to be strong global regulation of the tech and social media sectors. There are a number of options. It seems that there is an emerging consensus around cracking down on tax avoidance and protecting privacy. However, it is likely that some European countries will pass laws regulating the dissemination of free speech, just as Germany has done. Regulations will look different in different countries, as it will be hard to obtain a European-wide

consensus. Facebook, of course, argues that ultimately these regulatory mechanisms will be copied by authoritarian regimes.

Implications for Democracy

The implications of these developments for democracy are enormous. Is the Internet killing our democracy and paving the way for uninformed mob rule? Democracy rests upon the assumption of an educated populace; this is part of why public education is so important. Understanding the important issues of the day, as well as government representatives' positions on these issues, is necessary for citizens to participate actively in a democracy. Without this knowledge, voting decisions may be arbitrary, and government can be based on capturing voters or pandering, and can cease to be a truly functioning democracy.

The problem of misinformation on the Internet has come at a dangerous time, when growing resentment over inequality and the worsening state of the American middle class have stoked a deep mistrust in institutions of education, science, and media that have traditionally served to keep "false facts" and demagoguery at bay. Citizens are increasingly turning to the Internet, a forum for distributing information that does not adhere to typical standards of truth, scientific inquiry, and evidence-based news and information. At the same time, the institutions that have typically distributed information to citizens are being usurped. Consequently, for the average American citizen, distinguishing between true and false information has only become more difficult.[20]

Polling data suggests that we live in a country where a large part of the populace is either unable or unwilling to accurately educate themselves on the reality of their country and leaders. An uninformed citizenry of this type is unable to act in its own best interest when electing leaders and representatives. If misinformation and fake news campaigns truly do frustrate citizens' attempts to educate themselves—or, even worse, actively manipulate citizens into believing false information—then the very foundations of democracy are at risk.♕

Anya Schiffrin is the director of the Technology, Media, and Communications specialization at Columbia University's School of International and Public Affairs where she teaches courses on media and development and innovation. Among other topics, she writes on journalism and development, as well as the media in Africa and the extractive sector. Schiffrin spent 10 years working overseas as a journalist in Europe and Asia and was a Knight-Bagehot Fellow at Columbia University's School of Journalism in 1999–2000. Schiffrin is on the Global Board of the Open Society Foundations and the advisory board of the Natural Resource Governance Institute and the American Assembly. Her most recent books are African Muckraking: 75

Years of African Investigative Journalism *(Jacana 2017) and* Global Muckraking: 100 Years of Investigative Reporting from Around the World *(New Press, 2014).*

NOTES

[1]   Acknowledgments: Anamaria Lopez (CC 2018) assisted with research and writing the conclusion.

[2]   Evgeny Morozov, *The Net Delusion: The Dark Side of Internet Freedom* (New York: Public Affairs, 2011); Cass Sunstein, *#Republic: Divided Democracy in the Age of Social Media* (Princeton: Princeton University Press, 2017); Markus Prior, *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections* (New York: Cambridge University Press, 2007).

[3]   Heather Brookes in conversation, 2016.

[4]   Jennifer Earl and Katrina Kimport, *Digitally Enabled Social Change* (Cambridge, MA: MIT Press, 2011); Zeynep Tufekci, *Twitter and Tear Gas: The Power and Fragility of Networked Protest* (New Haven: Yale University Press, 2017).

[5]   Nathan Gardles, "Wael Ghonim: We Have a Duty to Use Our Social Media Power to Speak the Truth," *Huffington Post*, 29 October 2016, https://www.huffingtonpost.com/entry/wael-ghonim-social-media_us_580e364ae4b000d0b157b53a.

[6]   Lisa-Maria Neudert, Bence Kollanyi, and Philip N. Howard, "Junk News and Bots during the German Parliamentary Election: What are German Voters Sharing over Twitter?" (Oxford, UK: Project on Computational Propaganda, 2017); Craig Timberg, "Propaganda Flowed Heavily Into Battleground States Around Election, Study Says," *Washington Post*, 28 September 2017, https://www.washingtonpost.com/business/technology/2017/09/27/32855bba-a3a0-11e7-ade1-76d061d56efa_story.html?utm_term=.03a84c04b6f5.

[7]   Phillip M. Howard et al., "Social Media, News and Political Information during the US Election: Was Polarizing Content Concentrated in Swing States?" (Oxford, UK: Project on Computational Propaganda, 2017).

