CONVENED BY GEORGE P. SHULTZ

GOVERNANCE IN AN EMERGING NEW WORLD

THE INFORMATION CHALLENGE TO DEMOCRACY

Niall Ferguson
Joseph S. Nye

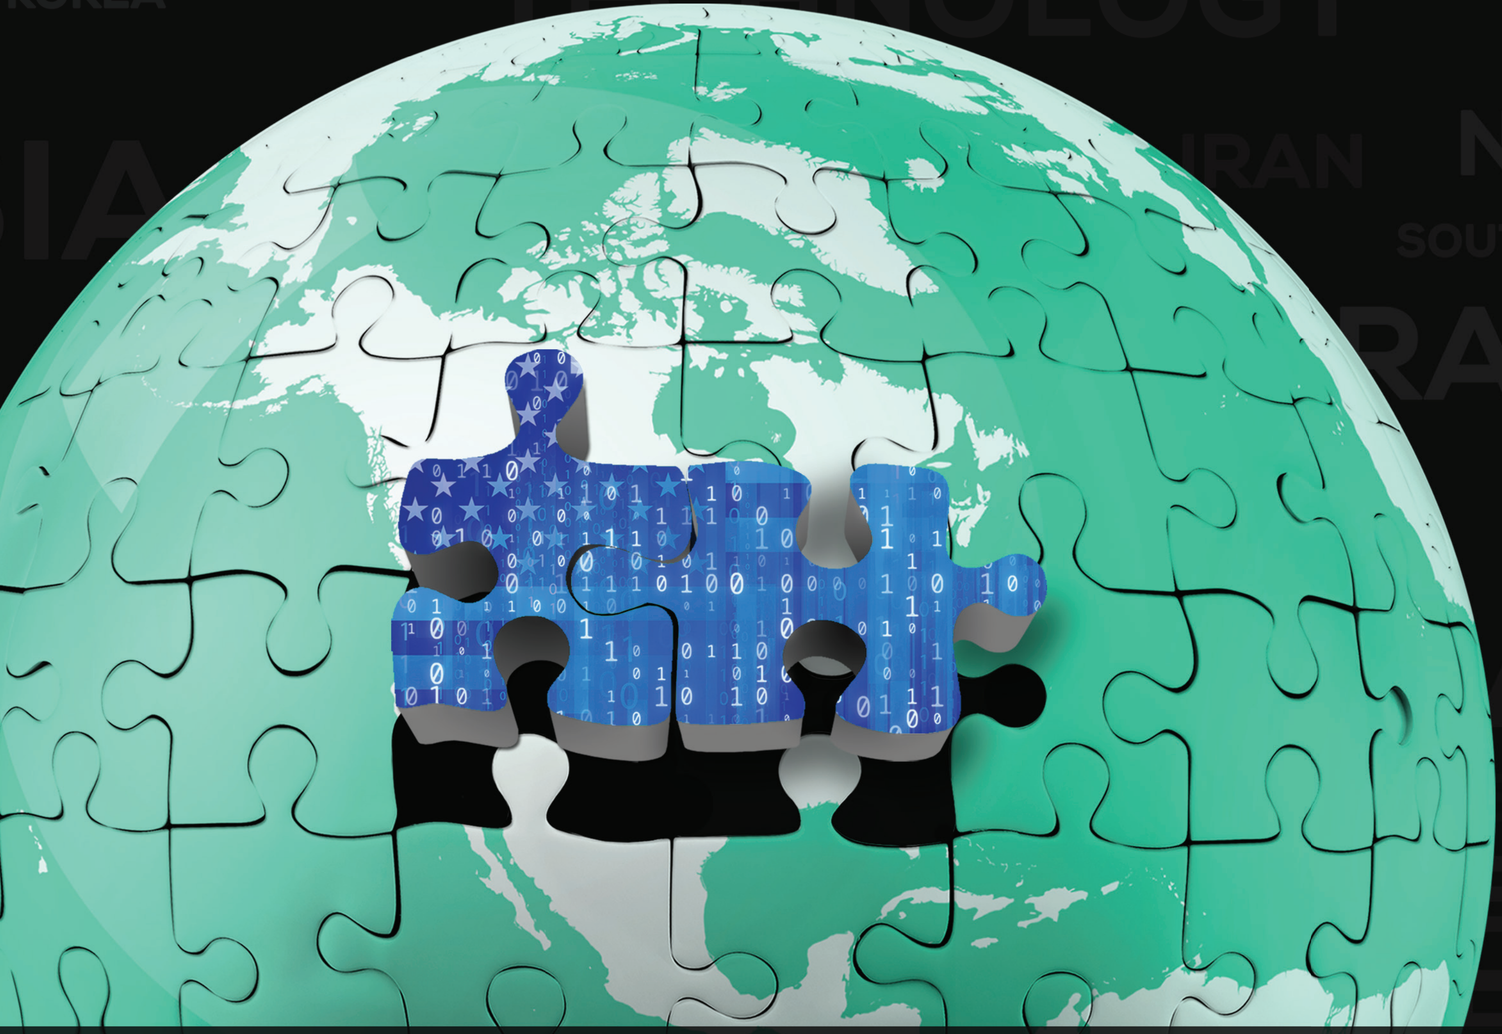

A PUBLICATION OF THE HOOVER INSTITUTION

HOOVER
INSTITUTION

GOVERNANCE IN AN EMERGING NEW WORLD

Convened by George P. Shultz
with James Cunningham, David Fedor, and James Timbie

Table of Contents

FALL SERIES, ISSUE 318

HOOVER INSTITUTION

A Letter from the Conveners

Sharp changes are afoot throughout the globe. Demographics are shifting, technology is advancing at unprecedented rates, and these changes are being felt everywhere.

How should we develop strategies to deal with this emerging new world? We can begin by understanding it.

First, there is the changing composition of the world population, which will have a profound impact on societies. Developed countries are experiencing falling fertility and increasing life expectancy. As working-age populations shrink and pensions and care costs for the elderly rise, it becomes harder for governments to afford other productive investments.

At the same time, high fertility rates in Africa and South Asia are causing both working-age and total populations to grow, but that growth outpaces economic performance. And alongside a changing climate, these parts of the world already face growing impacts from natural disasters, human and agricultural diseases, and other resource constraints.

Taken together, we are seeing a global movement of peoples, matching the transformative movement of goods and of capital in recent decades—and encouraging a populist turn in world politics.

Second is automation and artificial intelligence. In the last century, machines performed as instructed, and that "third industrial revolution" completely changed patterns of work, notably in manufacturing. But machines can now be designed to learn from experience, by trial and error. Technology will improve productivity, but workplace disruption will accelerate—felt not only by call center responders and truck drivers but also by accountants, by radiologists and lawyers, even by computer programmers.

All history displays this process of change. What is different today is the speed. In the early 20th century, American farm workers fell from half the population to less than five percent alongside the mechanization of agriculture. Our K-12 education systems helped to navigate this disruption by making sure the next generation could grow up capable of leaving the farm and becoming productive urban workers. With the speed of artificial intelligence, it's not just the children of displaced workers but the workers themselves who will need a fresh start.

Underlying the urgency of this task is the reality that there are now over 7 million "unfilled jobs" in America. Filling them and transitioning workers displaced by advancing technology to new jobs will test both education (particularly K-12, where the United States continues to fall behind) and flexibility of workers to pursue new occupations. Clearly, community colleges and similarly nimble institutions can help.

The third trend is fundamental change in the technological means of production, which allows goods to be produced near where they will be used and may unsettle the international order. More sophisticated use of robotics alongside human colleagues, plus additive manufacturing and unexpected changes in the distribution of energy supplies, have implications for our security and our economy as well as those of many other trade-oriented nations who may face a new and unexpected form of deglobalization.

This ability to produce customized goods in smaller quantities cheaply may, for example, lead to a gradual loss of cost-of-labor advantages. Today, 68 percent of Bangladeshi women work in sewing, and 4.5 million Vietnamese work in clothing production. Localized advanced manufacturing could block this traditional route to industrialization and economic development. Robots have been around for years, but robotics on a grand scale is just getting started: China today is the world's biggest buyer of robots but has only 68 per 10,000 workers; South Korea has 631.

These advances also diffuse military power. Ubiquitous sensors, inexpensive and autonomous drones, nanoexplosives, and cheaper access to space through microsatellites all empower smaller states and even individuals, closing the gap between incumbent powers like the United States and prospective challengers. The proliferation of low-cost, high-performance weaponry enabled by advances in navigation and additive manufacturing diminishes the once-paramount powers of conventional military assets like aircraft carriers and fighter jets. This is a new global challenge, and it threatens to undermine U.S. global military dominance, unless we can harness the new technologies to serve our own purposes. As we conduct ourselves throughout the world, we need to be cognizant that our words and deeds are not revealed to be backed by empty threats. At the same time, we face the challenge of proliferation of nuclear weapons.

Finally, the information and communications revolution is making governance everywhere more difficult. An analogue is the introduction of the printing press: as the price of that technology declined by 99 percent, the volume grew exponentially. But that process took ten times longer in the 15th, 16th, and 17th centuries than we see today. Information is everywhere—some accurate, some inaccurate, such that entire categories of news or intelligence appear less trustworthy. The "population" of Facebook now exceeds the population of the largest nation state. We have ceaseless and instantaneous communication to everybody, anybody, at any time. These tools can be used to enlighten, and they can also be used to distort, intimidate, divide, and oppress.

On the one hand, autocrats increasingly are empowered by this electronic revolution, enabled to manipulate technologies to solidify their rule in ways far beyond their fondest dreams in times past. Yet individuals can now reach others with similar concerns around the earth. People can easily discover what is going on, organize around it, and take collective action.

At present, many countries seek to govern over diversity by attempting to suppress it, which exacerbates the problem by reducing trust in institutions. Elsewhere we see governments unable to lead, trapped in short-term reactions to the vocal interests that most effectively capture democratic infrastructures. Both approaches are untenable. The problem of governing over diversity has taken on new dimensions.

The good news is that the United States is remarkably well-positioned to ride this wave of change if we are careful and deliberate about it. Meanwhile, other countries will face these common challenges in their own way, shaped by their own capabilities and vulnerabilities. Many of the world's strongest nations today—our allies and otherwise—will struggle more than we will. The more we can understand other countries' situations, the stronger our foundation for constructive international engagement.

This is why we have set off on this new project on Governance in an Emerging New World. Our friend Senator Sam Nunn has said that we've got to have a balance between optimism about what we can do with technology and realism about the dark side. So we aim to understand these changes and inform strategies that both address the challenges and take advantage of the opportunities afforded by these transformations.

To do so, we are convening a series of papers and meetings examining how these technological, demographic, and societal changes are affecting the United States (our democracy, our economy, and our national security) and countries and regions around the world, including Russia, China, Latin America, Africa, and Europe.

***

The two papers included in this volume take on the challenges the information and communications revolution pose to governance, particularly here in the United States. The rapid spread of information can enlighten, but it also can confuse, sow discord, and endanger democratic institutions. In the hands of autocrats, the new means of communicating become weapons of coercion, removing the information transparency that is so critical to democracy. And in the hands of individuals, these technologies enable never before seen forms of social and political organization, bringing new dimensions to the old problem of governing over diversity.

Recent events have made clear the current path is unsustainable. What can the United States do to protect its political process from the conflict and polarization catalyzed by social media and other network platforms? And what are some rules of the road for information warfare that might secure democratic institutions both at home and abroad? To begin addressing this challenge of information, communications, and governance, we have asked two eminent scholars to contribute their thoughts:

Hoover Institution senior fellow Niall Ferguson argues that network platforms, left unregulated, have damaged the democratic process in America. He proposes a multidomain approach to redress what he calls an indefensible status quo.

Joseph Nye of Harvard University writes that, although information warfare is not new, the communications revolution has changed its very nature, making it faster and cheaper than ever before. A national strategy to address that change and secure democracy against it must be both long-term in view and focus on resilience, deterrence, and diplomacy.

The authors came together this fall for a roundtable at the Hoover Institution to discuss their ideas, challenge each other's perspectives, and carry the conversation to the broader Stanford University and Silicon Valley community. We conclude this volume with our observations from that discussion, prepared along with Hoover research analysts David Fedor and James Cunningham, and we thank our colleagues at the Hoover Institution who have supported this project, particularly Shana Farley and Rachel Moltz for their work on this volume.

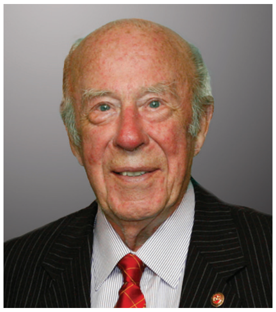

George P. Shultz
Thomas W. and Susan B. Ford Distinguished Fellow

James Timbie
Annenberg Distinguished Visiting Fellow

What Is to Be Done? Safeguarding Democratic Governance in the Age of Network Platforms

By Niall Ferguson, Hoover Institution

Introduction

Once upon a time, only the elite could network globally.[1] David Rockefeller—the grandson of the oil tycoon John D. Rockefeller—was a pioneer networker. According to a recent report, "He recorded contact information along with every meeting he had with about 100,000 people world-wide on white 3-by-5-inch index cards. He amassed about 200,000 of the cards, which filled a custom-built Rolodex machine, a 5-foot high electronic device." Rockefeller's contacts ranged from President John F. Kennedy to the shah of Iran, Pope John Paul II, and the astronaut Neil Armstrong. Henry Kissinger generated the most cards—35 in all, describing hundreds of encounters over sixty years. Not far behind was Gianni Agnelli, the Italian industrialist who ran Fiat. In his memoirs, Rockefeller recalled how, while serving as an Army intelligence officer during World War II, his "effectiveness [had] depended on my ability to develop a network of people with reliable information." It was an experience that he took home with him when he joined Chase Manhattan after the war.[2]

What was once the preserve of a tiny elite is now available to everybody with an Internet connection and a smartphone, tablet, or laptop computer. Facebook was founded at Harvard in 2004, before David Rockefeller's 90th birthday, and rose rapidly to become the world's dominant social media platform. The company's foundational premise was and remains that "Simply through sharing and connecting, the world gets smaller and better."[3] Connecting people on Facebook, Mark Zuckerberg declared in 2015, was building a "common global community" with a "shared understanding." A year later he summed up his utopian vision in a Facebook Live session (with Jerry Seinfeld):

> We're at this next point in human civilization, where we have the next set of tools that we need, things like the internet, that can be this global communication infrastructure ... Just like we went from hunter-gatherers to villages and cities and then nations, I think we now need to come together as a global community. Because a lot of the problems that you're talking about, whether it's terrorism or the refugee crisis or climate change or global diseases spreading around the world—these are not things that can be solved by any one city or one nation or one small group of people.

Humanity is certainly connected as never before. One in three people used a social network in 2017, according to eMarketer, nearly nine percent more than the previous year, thanks to rapid expansion in Asia-Pacific, Latin America, the Middle East, and Africa. Three quarters of all users of the Internet on mobile phones used their devices to access social media. In the United States, an estimated 60 percent of the population will use a social network at least once a month this year, up 2.6 percent since 2017. The number of U.S. Facebook users is expected to reach 169.5 million in 2018, more than 60 percent of all Internet users. Penetration is similarly high in the United Kingdom, where half the population is on Facebook.

How far the resulting network can be regarded as a problem-solving global community remains an open question, as we shall see. What is beyond doubt is that network platform companies are astonishingly profitable businesses—not least because users have handed them so much of their personal data for nothing, allowing advertisements to be targeted more precisely than ever before. As a commenter on the website MetaFilter memorably observed in 2010: "If you are not paying for it, you're not the customer; you're the product being sold." That was neat but not quite true. Users of network platforms enjoy access to numerous very useful services for which they pay nothing, aside from the distraction of on-screen advertisements—and the numerous negative externalities to be discussed below. It did not have to be this way, as network platforms might have opted to finance themselves through fees, subscriptions, or donations.[4] But the decisions were taken to monetize this way. In terms of its revenues, Facebook today is primarily a vast billboard, as is Google's parent company Alphabet. As a half-smirking Zuckerberg explained to Senator Orrin Hatch at a congressional hearing in 2018, "We sell ads, Senator." He meant ad space or, to be more precise, the potential interest of Facebook users in advertisements aimed at them on the basis of Facebook's data.

Eight of the world's most highly valued companies in 2017 were technology businesses, with a market capitalization equal to nearly a third of the market capitalization of the other 92 companies in the global top hundred. Of these eight companies, five (Apple, Alphabet, Microsoft, Amazon, and Facebook) were American, two Chinese (Alibaba) and Tencent), and one South Korean (Samsung). It was not strictly speaking software that "ate the world," in Marc Andreessen's famous phrase, because two of these companies sell hardware. It would be more accurate to say that network platforms ate the world, as the market dominance of all these companies arises from network effects and the operation of Zipf's Law (which can be summed up as "winner takes nearly all").

Amazon ate bookselling. Jeff Bezos's company today sells 55 percent of all the books sold in the United States, 82 percent of all the e-books, and 99 percent of all the audiobooks. Nor is that all. Although its share of total U.S. retail remains small (around 4 or 5 percent, or half of Walmart's share), the company's share of the global cloud business is 34 percent, its share of U.S. online commerce is 44 percent, and its share of the voice-activated device market is 71 percent.[5] Google ate search. It accounts for between 87 and 92 percent of online searches worldwide, processing 63,000 queries a second. Apple ate music (along with Alphabet's YouTube and Spotify). YouTube ate television. Above all, Google and Facebook ate advertising. According to eMarketer, the two companies will capture a combined 56.8 percent of U.S. digital ad spending in 2018, though Amazon and Snapchat have been increasing their shares. The revenues from this source seem certain to continue growing as a rising share of total advertising expenditure goes to the Internet. Together, all these companies also dealt heavy blows to all those businesses (for example, travel agents) that attracted customers through store window displays.

"For many years," according to a *New Yorker* profile, "[Mark] Zuckerberg ended Facebook meetings with the half-joking exhortation 'Domination!'" He stopped doing this because in European legal systems "dominance" is a term used to describe a business monopoly. However, he remains unabashed about Facebook's appetite for market share. "There's a natural zero-sumness," he told an interviewer in September 2018. Revealingly, the figure he most admires in history is the Emperor Augustus:

> You have all these good and bad and complex figures [in ancient Rome]. I think Augustus is one of the most fascinating. Basically, through a really harsh approach, he established two hundred years of world peace. What are the trade-offs in that? On the one hand, world peace is a long-term goal that people talk about today. Two hundred years feels unattainable. [But] that didn't come for free, and he had to do certain things.[6]

Facebook is indeed an empire, with as many users as Christianity has adherents, but a workforce of just 23,000. Nor does the new Caesar render up much to the old one: between 2007 and 2015, according to an estimate by S&P Global Market Intelligence, Facebook paid just 4 percent of its profits in federal, state, local, and foreign taxes. Amazon paid only 13 percent, Google 16 percent and Apple paid 17 percent. (The average S&P 500 company paid 27 percent.)[7]

Market dominance is seldom without its political aspect. Prior to 2016, remarkably little attention was paid to the growing political power of the big technology companies. The posture they adopted might best be summed up as *faux naïf*. "Don't be evil" was the motto adopted by Google in July 2001, after a brainstorming session between Eric Schmidt and early employees shortly before he took over as chief executive. As Schmidt recalled in 2006, referring to the company's decision to offer a censored version of its search services in China, "We actually did an 'evil scale' and decided [that] not to serve at all was worse evil."[8] By contrast, close involvement in the administration of Barack Obama—including energetic efforts to secure his reelection in 2012—was deemed to lie at the other end of the evil scale. According to one estimate, there were 252 job moves between Google and the Obama administration from its inception to early 2016, and 427 meetings between White House staff and Google employees from 2009 to 2015. In 2012, Google was the nation's second-largest corporate spender on lobbying, behind General Electric. By 2017 Google was number one, spending $18 million.[9] The closeness of the relationship between Silicon Valley and the Democratic Party was not unknown at the time, but it was remarkably uncontroversial. Little, if any, attention was paid to the political activities of the other big technology companies.

Donald Trump's victory in the November 2016 presidential election changed that, not just because the leaders of the big technology companies had confidently expected him to lose, but also because it was immediately apparent to them that their core products had either helped Trump win or failed to avert Clinton's defeat. The fact that Trump had dominated Clinton on both Facebook and Twitter throughout the campaign had been overlooked by most political pundits because his huge lead in follower numbers was at odds with all opinion polls. Likewise, experts ignored or discounted his even larger lead over her in terms of Google searches. Brad Parscale, Trump's digital media director, put it well: "These social platforms are all invented by very liberal people on the west and east coasts. And we figure out how to use it to push conservative values. I don't think they thought that would ever happen." A video recording of a post-election internal meeting at Google confirms this. Senior executives lined up to express their dismay at Trump's victory and their allegiance to Clinton.[10] As President Obama told

David Letterman, his own successful use of social media in 2008 had left him and other Democrats with "a pretty optimistic feeling about it. … What we missed was the degree to which people who are in power [*sic*], special interests, foreign governments, etcetera, can in fact manipulate that and propagandize." Obama's analysis of what had gone wrong for his party and its presidential candidate in 2016 is worth quoting in full:

> One of the biggest challenges we have to our democracy is the degree to which we don't share a common baseline of facts. … What the Russians exploited but it was already here is we are operating in completely different information universes. If you watch Fox News, you are living on a different planet than you are if you listen to NPR. … That's what's happening with these Facebook pages where more and more people are getting their news from. At a certain point you just live in a bubble. And that's part of why our politics is so polarized right now. I think it is a solvable problem but it's one we have to spend a lot of time thinking about.[11]

Speaking at MIT in February 2017, Obama suggested that "the large platforms—Google and Facebook being the most obvious, but Twitter and others as well that are part of that ecosystem—have to have a conversation about their business model that recognizes they are a public good as well as a commercial enterprise." It was, he said, "very difficult to figure out how democracy works over the long term" when "essentially we now have entirely different realities that are being created, with not just different opinions but now different facts—different sources, different people who are considered authoritative."[12]

Obama was right that we have a problem, and it is not a problem we were prepared for by the masters of Silicon Valley. "I thought once everybody could speak freely and exchange information and ideas, the world is automatically going to be a better place," Evan Williams, one of the founders of Twitter, told the *New York Times* in 2017. "I was wrong about that."[13] Indeed, he was. The impacts of the internet and the personal computer are akin to those of the printing press after it spread throughout Europe from the late 15th Century. The benefits of much cheaper, faster, and wider dissemination of ideas were offset by the costs of 130 years of religious conflict. Yet the network of printing presses has remained, to the present day, relatively distributed, with only limited concentrations of ownership (for example of newspapers or magazines). Today, by contrast:

- A handful of very large and very profitable corporations dominate the online public sphere in most countries in the world, raising questions about monopoly power at the international as well as the national level.
- Subject to the most minimal regulation in their country of origin—far less than the terrestrial television networks in their heyday—they tend not only to divide it but also to pollute national discourse with a torrent of fake news[14] and extreme views. The effects on the democratic process all over the world are potentially destabilizing.
- The negative effects on large numbers of individuals' privacy, safety and psychological health are also not trivial.
- Moreover, the vulnerability of the network platforms to outside manipulation—so called "information warfare"—poses a serious new challenge to national security.
- Yet attempts by the network platforms to clean up their act run the risk of restricting free speech, as their terms of service and growing army of content monitors seek to root out "hate speech."

Something needs to change. But what? What exactly does it imply to say that the network platforms are "a public good as well as a commercial enterprise"? Should the tech giants be broken up, as proponents of a revamped antitrust law argue? Should they be subject to tighter regulation, of the sort being pioneered by the European Union? Or should they be more exposed than they currently are to litigation by those harmed by the content they host? Finally, how can we mitigate the vulnerabilities unwittingly created by network platforms—in particular, the exposure of democracies to the disruptive tactics of information warfare? In short, to quote Nikolai Chernyshevsky (from whom Lenin stole the famous title), "What is to be done?"

The Self-Non-Regulation of the Internet

The starting point for any serious analysis must be the business model of the network platforms. They are, in Tim Wu's phase, "attention merchants," the heirs of the big 20th-century media companies, such as Hearst, which used news and other mostly non-fictional content to attract readers' attention, selling space alongside articles and photographs to advertisers.[15] They have a lot of attention to sell—more than any print publishing company in history. The average American spends 5.5 hours a day using digital media, more than television, radio and print put together.[16] More than half that digital media time is spent on mobile devices. An average smartphone user clicks, taps and swipes 2,617 times per day. In April 2016, Facebook said it was capturing on average 50 minutes of every American's day, up from 40 minutes in July 2014, though usage may have declined slightly in 2017.[17]

Most of what is said online is inane. The ten most common words used in Facebook status updates are day, loud, word, ticket, nice, long, light, hangover, good, and vote.[18] A little of what is said is highly cerebral. Some is old-fashioned fiction: stories that do not purport to be true. But a significant proportion is "news," i.e., content that purports to be true information about current affairs. In 2017, two thirds of American adults said they got news from social media sites. Around three quarters of Twitter users got news from the application, around two thirds of Facebook users, and around a third of YouTube users. In all, 45 percent of American adults get news from Facebook, 18 percent of them from YouTube, and 11 percent from Twitter. A significant share of younger users also gets news from Instagram and Snapchat.[19] It is a startling fact that Facebook and Google are still responsible for two thirds of news publishers' referral traffic, even after a deliberate effort by Facebook to reduce the importance of news in users' News Feeds.[20]

The network platforms set out to act as aggregators of news for the simple reason that it engages users' attention. If extracts from the novels of Dickens or Pascale's *Pensées* had the same appeal, these too would feature in users' feeds or search results. The key point is that the network platforms customize the news that users see in order to maximize their engagement. When Mark Zuckerberg talked in 2013 of making Facebook "the best personalized newspaper in the world," this was what he meant. News Feed is a "personalized collection of stories," and a user sees on average of 220 per day. Advertising and Pages, dedicated profiles for groups or causes, are sources of stories in News Feed. Anyone can buy ads to promote Pages, using an automated interface. Whenever one of Facebook's users opens the Facebook app, a personalization algorithm sorts through all the posts that a person could potentially see, serving up and sorting the fraction it thinks he or she would be most likely to share, comment on, or like. (Shares are worth more than comments, which are both worth more than likes.) Around two thousand pieces of user data ("features") are used by Facebook's machine-learning system to make those predictions. A somewhat similar process works when a user enters words in the Google search box. The user's individual search history, geographic location, and other demographic information affect the content and ranking of the results.