[8]   Phillip Howard in conversation, 29 July 2017.

[9]   Hunt Alcott and Matthew Gentzkow, "Social Media and Fake News in the 2016 Election," *Journal of Economic Perspectives* 31, no. 2 (2017), 211-236.

[10]   The paper includes the placebo headlines but does not explicitly mention the 1:36 fake news to campaign ad ratio, stating instead, "How much this affected the election results depends on the effectiveness of fake news exposure in changing the way people vote. As one benchmark, Spenkuch and Toniatti (2016) show that exposing voters to one additional television campaign ad changes vote shares by approximately 0.02 percentage points. This suggests that if one fake news article were about as persuasive as one TV campaign ad, the fake news in our database would have changed vote shares by an amount on the order of hundredths of a percentage point. This is much smaller than Trump's margin of victory in the pivotal states on which the outcome depended."

[11]   Man-pui S. Chan, Christopher R. Jones, Kathleen Hall Jamieson, and Dolores Albaraccín, "Debunking: A Meta-Analysis of the Psychological Efficacy of Messages Countering Misinformation," *Psychological Science* 28, no. 1 (2017), 1–16. This paper conducted a meta-analysis of studies on how to effectively debunk misinformation. Among other things it found that simply printing a correction or warning does not bring people to change their minds but rather can result in "misinformation persistence." The authors point to Schwarz et al. (2007), who found that corrections often inadvertently strengthen the misinformation they intend to contest when they merely ask people to "consider the opposite" of stated facts. This risk is lowered only when a well argued, detailed debunking message is offered (Jerit 2008). The authors concluded that detailed corrections produce a stronger debunking effect than non-detailed ones. However, they can also inadvertently perpetuate misinformation.

[12]   Stefano DellaVigna and Ethan Kaplan, "The Fox News Effect: Media Bias and Voting," NBER Working Paper 12169 (April 2006), http://www.nber.org/papers/w12169.

[13]   Gregory J Martin and Ali Yurukoglu, "Bias in Cable News: Persuasion and Polarization," *American Economic Review* 107, no. 9 (2017), 2565-99.

[14]   Craig Silverman and Lawrence Alexander, "How Teens in the Balkans are Duping Trump Supporters with Fake News," *Buzzfeed News*, 3 November 2016, https://www.buzzfeed.com/craigsilverman/how-macedonia-became-a-global-hub-for-pro-trump-misinfo?utm_term=.ygVV9NJpA#.jqb-

nQ58jk.

[15]   Carole Cadwalladr, "The Great British Brexit Robbery: How Our Democracy Was Hijacked," *The Guardian*, 7 May 2017, https://www.theguardian.com/technology/2017/may/07/the-great-british-brexit-robbery-hijacked-democracy.

[16]   Kenneth P. Vogel and Cecilia Kang, "Senators Demand Online Ad Disclosures as Tech Lobby Mobilizes," *New York Times*, 19 October 2017, https://www.nytimes.com/2017/10/19/us/politics/facebook-google-russia-meddling-disclosure.html.

[17]   Courtney C. Radsch, "Deciding Who Decides Which News is Fake," Committee to Protect Journalists, 14 March 2017, https://cpj.org/blog/2017/03/deciding-who-decides-which-news-is-fake.php.

[18]   Claire Wardle and Hossein Derakhshan, "Information Disorder: Toward an Interdisciplinary Framework for Research and Policy Making" (Strasbourg, France: Council of Europe, 2017).

[19]   Pierre Omidyar, "Pierre Omidyar: 6 Ways Social Media Has Become a Direct Threat to Democracy," *Washington Post*, 9 October 2017, https://www.washingtonpost.com/news/theworldpost/wp/2017/10/09/pierre-omidyar-6-ways-social-media-has-become-a-direct-threat-to-democracy/?utm_term=.194674e5e885.20, Pharr, Putnam, and Dalton, "A quarter century of declining confidence," 5-25.

[20]   Camila Domonoske, "Students Have 'Dismaying' Inability to Tell Fake News From Real, Study Finds," *National Public Radio*, 23 November 2016, https://www.npr.org/sections/thetwo-way/2016/11/23/503129818/study-finds-students-have-dismaying-inability-to-tell-fake-news-from-real.