The problem is that the algorithms are not prioritizing truthfulness or accuracy but user engagement. For example, on October 1, 2017, Google directed users towards a false story alleging that the perpetrator of the Las Vegas massacre on that date was a member of the far-left group Antifa. A study by the *Wall Street Journal* and former YouTube engineer Guillaume Chaslot showed that a user searching for "The Pope" on YouTube was directed to videos with titles such as "How Dangerous is the Pope?", "What if the Pope was assassinated?" and "BREAKING: They caught the Pope."[21] As we shall see, these and other features of the network platforms had historic consequences in 2016 when they played a decisive role in the election of Donald Trump as the U.S. president.

How did we arrive at this state of affairs—when such important components of the public sphere could operate solely with regard to their own profitability as attention merchants? To answer this question, we must briefly review the history of Internet regulation prior to the crisis unleashed by the 2016 election. A decisive early decision was to define the Internet as a Title I information service, and therefore fundamentally different from the old telephone network, which was governed by Title II's intrusive monopoly utility regulations. (The Internet was briefly re-classified as a Title II service between 2015 and 2017, but no major regulatory change occurred in that period.) Another important decision was to give Internet companies very lenient treatment when they violated copyright. The Digital Millennium Act's notice-and-takedown provisions minimized the penalties to the network platforms of making the intellectual property of others available gratis to their users. A third vital decision—arguably the most important of all—was enshrined in Section 230 of Title V of the 1996 Telecommunications Act,[22] which was enacted after a New York court held online service provider Prodigy liable for a user's defamatory posts. Previously, managing content had triggered classification as a publisher—and hence civil liability—creating a perverse incentive not to manage content at all. Thus, Section 230c, "Protection for 'Good Samaritan' blocking and screening of offensive material," was written to encourage nascent firms to protect users and prevent illegal activity without incurring massive content management costs. It states:

1. No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.
2. No provider or user of an interactive computer service shall be held liable on account of:
    A. any action voluntarily taken in good faith to restrict access to or availability of material that the provider or user considers to be obscene [etc.]; or
    B. any action taken to enable or make available to information content providers or others the technical means to restrict access to material in paragraph 1.

In essence, Section 230 gives websites immunity from liability for what their users post—or, to be more precise, it "immuniz[e] platforms from liability both for underfiltering under Section 230(c)(1) and for "good faith" over-filtering under Section 230(c)(2)." The net result of this regulatory framework is that technology companies are neither communication utilities nor content publishers.

The argument for Section 230, as articulated by the Electronic Frontier Foundation, was that, "given the sheer size of user-generated websites … it would be infeasible for online intermediaries to prevent objectionable content from cropping up on their site. Rather than face potential liability for their users' actions, most would likely not host any user content at all or would need to protect themselves by being actively engaged in censoring what we say, what we see, and what we do online." Senator Ron Wyden put it even more strongly: "If websites, ISPs, text message services, video game companies and any other type of platform were held liable for every word and deed they facilitated or somehow enabled, the entire system would shut down … collaboration and communication on the internet would simply cease."[23] In effect, Section 230 split the difference between liability, which would have meant restriction, or complete lack of curation, which would have led to a torrent of "filth, racism, insults, and pornography." Thus, "hobbling 230" would "stifle the competition that got us to today's rich internet in the first place."[24] According to one recent and influential account:

> Platforms moderate content because of a foundation in American free speech norms, corporate responsibility, and the economic necessity of creating an environment that reflects the expectations of their users. Thus, platforms are motivated to moderate by both of §230's purposes: fostering Good Samaritan platforms and promoting free speech. … [They] should be thought of as operating as the New Governors of online speech. These New Governors are part of a new triadic model of speech that sits between the state and speakers-publishers. They are private, self-regulating entities that are economically and normatively motivated to reflect the democratic culture and free speech expectations of their users.[25]

Note that under the present dispensation, the network platforms have the power (not the obligation) to "curate" content that they host. They do so, it is argued, "out of a sense of corporate social responsibility, but also, more importantly, because their economic viability depends on meeting users' speech and community norms." This curation began some time ago with the exclusion of content that no one would publicly condone. For years, the big technology companies have filtered out child pornography using an automated hash database assembled by the National Center for Missing and Exploited Children, so that, as soon as an illegal photo or a video is uploaded to one site, it is detected and excluded from all platforms. Facebook, Twitter, YouTube, and Microsoft have a global working group that applies somewhat similar technology to find and filter out terrorist content. In November 2017, for example, YouTube took down videos of Anwar al-Awlaki, the jihadist cleric killed by a U.S. drone strike in Yemen in 2011. Video "fingerprinting" has not removed al-Awlaki altogether from the platform, but it has substantially reduced the number of videos relating to him.

However, the process of removing or at least downgrading offensive content has not stopped with recognized advocates of pedophilia or jihad. In January 2018, for example, YouTube removed from its Google Preferred platform the channels of Logan Paul, a YouTube star with almost 16 million subscribers, after he posted a video showing a suicide victim in Japan. "Demonetizing" YouTube videos, so that they are not promoted on the platform and their creators receive no share of any advertising revenue, is a powerful sanction short of outright prohibition. Twitter set out to be the "free speech wing of the speech movement," but in 2015 added a new line to its Twitter Rules that barred "promot[ing] violence against others … on the basis of race, ethnicity, national origin, religion, sexual orientation, gender, gender identity, age, or disability." Any concern that the "new governors" might abuse their power of moderation was dismissed with a promise that all problems could be addressed by making "changes to the architecture and governance systems put in place by these platforms," with regulation as a last resort.[26] Yet platforms' content moderation policies are not public, only their terms of service and usually vague community standards. Under the current, sweeping interpretation of Section 230, the network platforms can rely on judges to dismiss most litigation whether they under-filter or over-filter.[27]

The network platforms have been left blissfully unmolested by the Federal Communications Commission, even as they have moved into direct competition with the television radio stations that it regulates and the telecommunications firms that provide most consumers with access to the Internet. The Federal Election Commission until recently considered digital platforms exempt from disclosure rules on political ads. In 2011 Facebook asked the FEC for an exemption to rules requiring the source of funding for political ads to be disclosed, arguing that the agency "should not stand in the way of innovation."[28] The only U.S. regulator that Silicon Valley has had to contend with is the Federal Trade Commission, which has powers to enforce consumer-protection laws. In 2011 the FTC cited Facebook for "engaging in unfair and deceptive practices" with regard to the privacy of user data after it became clear that the company had changed users' privacy settings in 2009 and shared users' locations and religious and political leanings with Microsoft and others. The company signed a consent decree pledging to establish a "comprehensive privacy program" and to evaluate it every other year for twenty years—a commitment it seems to have honored mainly in the breach, as we shall see, in its reckless pursuit of new users through partnerships with other big tech companies including Microsoft, Netflix and Amazon.[29]

In short, to characterize the U.S. regulation of network platforms as *laissez faire* or "light touch" would be a considerable understatement. Until very recently, the system has been, in essence, one of self-regulation—or, to be more precise, self-non-regulation.

The Political Consequences

The consequences have been profound. The first and best known has been to favor fake news. It has been claimed that social media mainly tend to amplify the content produced by traditional media.[30] But that overlooks two things. First, if network platforms choose to promote a message of their own—for example that users should consider organ donation for transplants—the results are far more impressive than when newspapers or television channels make similar appeals.[31] Second, social media also disseminate fake news, which traditional media tend not to do (with the exception of sensationalist publications such as the *National Enquirer*, which are generally understood to be factually unreliable). Unfortunately, it appears that false information on Twitter is typically retweeted by many more people, and far more rapidly, than true information, especially when the topic is politics. Researchers at the Massachusetts Institute of Technology tracked 126,000 stories—some true, some false—tweeted by roughly three million people more than 4.5 million times from 2006 through 2017. They then used six different fact-checking sites—including Snopes, Politifact, and FactCheck.org—to rate the truthfulness of each story. It turned out that false claims were 70 percent more likely than the truth to be shared on Twitter. True stories were rarely retweeted by more than a thousand people, but the top one percent of false stories were routinely shared by between 1,000 and 100,000 people. And it took true stories about six times as long as false ones to reach 1,500 people. No accurate news item was able to chain together more than ten retweets, whereas fake news could put together a retweet chain 19 links long—and do it ten times faster than the accurate news item put together its ten retweets. This finding is especially startling, as Twitter users who share accurate information typically have more followers, and send more tweets, than fake-news sharers.[32] An important role is evidently played in the dissemination of fake news by "bots"—automated accounts purporting to be humans—who are believed to account for between 9 and 15 percent of active Twitter accounts and as many as 60 million Facebook users.[33] In the words of a recent large-scale study by the Knight Foundation, "A supercluster of densely interlinked, heavily followed accounts plays a large role in the spread of fake news and disinformation on Twitter. Social bots likely make up the majority of the accounts in the supercluster, and accounts in the cluster participate in what appear to be coordinated campaigns to push fake news stories."[34]

A second consequence of leaving the network platforms to their own devices has been polarization. Homophily—the tendency of birds of a feather to flock together—has long been a recognized feature of social networks, even those of modest size. Giant online networks were therefore always likely to self-segregate into more or less homogeneous clusters. This was true of the "blogosphere," for example, prior to the ascendancy of the network platforms.[35] We are innately inclined to form into opposing sides over any bone of contention, as the case of "The Dress"—a photograph of a dress taken in England at 3:30pm on a February afternoon—illustrates. (For days the Internet was rent asunder: Was the dress black and blue or white and yellow?)[36] Yet the network platforms do more than merely reveal our innate divisions. Because their algorithms are designed to maximize our engagement and because confirmation bias is one of our many cognitive frailties, they tend to accentuate it.

A good illustration of the point is the way that Twitter works. If the network of political retweets is graphed using a force-directed algorithm, two highly segregated communities of users are revealed: liberals and conservatives.[37] A similar depiction of the retweet activity of messages containing moral and emotional language on a range of political topics (gun control, same-sex marriage, climate change) looks very similar. The presence in tweet of words that researchers classified as "moral-emotional" increased its diffusion by a factor of 20 percent for each such word. Moreover, "moral contagion" was bounded by group membership, in that "moral-emotional language increased diffusion more strongly within liberal and conservative networks, and less between them."[38] Another paper, based on 3,938 Twitter users who together generated 4.8 million tweets in August 2016, suggested that it is politically extreme people who tweet about politics more than centrist users of Twitter.[39] In Congress, more ideologically extreme lawmakers get more Facebook followers than centrists.[40]

This is not to say that people consciously choose to inhabit filter bubbles or echo chambers. A 2015 study found that "most social media users [were] embedded in ideologically diverse networks, and that exposure to political diversity has a positive effect on political moderation." Nor is it to blame all polarization on social media. The division between liberals and conservatives has deep historical roots. Its reflection in the media—from newspapers to cable television—is not new. And the increase in polarization in recent years seems to have been greatest amongst the elderly, the group least likely to use the internet and social media.[41] Yet the network platforms are clearly making polarization worse because of the way they work. A good example is the way the YouTube suggestions algorithm works. In the words of former Google engineer Guillaume Chaslot, "Videos about vegetarianism led to videos about veganism.

Videos about jogging led to videos about running ultramarathons. Given its billion or so users, YouTube may be one of the most powerful radicalizing instruments of the 21st Century. ... YouTube leads viewers down a rabbit hole of extremism, while Google racks up the ad sales."[42] YouTube assists extremists in other ways, notably by placing advertisements for respectable corporations and other organizations on the websites of Nazis and so-called white nationalists. YouTube channels with over 1,000 subscribers and 4,000 watch-hours over a twelve-month period can apply to receive a portion of YouTube's revenue from the advertisements running on their videos. Until it was shut down, Brian Ruhe's Nazi channel featured ads—and therefore earned money—from Nissan, Disney, Mozilla and 20th Century Fox.[43] There is also evidence from the 2015 protests against police violence in Baltimore that the degree and volume of moral rhetoric used on social media predicted the number of arrests during protests.[44] The "culture war" in the United States—over campus free speech, climate change, healthcare reform, racism, gender fluidity, and sexual harassment—would no doubt be happening without social media. But the network platforms have surely intensified the conflict.

The elections of 2016 brought to light the full extent of the network platforms' role in the modern public sphere. Like the culture war, the elections would no doubt have been bitterly contested without social media. But the network platforms in 2016 did more than merely intensify the contests. They decisively influenced the results.

There is of course nothing new about a change in the structure of the public sphere having political consequences. Each new communications technology—the newspaper, the telegraph, the radio, the television—has had its effect on the political process. William Randolph Hearst was first a hero and then a pariah for his accumulation of mass-circulation newspapers and magazines, each of which toed his political line. Most recently, cable news stations such as CNN and Fox have had a measurable impact on American elections. Biased coverage by Fox News during the 2000 presidential election is said to have helped George W. Bush win an additional 11,000 votes in the crucial state of Florida, where his margin of victory was just 537 votes.[45] Social media played a modest role in the 2008 election because at that time Facebook and Twitter were still in their infancy, though there is no question that Barack Obama's use of them was far superior to John McCain's. But they were more important in the 2012 election. This shift had its roots in 2009, when Dan Wagner began work on a sophisticated model (known as the Survey Manager), which before long was accurately forecasting the "shellacking" inflicted on the Democrats by Tea Party Republicans in the 2010 midterms. At first, much of the data Wagner used were taken from surveys that he and his team carried out. But by 2012 the Democratic National Committee had invested heavily in technology designed to integrate as much information as possible about individual voters. The Obama campaign spent twice as much on online advertisements as Mitt Romney's. Its approach to television advertising was based on far superior data on the viewing habits of potentially persuadable voters. The Romney campaign mostly out-sourced its data operation; in effect, it relied on the technology and methods Obama had used in 2008.[46] In 2012, by contrast, the Obama campaign was able to create a voter-outreach app that analyzed users' Facebook connections and encouraged users to reach out to potential Obama supporters among their Facebook friends. Google also made its data tools (Google Analytics) available to Obama's reelection campaign.

Google is first and foremost a search engine. So well established is the company's reputation for fast and efficient ranking of web pages that most users assume more or less uncritically that its search results are an objective measure of citation frequency, if not credibility. That is why nine out of ten clicks are on the first page of Google search results. In fact, the search engine can be manipulated and with significant effects. Three experiments relating to the 2010 Australian election found that the influence of a Google-like search engine on voter behavior was very great. After subjects were left to conduct their own search-based web research, seeing results that were biased by design, the likelihood of their voting for a particular candidate diverged from their initial preferences by between two and four percentage points. The authors term this differential "voter manipulation power" or VMP. Between three quarters and all of the subjects in these experiments showed no awareness of the manipulation. A similar experiment with U.S. voters had similar results. The VMP of voters in the 2014 Lok Sabha elections in India was 9.4 percent and as much as 73 percent amongst unemployed males from Kerala. The authors of a pioneering 2015 paper on the "search engine manipulation effect" concluded: "If a search engine company optimized rankings continuously and sent customized rankings only to vulnerable undecided voters, there is no telling how high the VMP could be pushed. ... Search results favoring one candidate could easily shift the opinions and voting preferences of real voters in real elections by up to 80 percent in some demographic groups with virtually no one knowing they had been manipulated."[47] The same author's estimated that "Google's search engine—with or without any deliberate planning by Google employees—was currently determining the outcomes of upwards of 25 percent of the world's national elections … because Google's search engine lacks any kind of equal-time rule, so it virtually always favors one candidate over another."[48] Comparably powerful is the "Search Suggestion Effect," whereby Google search suggestions ("autocomplete" suggestions) have the power to shift opinions and voting preferences:

> Negative ("low valence") search terms can attract 10-to-15 times as many clicks as neutral or positive terms can (an example of "negativity bias"), which means that a simple yet powerful way for a search engine company to manipulate elections is to suppress negative search suggestions for the candidate it supports, while allowing one or more negative search suggestions to appear for the opposing candidate (the "differential suppression of negative search suggestions") … [T]he higher a suggestion appears in a list of search suggestions, the more impact it has on search, and overall, manipulating search suggestions can shift a 50/50 split among people who are undecided on an issue to a 90/10 split without people's awareness and without leaving a paper trail for authorities to follow.[49]

Nearly everything that happened in 2016 had a pre-history. Cambridge Analytica was established in June 2014 by Robert and Rebekah Mercer, who had previously invested in Alexander Nix's Strategic Communications Laboratories (SCL) Group. The new company was set up as a joint venture with SCL; its name was chosen by Stephen K. Bannon, executive chairman of the right-wing website Breitbart, who became a director. The company's voter database was purchased from the Cambridge University researcher Aleksandr Kogan, whose "this is your digital life" app represented itself as a research tool used by academic psychologists. Its online questionnaire was hosted by a company called Qualtrics. Respondents were asked to authorize access to their Facebook profiles and, when they did, Kogan's app harvested their data as well as the data of all their Facebook friends—not only their names, birth dates and location data, but also lists of every Facebook page they had ever liked. The small print accompanying Kogan's questionnaire told users that their data could be used for commercial purposes, a violation of Facebook's rules at the time. In 2015, Facebook said that, having learned that Kogan had passed user data to Cambridge Analytica, it had excluded him from access to Facebook data and demanded assurances that the data had been deleted. But the company did not disclose to users—or anyone else—that their data had been misused. Moreover, in November 2015, Facebook brought Kogan in as a consultant to explain the technique he had used.[50] In March 2018, when Facebook was forced by news revelations to comment, the company insisted there had not been a data breach but merely "a scam—and a fraud." In all, some 270,000 people downloaded Kogan's app. That was enough to give him access to the data of 87 million Facebook users, 71 million of whom were Americans. Around 30 million profiles contained enough information, including places of residence, to enable Cambridge Analytica to build the psychographic profiles that were its main selling point.[51] In 2014 the company did work for Texas Senator Ted Cruz, the John Bolton Super PAC, conservative groups in Colorado and the campaign of Senator Thom Tillis, the North Carolina Republican.[52]

To date, the most controversial aspect of social media's involvement in the 2016 U.S. election was the way Russian entities, principally the Internet Research Agency (IRA), disseminated inflammatory content through Twitter and Facebook. This, too, had a prelude. According to special counsel Robert Mueller's indictment of 25 Russian individuals (12 of them identified as Russian military intelligence agents), Moscow began building a U.S. influence operation in 2014. It had already experimented with "information warfare" in a number of countries, including Ukraine, which Russian forces invaded in 2014. During the UK referendum on membership of the European Union, at least 2,752 Twitter profiles appear to have been created and managed by the IRA. Russian content found itself into mainstream media on numerous occasions: 29 different Russian-run accounts were quoted across 73 different news stories. The *Daily Telegraph* embedded posts from Russian accounts such as @Ten_GOP and @Pamela_Moore13 15 times, for example. Another Russian account, @WarfareWW, accounted for four of the *Daily Mail*'s seven citations. The extent of the cooperation between the Leave campaign, its financial backers, Cambridge Analytica, and the Russians remains controversial.[53] What is clear is that there was a "massive volume" of tweets from Russian-language accounts in the few days before the referendum. According to a recent NBER study, automated English-language tweets added as much as 1.76 percentage point to the "Leave" share of the vote.[54] When, in October 2018, Twitter released datasets of more than 10 million tweets and more than 2 million images and videos from 3,841 IRA-affiliated (plus 770 others suspected of being Iranian), the material dated back as far as 2009.[55]

Yet Russian disinformation needs to be seen as part of a wider problem that had become apparent before 2016. Already in 2013 Kate Starbird identified the use of fake Twitter accounts and bots to disseminate fake news and conspiracy theories in the wake of the Boston Marathon bombings and the Umpqua Community College shooting. Although the Russians had long been interested in information warfare, there was only minimal Russian involvement at this stage. Sites such as VeteransToday.com, BeforeItsNews.com, and NoDisinfo.com were created by home-grown purveyors of fake news, exploiting the vulnerabilities of the unregulated Internet. Those responsible had also understood, without Russian assistance, that mainstream media sites could in various ways be coopted into the dissemination of fake news.[56]

In short, there was ample reason to expect the network platforms to play both an important and a subversive role in the U.S. election of 2016. And yet almost no precautions appear to have been taken by either the authorities or the big technology companies themselves. On the contrary, Facebook executives continued to repeat the company's foundational mantra: "Move fast and break things." On June 18, 2016, a day after the fatal shooting of a Chicago man was captured on Facebook Live, the company's vice president Andrew "Boz" Bosworth, circulated an internal memorandum that epitomized a culture of indifference to negative externalities:

> So we connect more people.
>
> That can be bad if they make it negative. Maybe it costs a life by exposing someone to bullies. Maybe someone dies in a terrorist attack coordinated on our tools.
>
> And still we connect people.
>
> The ugly truth is that we believe in connecting people so deeply that anything that allows us to connect more people more often is *de facto* good. It is perhaps the only area where the metrics do tell the true story as far as we are concerned.
>
> That isn't something we are doing for ourselves. Or for our stock price (ha!). It is literally just what we do. We connect people. Period. …
>
> … I know a lot of people don't want to hear this. Most of us have the luxury of working in the warm glow of building products consumers love. But make no mistake, growth tactics are how we got here. …
>
> That's our imperative. Because that's what we do. We connect people.

What Exactly Happened in 2016?

Since the November 2016 election, Facebook has admitted that the Internet Research Agency used up to 470 false identities and spent $100,000 on more than 3,000 Facebook and Instagram ads to spread politically divisive posts to Americans before and after the vote. The ads promoted about 80 Facebook Pages set up by the IRA and related groups. These Pages posted more than 60,000 pieces of content between January 2015 and August 2017, reaching up to 126 million Americans—a number only slightly smaller than the total number who voted. In all, an estimated 146-150 million users saw posts from accounts linked to the IRA, including at least 20 million Instagram users. The most thorough study to date shows that the Russians also disseminated content through Twitter (more than 10 million tweets across 3,841 accounts) YouTube (to which they uploaded around 1,100 videos using 17 channels), G+, Gmail, and Google Voice, as well as Vine, Gab, Meetup, VKontakte, and LiveJournal, not to mention Reddit, Tumblr, and Pinterest—even Pokémon Go.[57] Russia content covered numerous themes, not all of them explicitly political, but all calculated to exacerbate social—and especially racial—divisions. In addition, Russians with false identities used Facebook Events to promote political protests, including an August 27 anti-immigrant, anti-Muslim rally in a rural Idaho town that was known to welcome refugees. The event was "hosted" by "SecuredBorders," a phony anti-immigration group that was in fact a Russian front. (It had 133,000 followers when Facebook closed it down.) "Heart of Texas," a Russian-controlled Facebook group that promoted Texas secession, announced a rally to "Stop Islamification of Texas" in front of the Islamic Da'wah Center of Houston on May 21, 2016. A separate Russian-sponsored group, "United Muslims of America," advertised a "Save Islamic Knowledge" rally for exactly the same place and time. The Russians did not confine themselves to Facebook. Twitter has admitted that Russian bots tweeted 2.1 million times before the election. So far, 2,700 Twitter accounts have been identified as IRA-run. Of these, 65—notably @ WarfareWW, @TEN_GOP and @Jenn_Abrams—featured prominently in a large-scale study of the fake news phenomenon.[58]

How big an impact did this Russian effort have? Political scientist Kathleen Hall Jamieson goes so far as to argue that the election was in fact decided by Russian content aimed at discouraging potential Clinton voters from turning out, as this was the key variable that tipped Michigan, Pennsylvania, and Wisconsin over to Trump.[59] It is important, however, to recognize that the impact of Russian was not simply in one, pro-Trump direction. A study of retweet activity reveals, surprisingly, that liberals were more inclined to disseminate this content than conservatives (presumably because it confirmed their worst fears about conservatives).[60] Another analysis of the Twitter network suggests that "clusters of accounts affiliated with Russia serve a brokerage role, serving as a cultural and political bridge between liberal U.S. accounts and European far-right accounts."[61] Moreover, the Russians only accounted for a part of the fake news disseminated by foreign actors during the election. Buzzfeed traced a hundred pro-Trump sites to a small town in Macedonia. At least one of these was run by the Israeli private intelligence firm Psy-Group. The Iranians also appear to have abused Twitter.[62]

Yet the key point is that all this foreign-originated content was a drop in the ocean. Between March 23, 2015, and November 2016, an estimated 128 million people in America were responsible for nearly 10 *billion* Facebook posts, shares, likes, and comments about the election. Hindman and Barash identified 6.6 million tweets or retweets from 454,832 separate accounts that linked to at least one of more than 600 fake news or conspiracy websites during the month before the election. The tweets of Russian operatives and bots represented perhaps one percent of all electionrelated tweets. The same goes for fake news from any source was only a fraction of total election-related content. True, according to one estimate, the average American encountered between one and three stories from known publishers of fake news during the month before the 2016 election.[63] Another study of the browsing histories of 2,525 adult Americans during the run-up to the 2016 election suggested that one in four Americans saw at least one false story, with the most conservative 10 percent of the sample accounting for two thirds of visits to fake news sites. Yet false stories were a very small proportion of the total news people consumed, accounting for just 1 percent of the news intake of Clinton supporters and 6 percent in the case of Trump supporters.[64]

It could be argued that even that small a proportion of fake news—foreign and home-grown—might have sufficed to decide a very tight election. Research on the impact of traditional television advertising between 2004 and 2012 concluded that partisan imbalances in advertising had significant effects. (Exposure to an additional ad by one party or the other shifted the partisan vote of approximately two people out of 10,000.)[65] Moreover, false election stories tended to attract more attention than true ones. In August 2015, a rumor circulated on social media that Donald Trump had let a sick child use his plane to get urgent medical care. Snopes confirmed almost all of the story as true. However, only about 1,300 people shared or retweeted the story. In February 2016, by contrast, a rumor circulated that Trump's elderly cousin had recently died and that he had left a message publicly condemning his relative's presidential bid. Snopes rejected this story as false, but around 38,000 Twitter users shared it. Its retweet chain was three times longer than the one produced by the true sick child story. A false story alleging that the boxer Floyd Mayweather had worn a Muslim head scarf to a Trump rally also reached an audience more than ten times the size of the sick child story.

However, the Russians had real news, too, which they had obtained by means deemed legitimate by the mainstream media outlets that published it. Beginning in March 2016, Russian hackers Fancy Bear sought to phish their way into the Democratic party's emails, sending fake emails that seemed to come from Google and told recipients to change their passwords. Chairman John Podesta fell for the ruse on March 19, clicking on the fatal link and giving the Russians access to 50,000 emails. Former Trump foreign policy adviser George Papadopoulos said that he was told on April 26 by an academic closely connected to the Kremlin that the Russian government had obtained compromising information about Clinton. "They have dirt on her," Papadopoulos was told. "They have thousands of emails." On June 12 WikiLeaks founder Julian Assange told a British TV show that emails related to Clinton were "pending publication." On June 15 the supposedly Romanian (but in reality Russian) Guccifer 2.0 claimed to have hacked the DNC and began directing reporters to the newly launched DCLeaks site. WikiLeaks joined in on July 22. Guccifer 2.0, WikiLeaks and DCLeaks ultimately published more than 150,000 emails stolen from more than a dozen Democrats. Long-established newspapers did as much as network platforms to disseminate the leaked content. Hillary Clinton herself attaches more blame for her defeat to FBI James Comey's intervention on October 28—when he told Congress that he was reopening an investigation into her private email server—than to Russian disinformation. Still, it would be surprising if Robert Mueller's inquiry does not reveal at least some coordination between the Trump campaign, the Russian government and Wikileaks. The irony is that mainstream media did as much as social media to disseminate the more damaging content of the Democrats' emails.

Another point often overlooked is that not all online manipulation during the campaign was calculated to hurt Clinton. True, YouTube tilted its users towards Trump. Guillaume Chaslot discovered that, regardless of whether users started with a pro-Clinton or a pro-Trump video, they were many times more likely to have a pro-Trump video recommended by YouTube. Twitter also helped Trump as most Twitter bots (not only the Russian ones) favored him.[66] However, there is striking evidence that the Search Engine Suggestion Effect was exploited by Google in Clinton's favor. Search suggestions on Google, Yahoo and Bing on August 3, 2018, differed strikingly. The user who consulted Yahoo was promoted to search for "Hillary Clinton is a liar" or "Hillary Clinton is a criminal." On Bing the top two suggestions were "Hillary Clinton is a filthy liar" and "Hillary Clinton is a murderess." But the user who typed "Hillary Clinton is" into Google's search box was prompted to search for "Hillary Clinton is winning" or "Hillary Clinton is awesome."[67] Robert Epstein argues that Google search rankings favored Clinton over most of the six-month period that he and his research associates monitored. Between October 15 and Election Day, "search rankings favored Mrs. Clinton in all 10 of the search positions on the first page of search results."[68]

Yet the truly decisive factor in the 2016 was probably none of the above. It was the different ways the Trump and Clinton campaigns themselves used the network platforms, particularly Facebook. According to the *New Yorker*, Facebook "offered to 'embed' employees, for free, in presidential campaign offices to help them use the platform effectively. Clinton's campaign said no. Trump's said yes." Trump used Facebook to raise $280 million. More than a third of that money—vastly more than the Russians spent—went to pay for targeted Facebook ads. These included a voter-suppression drive in the days before the election, targeting "idealistic white liberals, young women, and African Americans." Theresa Hong, the Trump campaign's digital-content director, later told an interviewer, "Without Facebook we wouldn't have won."[69] Her verdict has been echoed by Gary Coby, the director of advertising at the Republican National Committee and director of digital advertising and fundraising for Trump's campaign, and by Brad Parscale in an interview with Lesley Stahl for *60 Minutes* and in a conversation with Michael Isikoff in Lisbon in November 2017. Ali-Jae Henke, the head of elections at Google, explained how his company also provided help to both campaigns:

> They say like "look, we really want to get attention and we want to reach as many people as possible and these are kind of the areas politically where we might have challenges or the different types of voting blocs we need to reach" … and so then I am able to in that advisory capacity be like, "well this is what moms look like online, this is how we find them ..."

Parscale explained the role of Cambridge Analytica in this process:

> [They] didn't play a role in crafting ads [but] helped with a research strategy to help us raise money. We needed to build an infrastructure. [Cambridge Analytica] provided staff, resources, because we had to grow a large organization, fast. They did a lot of polling, and they did a lot of building some directional arrows for us [regarding] where to place the money, being able to provide reporting back that says, "Here are trends that are happening," so I could move the budget around in a way and I could make recommendations to [then candidate Trump] and to leadership, saying, "Here's an opportunity. We should go into this part of Michigan. We should go into this part of Wisconsin." [Cambridge was] able to drive that kind of information … and in a simple consumption model, daily.[70]

Also important was the way Facebook directed users to content on the Breitbart website, run since March 2012 by Steve Bannon. On August 17, 2016, Bannon was appointed chief executive of Trump's presidential campaign. "I wouldn't have come aboard, even for Trump," he later said, "if I hadn't known they were building this massive Facebook and data engine. Facebook is what propelled Breitbart to a massive audience. We know its power."

This was the crucial difference between 2012 and 2016. By the time of the later election, Facebook had acquired a database of American voters far superior to anything either party could possibly have built on its own. It had done so not only by persuading a majority of Americans to join Facebook, but also by (among other things) logging the phone call and messaging histories of Android smartphone users who installed Messenger or Facebook Lite and then synced their phone contacts with the app.[71] Facebook had also paid or otherwise persuaded third-party websites and apps to let it place cookies, invisible pixels, "like" and "share" buttons on them, thereby acquiring data on people who were not Facebook users.[72] Facebook's approach to data-gathering was even more ruthless than Google's.[73] Its attitude towards how that data got used by third parties was, at best, cavalier. And only one campaign made full use of Facebook's data.

The evidence that Facebook played a decisive role in the election is compelling. According to a pioneering European study published in November 2018, Facebook "had a significant effect in persuading undecided voters to support Trump and in persuading Republican supporters to turn out on election day, but had no effect on Clinton's side. … Exposure to political ads on Facebook increased the likelihood of voting by between 5% and 10%. … Targeted Facebook campaigning increased the probability that a previously nonaligned voter would vote for Trump; … if the voter used Facebook regularly, the probability increased by at least 5%."[74]

We may ask counterfactual questions to our heart's content about how the election would have turned out if the Russians had played no role. Yet the crucial point is that Russian meddling was a subplot in a much bigger crisis of the American political system produced by the unregulated and generally reckless operation of the network platforms. Without the Russians, Trump might not have won. It would still have been close. Without Facebook, he would have stood no chance.

The Backlash

That there would be a backlash against the network platforms after their role in the 2016 election was easily predictable.[75] A few writers—such as Jonathan Zittrain and Tim Wu—had been warning about their growing power for some time. Established content publishers such as Rupert Murdoch and Michael Bloomberg had obvious commercial reasons for going on the offensive. For Wu, Facebook was like a television network but with no "sense of responsibility. No constraints. No regulation. No oversight. Nothing." Robert Thomson of News Corp talked about "tech tapeworms in the intestines of the internet." Speaking at Davos in January 2017, George Soros warned of "a web of totalitarian control the likes of which not even Aldous Huxley or George Orwell could have imagined," and called on European Union Competition Commissioner Margrethe Vestager to be the "nemesis" of the network platforms.[76]

The most credible critics have been the insiders—former Facebook employees such as Antonio Garcia Martinez, author of *Chaos Monkeys*, or Sandy Parakilas, a former operations manager, who publicly criticized the company's handling of privacy issues in 2017, warning that "The company won't protect us by itself, and nothing less than our democracy is at stake."[77] A former vice president for user growth, Chamath Palihapitiya, told an audience at Stanford's Graduate School of Business: "I think we have created tools that are ripping apart the social fabric of how society works. … The short-term, dopamine-driven feedback loops that we have created are destroying how society works. No civil discourse, no co-operation: misinformation, mistrust." He felt "tremendous guilt" about his own part in this because, deep down, he and his former colleagues "kind of knew something bad would happen." In a similar vein, Facebook's first president Sean Parker admitted that the platform was consciously designed to take advantage of "a vulnerability in human psychology" by delivering "a little dopamine hit every once in a while." Parakilas joined forces with Dave Morin, Justin Rosenstein, and Roger McNameee—all former Facebook employees or early investors—as well as Tristan Harris and Lynn Fox (both ex-Google) to establish the Center for Human Technology and launch a campaign with the title: "The Truth about Tech."

Negative publicity in the wake of the election has led to a decline in public trust in Facebook and Twitter.[78] Facebook also now lags quite far behind Amazon, Google, Apple, and Microsoft in the Small Business Trust Index. A Gallup Knight survey published in January 2018 showed that 57 percent of Americans regarded the way tech companies chose which stories to show to users as "a major problem" for democracy, while 73 percent said the same about the spread of inaccurate information on the Internet.[79] At that time only 49 percent favored regulation of how websites provide news. By February 2018, however, the proportion "concerned that government would do too little" to address the problem had risen from 40 to 55 percent.[80] Young Americans in particular have lost trust in Facebook (though they trust it more than Washington and Wall Street).[81] The *Economist* summed up the prevailing mood: the big tech companies were "too big, anti-competitive, addictive and destructive to democracy"—BAADD.[82]

At Facebook's second annual Social Good Forum in December 2017, Mark Zuckerberg described how his company uses artificial intelligence to identify users who might be contemplating self-harm or suicide. He did not discuss the possibility that Facebook might itself be driving people to self-harm or suicide. As Deborah M. Gordon has suggested, online social networks replicate on a vast scale many of the more insidious features of friendship circles amongst girls in a middle school.[83] Using data from 5,208 adults over two years from a national longitudinal panel, Holly Shakya and Nicholas Christakis argue that "the more you use Facebook, the worse you feel."[84] They found that "most measures of Facebook use in one year predicted a decrease in mental health in a later year … [B]oth liking others' content and clicking links significantly predicted a subsequent reduction in self-reported physical health, mental health, and life satisfaction." The authors suggest that use of social media gives the impression of "meaningful social interaction" but is in fact no substitute for the real thing and therefore undermines wellbeing and health. Even Facebook's own research comes to similar conclusions about the effects of overuse of social media by students.[85]

Especially troubling is the effect of social media on children. As James Bridle has shown, Kid's YouTube seems designed to lure young users towards disturbing videos with titles like: "Surprise Play Doh Eggs Peppa Pig Stamper Cars Pocoyo Minecraft Smurfs Kinder Play Doh Sparkle Brilho" or "BURIED ALIVE Outdoor Playground Finger Family Song Nursery Rhymes Animation Education Learning Video." A search for "Peppa Pig dentist" leads to a video on which Peppa Pig is "tortured, before turning into a series of Iron Man robots and performing the Learn Colours dance." Others feature Peppa eating her father or drinking bleach.[86] Facebook Messenger Kids, launched in 2017, seems calculated to introduce children under 13 to the app in order to get them hooked as early as possible. The analogies with cigarettes and corn syrup in the 20th Century and opioids in the twenty-first are not fanciful, as Marc Benioff of Salesforce and Aza Raskin of Mozilla and Firefox have acknowledged.[87] There is some evidence that younger users of Facebook have been kicking the habit, but only in favor of Instagram (which Facebook owns) and Snapchat (which it would have liked to own).[88]

The revulsion against the power of the network platforms has not been confined to the United States. In Britain, Facebook has suffered a reputational hit, especially with older people. Although the evidence is less compelling that the company played both a malign and a decisive role in the Brexit referendum, journalists such as Carole Cadwalladr have done their utmost to make that case.[89] There was controversy over the role of social media in the 2018 Irish referendum on abortion, too, as Facebook and Google restricted advertisements in moves widely interpreted to be helpful to the proponents of constitutional change. Did Facebook interfere in the 2017 election in Iceland by selectively displaying its "I Voted" button on some voters' pages but not others?[90] Perhaps. In India, fake news stories on WhatsApp have certainly triggered riots, lynchings, and fatal beatings. In Sri Lanka, after a Buddhist mob attacked Muslims over a false rumor, a presidential adviser put it nicely: "The germs are ours, but Facebook is the wind." In Myanmar, too, violence against the Rohingya minority has been fueled, in part, by disinformation and incendiary content systematically spread on Facebook by the military. The United Nations investigator in charge of examining the persecution of the Rohingya, told the *New Yorker*, "I'm afraid that Facebook has now turned into a beast, and not what … was originally intended."[91] A Burmese legislator called the company "dangerous and harmful for our democratic transition."[92] It would be easy to give other examples. In Egypt, the Arab Spring was supposed to be a democratic revolution propelled by social media. Eight years later, Wael Ghonim—a leading figure in the Tahrir Square protests of 2010—is pessimistic. "We wanted democracy," he has said, "but got mobocracy." It seems unlikely that Kenyan democracy benefited from Cambridge Analytica's work for Uhuru Kenyatta in the 2013 and 2017 Kenyan elections. According to Freedom House, online manipulation and disinformation tactics played an important role in elections in 18 countries in 2016. In undemocratic regimes, too, the manipulation of social media is now standard practice, from China to Saudi Arabia.

Yet the analogies offered by the critics of the network platforms are not consistent. As one journalist complained with respect to Facebook:

> I've heard government metaphors (a state, the E.U., the Catholic Church, Star Trek's United Federation of Planets) and business ones (a railroad company, a mall); physical metaphors (a town square, an interstate highway, an electrical grid) and economic ones (a Special Economic Zone, Gosplan). For every direct comparison, there was an equally elaborate one: a faceless Elder God. A conquering alien fleet. … Maybe Facebook is a church and Zuckerberg is offering his benedictions. Maybe Facebook is a state within a state and Zuckerberg is inspecting its boundaries. Maybe Facebook is an emerging political community and Zuckerberg is cultivating his constituents. Maybe Facebook is a surveillance state and Zuckerberg a dictator undertaking a propaganda tour. Maybe Facebook is a dual power—a network overlaid across the U.S., parallel to and in competition with the government to fulfill civic functions—and Zuckerberg is securing his command. Maybe Facebook is border control between the analog and the digital and Zuckerberg is inspecting one side for holes. Maybe Facebook is a fleet of alien spaceships that have colonized the globe and Zuckerberg is the viceroy trying to win over his new subjects.[93]

This kind of muddle helps explain the inconsistencies in the global debate on regulation.

The European Commission, as George Soros foresaw, has taken the lead in seeking to regulate the U.S.-based network platforms. In June 2017, the Commission's antitrust division fined Google $2.7 billion for "anticompetitive practices" related to Google Shopping, the company's product comparison tool (specifically, for favoring the company's own site over competitors). The European Union's General Data Protection Regulation (GDPR), which came into force in May 2018, requires online services to make it easier for customers to transfer their information to other providers and even competitors, as well as strengthening people's control over their data. At the same time, the European authorities have also been active in regulating online "hate speech." In May 2016, Facebook, Microsoft, Twitter, and YouTube signed an agreement with the European Commission to "prohibit the promotion of incitement to violence and hateful conduct" by removing content inciting violence or hatred against protected groups within 24 hours of its being posted. On December 5, 2016, the companies announced plans for an industry database of "hashes"—unique digital signatures—of all extremist material banned on their platforms. National governments have added to this pressure. In 2017 Germany passed a law threatening Facebook, Twitter, and other social media companies with fines of $50 million if they failed to give users the option to complain about hate speech and fake news or refused to remove illegal content within 24 hours. After a series of terrorist attacks in London in 2017, British Prime Minister Theresa May and French President Emmanuel Macron threatened to impose steep fines on companies that failed to remove extremist propaganda from online platforms. Shortly thereafter, Google announced a four-part plan to address terrorist propaganda that included the increased use of technology to identify terrorist-related videos, the hiring of additional content moderators, the removal of advertising on objectionable videos, and the directing of potential terrorist recruits to counter-radicalization videos.[94]

These regulations are seen by some as a model for the United States. Another way of looking at them is as a halfway house between the American *laissez faire* regime and the Chinese system of much stricter state control. An important difference between Europe and China is that, unlike the Europeans, the Chinese have succeeded in building their own technology giants: the online retail site Alibaba, the search engine Baidu, and the Internet conglomerate Tencent. That these companies are subordinate to the Chinese state is clear. They are obliged to share their data with central authorities such as the People's Bank of China (PBoC), which therefore has access to users' payment history, creditworthiness, and contacts. This universal "back door" into the data represents the first step towards a comprehensive system of "social credit."[95] On June 1, 2017, a new cybersecurity law came into effect in China that requires technology companies to help the authorities remove content that "endangers national security, national honor and interests." The Chinese police are rapidly expanding their network of surveillance cameras and facial-recognition technology. “The political and legal system of the future is inseparable from the internet, inseparable from big data,” Alibaba’s Jack Ma told a Communist Party commission overseeing law enforcement in 2017. In future, he said, “Bad guys won’t even be able to walk into the square.”[96] In the context of a one-party state, the combination of ubiquitous smartphones, network platforms, big data and artificial intelligence makes possible a precision-targeted totalitarianism beyond the dystopian visions of Orwell and Huxley. At the same time, the immense attractiveness of the rapidly growing Chinese market to Western tech companies makes them susceptible to pressure from the Chinese government. For example, Twitter joined Facebook and YouTube in restricting social media accounts of popular dissident Chinese businessman Guo Wengui, in response to pressure from Beijing.

Regulation of the Internet is on the increase. Nearly half of the 65 countries assessed in *Freedom on the Net 2017* experienced declines in online freedom last year, while just 13 made gains, most of them minor. Less than one quarter of users reside in countries where the internet is designated “free.” Not only China, but also Venezuela, the Philippines, and Turkey were among 30 countries where governments were found to employ armies of “opinion shapers” to spread government views.[97] Yet it is not self-evidently obvious which is more dangerous: a regulated Internet, in which governments exercise at least some control over network platforms, or an unregulated one, in which private companies continue to gather and exploit the personal data of citizens for profit and without scruple.

Promises, Promises

Facebook’s response to the criticism directed against it over the past two years has been unconvincing. “Personally,” declared Zuckerberg two days after the 2016 election, “I think the idea that fake news on Facebook, which is a very small amount of the content, influenced the election in any way is a pretty crazy idea.” Despite warnings from Alex Tsamos about Russian disinformation—which dated back to the spring of 2016—Facebook’s leadership sought to play down the problem.[98] Zuckerbeg’s February 2017 manifesto, entitled “Building Global Community,” was long on aspirations, short on specifics. Four months later, Facebook unveiled a new mission statement to “give people the power to build community, to bring the world closer together.” The assertion that Facebook was engaged in community-building was disingenuous. The implication was that Facebook should not be held responsible, like a media company, for the content that appeared on its platform. “Things happened on our platform that shouldn’t have happened,” Sheryl Sandberg conceded. But “at our heart we’re a tech company. We hire engineers. We don’t hire reporters. No one is a journalist. We don’t cover the news.”[99] A former senior employee explained:

> The view at Facebook is that “we show people what they want to see and we do that based on what they tell us they want to see, and we judge that with data like time on the platform, how they click on links, what they like.” And they believe that to the extent that something flourishes or goes viral on Facebook—it’s not a reflection of the company’s role, but a reflection of what people want. And that deeply rational engineer’s view tends to absolve them of some of the responsibility, probably.[100]

This defense had lost credibility, however. In the course of 2017 it crumbled, despite the best efforts of Zuckerberg and Sandberg to lobby their way out of congressional scrutiny.[101]

In the course of 2017, a new strategy evolved, which might be characterized as preemptive self-regulation. On September 21, Zuckerberg pledged to increase the resources of Facebook’s security and election-integrity teams in order to work “proactively to strengthen the democratic process.” Facebook would henceforth require that all political ads disclose which Facebook page paid for them and ensure that every ad a given advertiser ran was accessible to anyone on the buyer’s page. Facebook would double the number of employees and contractors working on user safety and security issues to 20,000 by the end of 2018. It would also build new artificial-intelligence systems to detect what Zuckerberg described as “bad content and bad actors.” This could mean a significant increase in operating costs. But Zuckerberg told investors: “I am dead serious about this. I’ve directed our teams to invest so much in security on top of the other investments we’re making that it will significantly impact our profitability going forward.” These pledges amounted to an admission of responsibility for content.

Yet the practical consequences were confusing. In October 2017, Facebook introduced “Explore Feed,” which required media companies to pay for inclusion in the News Feed.[102] Further changes were introduced in January 2018, when Zuckerberg announced that News Feed would prioritize “meaningful interaction” over “passive consumption of low-quality content,” demoting “things like clickbait headlines and false news, even though people often click on those links at a high rate.” He added: “We want to make sure that our products are not just fun, but are good for people … good for the world.” At the same time, Facebook would start to boost certain publishers whose content was “trustworthy, informative, and local,” according to reader surveys. In the course of 2018, Zuckerberg made numerous such announcements:

- In January Facebook hired Nathaniel Gleicher, the former director for cybersecurity policy on President Obama's National Security Council, to counter "information operations."
- In July, it removed thirty-two accounts running disinformation campaigns that were traced to Russia. A few weeks later, it removed more than six hundred and fifty accounts, groups, and pages with links to Russia or Iran.
- In March, Zuckerberg pledged "dramatically [to reduce] the amount of data that developers have access to, so that apps and developers can't do what Kogan did" in providing data to Cambridge Analytica.[103]
- That same month, he said that "People should know who is buying the ads that they see on Facebook, and you should be able to go on any page and see all the ads that people are running to different audiences." He said that this would be in place by the November 2018 midterms.[104]
- Also in March, Facebook announced that it would "try to make privacy settings clearer by creating a central hub where users can examine the data they are sharing" with third-party developers."[105]
- In September he announced a "three-year project" to "rebuild all of our content enforcement systems to proactively find harmful content rather than wait for people to flag issues." He added: "It is our responsibility to amplify the good and mitigate the bad."[106]
- A new and extended 30-page version of Facebook's Community Standards was released in April 2018, defining more clearly what "hate speech" meant. At the same time, an appeals process was created for six content categories: nudity, sexual activity, hate speech, graphic violence, bullying, and harassment.
- In May 2018, Facebook issued its first transparency report, providing examples of different types of content takedowns.
- In November 2018, Zuckerberg proposed "an external appeals 'court'," to rule on disputed cases.
- In January 2019, Facebook announced the removal of 364 Facebook pages and accounts linked to former Soviet republics, as well as 107 Facebook pages, groups, and accounts and 41 Instagram accounts operating in Ukraine that it believed to be under Russian control.[107]

A number of similar measures of self-regulation were announced by the other network platforms:

- In April 2017, Google announced the release of "Fact Check": "For the first time, when you conduct a search on Google that returns an authoritative result containing fact checks for one or more public claims, you will see that information clearly on the search results page." This was another sop to established publishers whose revenue Google and Facebook had been devouring.
- In October 2017 Google ended its "first click free" policy, which required publishers to give away some stories in order to appear high in its search rankings. Google chief executive Pichai spoke of "a flight to quality."[108]
- In November 2017, Google announced the "Trust Project," hosted by Santa Clara University and developed in conjunction with more than 75 news organizations worldwide. The goal of the project would be to introduce eight "trust indicators," such as "author expertise," "citations and references," and "diverse voices," into the search ranking algorithm.
- In December 2017, Google announced its intention to increase to 10,000 the number of employees tasked with removing extremist content from YouTube.

In short, the network platforms' response to the crisis of 2016 has been a barrage of promises to regulate themselves. "I actually am not sure we shouldn't be regulated," Zuckerberg said in a CNN interview in March. "I actually think the question is more, what is the right regulation rather than 'Yes or no, should it be regulated?'"[109] He told Ezra Klein that he could imagine "some sort of structure, almost like a Supreme Court, that is made up of independent folks who don't work for Facebook, who ultimately make the final judgment call on what should be acceptable speech in a community that reflects the social norms and values of people all around the world."[110] This was an unusual concession, as Zuckerberg—like his Roman role model Augustus—rarely offers to limit his own power. "One of the things that I feel really lucky we have," he also told Klein, "is this company structure where, at the end of the day, it's a controlled company. We are not at the whims of short-term shareholders. We can really design these products and decisions with what is going to be in the best interest of the community over time." This was an allusion to the preference shares that give Zuckerberg control over Facebook. Advertising might be Facebook's principal source of revenue, he conceded, but that was only so that his platform could be free to users. "I think probably to the dissatisfaction of our sales team here, I make all of our decisions based on what's going to matter to our community and focus much less on the advertising side of the business."

In addition to numerous interviews, the leaders of the big technology companies have set out the case for self-regulation in a succession of hearings: on October 31-November 1, 2017, in April 2018—when Zuckerberg himself testified before two congressional committees—and again in September 2018. "Which are you?" Senator Dan Sullivan (R-Alaska) asked Zuckerberg in April. "Are you a tech company, or are you the world's largest publisher?"

> Zuckerberg: I view us as a tech company, because the primary thing that we do is build technology and products.
>
> Sullivan: You said you were responsible for your content, which makes you kind of a publisher, right?
>
> Zuckerberg: I agree that we are responsible for the content, but we don't produce the content. I think that when people ask us if we are a media company or a publisher, my understanding of what the heart of what they are really getting at, is "Do we feel responsibility for the content on our platform?" And the answer to that, I think, is clearly "Yes."

This recognition of responsibility for content was an important moment in the April hearings. Also important, however, was Zuckerberg's assertion that Facebook is "a system of different things: we compete with Twitter as a broadcast medium; we compete with Snapchat as a broadcast medium; we do messaging, and iMessage is default-installed on every iPhone." This was part of a wider argument he sought to make against any attempt to apply antitrust law to Facebook: Silicon Valley was the home of cut-throat competition not anti-competitive practices.

The scale of the self-defense effort by Silicon Valley since 2016 cannot be understated. Facebook has followed Google's lead by investing heavily in Washington lobbying, hiring Joel Kaplan, a former policy adviser to President George W. Bush, Sandy Luff, a former aide to Attorney General Jeff Sessions, and Kevin Martin, the former Federal Communications Commission chairman. According to data from the Center for Responsive Politics, Facebook had contributed a total of $641,685 since 2014 to the members of Congress that Zuckerberg faced during his visit to Capitol Hill. The top recipients of that money included Senators Cory Booker and Kamala Harris. Facebook also employed the public relations company Definers Public Affairs (and its tame news site NTK Network). At one and the same time, the company sought to link the activist group Freedom from Facebook to George Soros and appealed to the Anti-Defamation League to represent FfF's criticism of the company as antiSemitic.[111]

However, these efforts could not prevent a flurry of activity by legislators and regulators, as well as continued criticism by academics.[112] In August 2017, a rare bipartisan alliance of 27 Democratic and Republican senators introduced the Stop Enabling Sex Traffickers Act (SESTA), despite opposition from the Internet Association and Google. It was passed, along with the Allow States and Victims to Fight Online Sex Trafficking Act (FOSTA), in April 2018. In October Senators Amy Klobuchar, Mark Warner and John McCain introduced the Honest Ads Act (S. 1989), which required the network platforms to reveal the buyers and content of all campaign-related ads as well as to maintain a public list of all political advertisers spending $500 or more. Other legislative initiatives included Keith Ellison's 21st Century Competition Commission Act, an antitrust bill; Senator Edward J. Markey's bill to commission research on the impact of technology on children's health, and a Californian measure to prohibit the use of bots without identification (the so-called "Blade Runner" law SB-1001). Senator Klobuchar also proposed two bills to change the criteria for mergers, an idea appealing not only to her Massachusetts colleague, Elizabeth Warren, but also to Republican Senators Mike Lee of Utah and Josh Hawley of Missouri. Meanwhile, the Federal Trade Commission (FTC) continued its "open non-public investigation" into whether or not Facebook had violated the terms of the 2011 consent decree.[113] In the course of 2018, the FTC held a series of ten hearings on competition and consumer protection. During his confirmation hearings, Attorney General William P. Barr said he was "for vigorous enforcement of the antitrust laws to preserve competition," a sentiment echoed by the head of the Antitrust Division, Makan Delrahim.[114] By the beginning of 2019 calls for new regulation had gone beyond antitrust. A recurrent theme was that more must be done to ensure that users' privacy must be protected—and that the platforms could not be trusted to do it themselves. California's legislature passed a new privacy law that, from 2020, will empower customers to sue technology companies if they can prove that their data have been illegally accessed.[115]

Not to be outdone, President Trump himself began to attack Silicon Valley, though he turned its own weapons against it. "Facebook was always anti-Trump," he tweeted in October 2017. ("That's what running a platform for all ideas looks like," Zuckerberg responded in a Facebook post.) In 2018, Trump attacked Amazon for tax evasion and taking advantage of the U.S. Postal Service, telling reporters that "Amazon is just not on an even playing field. You know, they have a tremendous lobbying effort, in addition to having The Washington Post [which Jeff Bezos acquired in 2013], which is, as far as I'm concerned, another lobbyist." He was, he said, "going to take a pretty serious look" at Amazon because "the playing field has to be leveled." According to an unnamed source who had spoken to Trump, the president had "wondered aloud if there may be any way to go after Amazon with antitrust or competition law."[116]

Even Trump's preferred platform, Twitter, was not spared. On July 26, 2018, he accused Twitter of "shadow banning" Republicans and said he would look into what he called a "discriminatory and illegal practice." The big tech companies "better be careful because you can't do that to people," Trump said in August, shortly after the expulsion from Facebook of far-right conspiracy theorist Alex Jones. "I think that Google, and Twitter and Facebook, they are really treading on very, very troubled territory and they have to be careful. It is not fair to large portions of the population." Google, too, came under attack. In August 28, 2018, Larry Kudlow, Trump's economic adviser, said that the administration was "taking a look" at whether or not Google and its search engine should be regulated by the government. Trump complained that "Fake CNN" was "prominent" in search results for him. "They have it RIGGED, for me & others," he complained, "so that almost all stories & news is BAD." Conservative media were being "shut out." Was this "illegal?" he asked, accusing Google of "controlling what we can & cannot see." In a tweet on August 18, Trump accused social media of "totally discriminating against Republican/Conservative voices. … They are closing down the opinions of many people on the RIGHT, while at the same time doing nothing to others."

Perhaps the only surprising thing is that this last complaint—of political bias—was not made sooner. FEC disclosures from the 2016 presidential campaign showed that 95 percent of big tech employees' donations went to Hillary Clinton, and only 4 percent to Donald Trump. The liberal politics of Silicon Valley was regularly on display in 2017, beginning as early as January, when thousands of Google employees walked out of the office for a rally sanctioned by the company—and indeed addressed by chief executive Sundar Pichai—to protest against Trump's executive order banning travel from seven predominantly Muslim countries. Google's firing of James Damore, for writing an in-house essay that questioned the desirability of hiring and promoting more female engineers, seemed to furnish fresh evidence of a corporate culture skewed to the left. In March 2018 YouTube imposed "strikes" on the videos of some prominent far-right actors and conspiracy theorists, including Mike Cernovich, Infowars (Alex Jones), Atomwaffen and Sargon of Akkad (Carl Benjamin). The company maintained that its "reviewers remove content according to our policies, not according to politics or ideology."[117]

In a similar way, employees at Facebook appeared to have been dismissed for apparently political reasons. Benjamin Fearnow was fired from Trending Topics for leaking Zuckerberg's condemnation of an "All Lives Matter" sign at Facebook. In the words of Brian Amerige, a senior Facebook engineer, and founder of "FB'ers for Political Diversity," the company had "a political monoculture that's intolerant of different views. We claim to welcome all perspectives, but are quick to attack—often in mobs—anyone who presents a view that appears to be in opposition to left-leaning ideology." When Facebook imposed an outright ban on the anti-immigration, anti-Islam group "Britain First," it explained that the group had used language "designed to stir up hatred against groups in our society." On July 27, after a direct appeal from the parents of a child killed at Sandy Hook, Facebook took down four Infowars videos and suspended Alex Jones for a month. On August 5 Apple stopped distributing five podcasts associated with Jones on the ground that they purveyed "hate speech." Facebook also shut down four of Jones's pages for "repeatedly" violating rules against hate speech and online bullying. Zuckerberg's attempt to explain his reluctance to ban Jones backfired when he explained to the journalist Kara Swisher:

> The principles that we have on what we remove from the service are: If it's going to result in real harm, real physical harm, or if you're attacking individuals, then that content shouldn't be on the platform. [But] … The approach that we've taken to false news is not to say: You can't say something wrong on the internet. I think that that would be too extreme. Everyone gets things wrong, and if we were taking down people's accounts when they got a few things wrong, then that would be a hard world for giving people a voice and saying that you care about that. … I'm Jewish, and there's a set of people who deny that the Holocaust happened. I find that deeply offensive. But at the end of the day, I don't believe that our platform should take that down because I think there are things that different people get wrong. [118]

The resulting storm of criticism illustrated the shift in attitudes in Silicon Valley. The libertarian instincts of an earlier generation of Silicon Valley entrepreneurs were being forced to yield to the more censorious attitudes of more recently hired employees who had been schooled in the modern campus culture of "no platforming" any ideas deemed to be "unsafe."

For conservatives, as well as for right-wing populists, the alarm bells could no longer be ignored. Alex Marlow, editor-in-chief of Breitbart News, and film-makers Peter Schweizer and James O'Keefe were among those to add their voices to the growing chorus of complaint about Silicon Valley's bias. In *Prager University v. Google*, conservative broadcaster Dennis Prager accused YouTube of violating his first amendment rights by "regulat[ing] and censor[ing] speech as if the laws governing free speech and commerce do not apply to it." Facebook was forced to apologize to Prager for removing videos with the titles "Where Are the Moderate Muslims?" and "Make Men Masculine Again." Writing for Breitbart in late October 2018, Brad Parscale accused "Big Tech monsters like Google and Facebook" of having become "nothing less than incubators for far-left liberal ideologies and … doing everything they can to eradicate conservative ideas and their proponents from the internet." This was, Parscale argued, "an existential threat to our individual liberties as well as our system of government."[119]

Renee DiResta might insist that the platforms "need to be able to take down users and sites that fail the tests of authenticity, organic distribution and integrity reputation." Jonathan Albright might point out that, during the 2018 midterm elections, the suspect Facebook Pages (with foreign "manager" accounts) or the Facebook Groups used to spread scare stories—often by "gaming the platform's metrics"—were mostly right-wing in character and content.[120] But it was surely inevitable that those who fell foul of supposedly "viewpoint agnostic moderation" would complain of politically motivated censorship.[121] At the time of writing, allegations of anti-conservative bias on network platforms were being made with increasing frequency.[122] A related allegation is that Google's seemingly innocuous "Go Vote" message on election day 2018 was disproportionately helpful to Democratic candidates.[123]

So What Is to Be Done?

American lawmakers must by now realize that the status quo is indefensible. The network platforms currently enjoy unprecedented power over the public sphere not only in the United States but around the world. Yet they have shown themselves to be very poor custodians of their users' personal data. There may be other harms arising from their applications' addictive character. They have shown themselves to be vulnerable to abuse by malevolent foreign and domestic actors. And they can no longer plausibly claim not to be publishers or media companies, as they are increasingly under pressure to curate, sort and otherwise manage the content that they host, and to do so in ways that have significant political implications. Minor modifications of the law, such as FOSTA, do not address the fundamental question of how to limit the power and capacity for harm of the big tech companies.

Five different proposals for change have been put forward in the past two years of debate:

1. Scrap "net neutrality" in order to empower the Internet service providers (ISPs) relative to the network platforms;
2. Update antitrust doctrine and law so that the network platforms can be broken up;
3. Increase the regulation of the network platforms by either the FTC or the FCC, acknowledging that the platforms are now public utilities;
4. Repeal Section 230 largely or wholly, thereby making the network platforms legally liable for the content they host, and leave the rest to the courts; and
5. Impose the equivalent of First Amendment obligations on the network platforms, recognizing that they are too important a part of the public sphere—the modern "town square"—to be able to regulate access to it on the basis of their own privately determined and almost certainly skewed "community standards."

The first of these options, the ending of so-called net neutrality, is the only one to date that has been acted upon. Unfortunately, it is also the least likely to be effective. FCC chairman Ajit Pai repealed Obama-era rules that were intended to ensure equal access to the Internet by preventing ISPs from charging users more to see certain content and to curb access to some websites. Under Pai's new regime, ISPs are able to block access, slow down or speed up service as they see fit. For AT&T, Comcast, and Verizon this had obvious appeal; for the network platforms it was much less attractive. The two kinds of company adopted predictably opposing stances. From the point of view of the citizen-consumer, however, the net benefits of scrapping net neutrality are not obvious. The most likely outcome would seem to be that Internet users will end up paying more for certain kinds of content and enduring slower speeds for other kinds. That does not address any of the fundamental problems described above. It merely shifts power from the network platforms to the ISPs. There is no reason to believe that they are superior actors from the point of view of the public interest in a functioning market for ideas.

Of considerably more importance is the revival of interest in antitrust as a tool. This began on the left, with the so-called "antitrust hipsters" around Barry Lynn of the Open Markets Institute, who was effectively ousted from the New America Foundation in 2017 because of Silicon Valley donors' opposition to his work. In the past year, however, interest has grown more widespread in the idea that the big tech companies are simply too big. The analogy with Standard Oil, which was broken up more than a century ago, is drawn with increasing frequency. In Washington, proponents of antitrust refer to Amazon, Facebook, and Google as "Standard Commerce, Standard Social and Standard Data." Tim Wu of Columbia Law School argues that Facebook should relinquish Instagram, Messenger, and WhatsApp; Google should give up YouTube and DoubleClick; Amazon should spin off Amazon Web Services. Scott Galloway takes a similar view and includes Apple on his list.[124] Such arguments have ceased to be the preserve of progressives.[125]

Antitrust law was developed in the United States as a response to the rise of trusts and combinations—that is, cross-ownership and management structures facilitating collusion. The focus of early antitrust law was enterprise size, market share, and strategic market positioning. At first, John D. Rockefeller won plaudits for bringing down the price of kerosene from 26 cents in 1870, when his market share was 4 percent, to 7 cents in 1890, when his market share reached 90 percent. Standard Oil was feted for bringing efficiency to the oil industry, which in turn fueled the development of steel, railroads, and the technological industries associated with them. However, in the early 1900s, journalists such as Ida Tarbell helped turn public opinion against Rockefeller (the "King of Combinations") and his allegedly anti-competitive tactics. Trusts became a political lightning rod. Supreme Court Justice Louis D. Brandeis coined the phrase "the curse of bigness." This laid the foundation for the 1911 landmark judgment in *Standard Oil Co. of New Jersey v. United States*, which broke Standard Oil up into 34 separate firms.

For half a century, Brandeis's approach was dominant in the courts. In 1956, the biggest technology company of that era, AT&T, settled an antitrust case by signing a consent decree in which it made two important concessions: it agreed not to expand its business into new markets (such as computers) and it made all of its patents available to others at no charge. The consent decree did not break up AT&T—that did not happen until 1984—but it was an essential first step, and one with great consequences as one of the many patents in AT&T's portfolio was the transistor. In 1968 the Department of Justice's merger guidelines were that any acquisition of a company with a market share above 3 percent by one with a share above 15 percent should be challenged. However, this view of antitrust law was challenged by Robert H. Bork and Ward S. Bowman in their seminal 1965 article "The Crisis in Antitrust," which asserted that the intention of antitrust law was simply the protection of "consumer welfare." Bork and Bowman argued that antitrust law could in fact harm consumer welfare by punishing aggressive pricing and preventing mergers that would reduce costs and therefore prices. "Consumer welfare" came to be equated with economic efficiency, notably in the Reagan administration's 1982 merger guidelines, which aimed to proscribe the "ability of one or more firms profitably to maintain prices above competitive levels"—a radical departure from the previous 1968 guidelines, which had aimed "to preserve and promote market structures conducive to competition." Another change was the narrowing of the definition of "barriers to entry" to exclude incumbent advantages from economies of scale and capital requirements. The argument was that all firms were subject to the threat of potential competition, so market power was always fleeting, and antitrust enforcement rarely needed.

This helps explain why neither IBM in the 1960s nor Microsoft in the 1990s was split up, despite their dominance of, respectively, computer hardware and software. The most that happened was that IBM had to open its platform to independent software developers and Microsoft was obliged to disclose details about the workings of its Windows operating system to rivals. The failure of the antitrust action against Microsoft looks, with hindsight, like a major turning point. Microsoft had sought to bundle its Windows operating system with its own web browser, Internet Explorer, to the disadvantage of its rival, Netscape. On April 3, 2000, Judge Thomas Penfield Jackson ruled that Microsoft had committed monopolization, attempted monopolization, and tying, in violation of the Sherman Antitrust Act and on June 7, 2000, ordered a breakup of Microsoft. However, Jackson was removed from the case after he talked to reporters in an off-the-record discussion before his final decision. The DC Circuit Court of Appeals then overturned Jackson's rulings and in 2002 the company reached a settlement with the Department of Justice that left it intact. Microsoft survived, though it has been argued that the warning shot fired by Judge Jackson had a significant effect on the subsequent conduct of Bill Gates and his colleagues. Plans to program Internet Explorer so that it would redirect users away from Google to MSN Search, or simply warn them against Google, were quietly shelved.

The interpretation and enforcement of today's U.S. competition laws follow two main principles: the *per se* concept, whereby behaviors with no judicially redeeming characteristics are illegal *per se*, and the rule of reason, where the illegality of behaviors rests on their probable negative effect on market competition or in creating "restraints of trade." Antitrust laws and market regulations are enforced in three ways: criminal and civil enforcement actions brought by the Department of Justice's Antitrust Division, civil enforcement actions brought by the FTC, and lawsuits brought by private parties asserting damage claims. Yet in practice the law did almost nothing to stand in the way of the emergence of the network platforms. How could "consumer welfare" be harmed by "free" services such as Facebook? Where were barriers to entry to the Internet? The absence of good answers to such questions encouraged a sense of impunity in Silicon Valley. During a 2011 Senate Judiciary Committee antitrust hearing, Alphabet Executive Chairman Eric Schmidt observed: "It's also possible to not use Google search." If competition was "one click away," how was Google in breach of antitrust law? In his book *Zero to One*, Paypal founder and early Facebook investor Peter Thiel argued that competition was for losers, as it eroded profits; one should invest only in companies with a shot at establishing a monopoly. At the start-up incubator Y-combinator, according to its president Sam Altman, "We … ask how the company will one day be a monopoly … We're looking for businesses that get more powerful with scale and that are difficult to copy."

In reaction to such boasts, some attempt to revitalize antitrust law was highly likely. In "Amazon's Antitrust Paradox" (2017), Lina M. Khan argued that the big tech companies had created conflicts of interest by using their dominant platforms to promote their own services against those of their competitors. This was in line with the reasoning advanced by the European Commission when it fined Google for giving more prominence to its own shopping services in search results. Khan also revived the presumption of predation in cases of below-cost pricing, which others have argued creates pernicious long-term effects on productive investment, worker wages, product quality, and consumer choice.[126] Writing in 2018, Roger McNamee, one of Facebook's earliest investors, called for (among other things) a ban on further acquisitions by the network platforms, an insistence on more equitable end-user license agreements with meaningful opt-out clauses, a return of data to the ownership of consumers, and—last but not least—a revival of "the country's traditional approach to monopoly."[127]

There are two problems with the attempt to revive pre-Bork and Bowman antitrust, aside from the difficulty of achieving a rapid paradigm shift in the minds of judges. First, the software industry is prone to natural monopoly or oligopoly. Software reduces friction along existing value chains, often reconfiguring industries into two-sided markets, with platforms intermediating between owners of assets or providers of services and consumers. As demonstrated by the work of Nobel laureate Jean Tirole, two-sided markets exhibit powerful network effects and therefore produce just a few dominant firms, an effect compounded by the aggressive use of patents in software markets and a deliberately engineered lack of inter-operability. Preferential attachment models, such as those developed by the physicist Albert-László Barabási, describe markets where strength begets strength. For example, in social media networks, the users with the most friends or followers are the most likely to get any additional nodes added to the network. This produces a power-law-like distribution that makes it all but impossible for latecomers to succeed. These insights from network science challenge the Chicago School's belief that high capital requirements or regulation are the most important barriers to market entry; network effects may matter more. Breaking up network platforms would reduce these effects because the whole network is genuinely more valuable than the sum of its parts (Metcalfe's law).

In any case, the historian is bound to point out that even supposedly successful antitrust actions in reality achieved much less than was intended. The outcome of the breakup of Standard Oil was in fact to make Rockefeller even richer. The 34 "Baby Standards" were worth a combined $600 million in 1911, swelling to $2.9 billion in value and paying out $920 million in dividends by 1921, making Rockefeller one of the wealthiest men in history. The interoperability forced on Microsoft, which allowed rivals to make their products more compatible with Windows, only made Microsoft more central to the software ecosystem. After the original 2001 judgment ordering the breakup of Microsoft, many predicted the company's decline. But Microsoft retained more than 90 percent of the operating system market until 2017 and Apple did not surpass it in market capitalization until 2010—only to lose that lead in 2018.

If breaking up big tech is either impossible or pointless, the obvious answer is to regulate the network platforms as utilities. There are ample precedents. In the 1876 case of *Munn vs. Illinois*, the Supreme Court upheld the power of the government to regulate private industries such as the railroads, though the railroads defied any attempts to regulate prices and the Munn decision was eventually reversed. The creation in 1886 of the Interstate Commerce Commission gave the government real power to intervene, but enforcement was left to the courts, which usually sided with the railroads. It was not until the 20th Century that the railroads came to be regulated—in other words, after the sector had leveled off in innovation, growth, and profitability. History shows that, in the United States, government intervention often serves to cement the dominance of large players for the foreseeable future, and that equity holders generally fare quite well. As Gabriel Kolko demonstrated in two seminal books—*Railroads and Regulation* and *The Triumph of Conservatism*—late 19th-century regulation was often designed to stabilize profitability in oligopolistic industries that had reached maturity.[128] It is therefore hard to summon up much enthusiasm for the prospect of a more assertive FTC—or for that matter FCC—working in tandem with the biggest companies of Silicon Valley to create a regulatory framework very likely to entrench their market dominance. We have seen this movie before.

Senator Mark Warner's October 2018 paper on "Potential Policy Proposals" illustrates what a more heavily regulated tech sector might have to contend with.[129] Out of twenty proposals, two have already made it into law:

1. A duty to label bots ("Blade Runner" law).
2. Disclosure requirements for online political advertisements (Honest Ads Act).

Four of Warner's proposals envisage new legislation:

3. Comprehensive (GDPR-like) data protection legislation, to be enforced by a new enforcement agency.
4. Legislation banning so-called "dark patterns" such as Facebook nudges to induce users to upload their contacts.

5. A Data Transparency Bill that would "require companies to more granularly (and continuously) alert consumers to the ways in which their data was being used, counterparties it was being used with, and … what each user's data was *worth* to the platform."

6. A Data Portability Bill: "predicated on a legal recognition that data supplied by … users (or user activity) is the users'—not the service provider's."

Only one Warner proposal envisages reducing the scope of legislation, namely removing Section 230 immunity for state-level "dignitary torts" such as defamation, false light, public disclosure of private facts.

Four other proposals would, implicitly or explicitly, increase the power of existing federal agencies or create a new agency:

7. Restore the FTC's rulemaking authority with respect to privacy.

8. Create an interagency task force for countering asymmetric threats to democratic institutions.

9. Establish a public initiative for media literacy.

10. Increase deterrence against foreign manipulation.

A total of eight would impose new duties or responsibilities on network platforms, which presumably one or more of these regulatory agencies would enforce:

11. A duty to determine [the] origin of posts and / or accounts.

12. A duty to identify inauthentic accounts.

13. And information fiduciary duty, stipulating "not only that providers had to zealously protect user data, but also [that they had to] pledge not to utilize or manipulate the data for the benefit of the platform or third parties."

14. An obligation to make anonymized activity data available to "independent, public interest researchers" (Public Interest Data Access Bill).

15. A requirement to get first-party consent for data collection.

16. A requirement to provide consumers with the sources of data used to make algorithmic determinations or classifications.

17. A requirement for platforms to make their services interoperable with other platforms.

18. An obligation to make "essential facilities" (e.g., Google Maps) publicly available a "fair, reasonable, and non-discriminatory" (FRAND) terms.

Only one (the requirement to open federal datasets to university researchers and qualified small businesses) asks anything from the government.

This is a list fraught with technical difficulties. For example, it is often impossible to determine the origins of posts and accounts. Insisting on identification could have unintended consequences for users living in authoritarian states and even for the civil rights of those living in democracies. Is it entirely clear to whom data belongs, if it is data based on the platform's observation of user behavior? Would not increased portability of data go hand in hand with increased insecurity? But the biggest objection to Warner's approach is that, as in the past, increasing the powers of federal agencies would incentivize the already establish network platforms to collude with those agencies to raise entry-barriers. For example, data portability might entrench incumbents by obliging any new entrants to give them access to their data.

A fourth approach would go beyond Warner's more limited proposal on Section 230 and repeal it altogether, ending the exemption of network platforms from liability for the content they host. That exemption made sense when these companies were fledglings or did not yet exist. Today, it gives the biggest companies in the world both power and influence without responsibility or accountability. One appealing feature of getting rid of Section 230 is that it would be left to the courts to bring the network platforms to heel when plaintiffs could show that a harm had arisen from, say, a fake news story disseminated by Facebook's News Feed. Already the courts have established that the network platforms are not exempt from liability to warn and product liability, because they have a duty of care to warn users of potential dangers. Indeed, given social media's uniquely deep and wide knowledge of users' interactions and relationships, they have unprecedented abilities to foresee potential harms. In one important case, a model ("Jane Doe") had been lured by scammers on ModelMayhem, a social network for models and photographers, who then drugged and raped her, filming the incident for a pornographic video. Internet Brands, the owner of ModelMayhem, was aware of this rape ring, but did not warn any of its users. In *Jane Doe No. 14 v. Internet Brands, Inc.*, the 9th circuit court ruled that "Doe's negligent failure to warn claim did not seek to hold Internet Brands liable as the 'publisher or speaker of any information provided by another …' and therefore the Communications Decency Act did not bar the claim." This case has established an important limit on Section 230 immunity, but it also exposes the anachronistic nature of Section 230 itself. If Internet Brands should have warned Jane Doe of the dangers of using ModelMayhem, why should not Facebook have warned all its users of the dangers of Russian-generated fake news targeted at them through the News Feed?

There is, however, an important corollary. If we are to end the fiction that network platforms are not, in some respects, media companies or publishers, then we must at the same time end the equally dangerous fiction that they are not in many respects the modern public sphere. A first step has already been taken in this direction. In *Packingham v. North Carolina* (2017), the Supreme Court overturned a state law that banned sex offenders from using social media. In the opinion, Justice Anthony Kennedy likened internet platforms to "the modern public square," arguing that it was therefore unconstitutional to prevent sex offenders from accessing, and expressing opinions, on social network platforms. In other words, despite being private companies, the big tech companies have, in some cases, a public function. "While in the past there may have been difficulty in identifying the most important places (in a spatial sense) for the exchange of views," Justice Kennedy wrote, "today the answer is clear. It is cyberspace—the vast democratic forums of the Internet in general, and social media in particular." In May 2017 the Southern District of New York gave a similar ruling in *Knight First Amendment Institute v. Donald J. Trump, Hope Hicks, Sarah Huckabee Sanders and Daniel Scavino*:

> We hold that portions of the @realDonaldTrump account—the "interactive space" where Twitter users may directly engage with the content of the President's tweets—are properly analyzed under the "public forum" doctrines set forth by the Supreme Court, that such space is a designated public forum, and that the blocking of the plaintiffs based on their political speech constitutes viewpoint discrimination that violates the First Amendment.[130]

As president, Donald Trump could not therefore block Twitter users from seeing his tweets.

If the network platforms are the new public sphere, then it cannot be their responsibility to remove "hateful content," as 19 prominent civil right groups demanded of Facebook in October 2017, because hateful content—unless it explicitly instigates violence against a specific person—is protected by the First Amendment. To be sure, Kate Klonick has argued that tech companies should not "be held to a First Amendment standard," because that would mean "porn stays up, spam stays up, everything stays up." But this is not convincing. It is better that porn and spam "stay up" than that our freedom of speech be circumscribed by the community standards of unaccountable private companies, run by men who imagine themselves to be emperors. The danger of a piecemeal erosion of Section 230 is that it leads to creating "censorship creep," by encouraging platforms to "over-moderate." If outright repeal is too bold a step, with too many unforeseeable consequences, then a better compromise would be to create a blanket exception to 230 for "bad actors" who "knowingly and intentionally leave up unambiguously unlawful content that clearly creates a serious harm to others" (as proposed by Geoffrey Stone) or for "online service providers that intentionally solicit or induce illegality or unlawful content" (Stacey Dogan's formulation) or for platforms that "can[not] show that their response to unlawful uses of their services is reasonable." At the very least, a new Section 230 might read: "No provider or user of an interactive computer service *that takes reasonable steps to prevent or address unlawful uses of its services* shall be treated as the publisher or speaker of any information provided by another information content provider *in any action arising out of the publication of content provided by that information content provider*."[131]

Other modifications are also conceivable. Jonathan Zittrain has proposed that "companies below a certain size or activity threshold could benefit from [Section 230'] s immunities, while those who grow large enough to facilitate the infliction of that much more damage from defamatory and other actionable posts might also have the resources to employ a compliance department." Alternatively, a distinction could be drawn "between damages for past acts and duties for future ones … leading only to responsibility once the knowledge is gained and not timely acted upon." Or "a refined CDA could take into account the fact that Facebook and others know exactly whom they've reached," so that the new remedy for defamation "would less be to assess damages against the company for having abetted it, but rather to require a correction or other follow up to go out to those who saw—and perhaps came to believe—the defamatory content."[132] Even these weaker modifications of Section 230 would meaningfully increase the legal costs of network platforms.[133] Combined with a requirement to act as if the First Amendment applies in cyberspace, these additional liabilities seem a perfectly justifiable way of countering the various negative externalities currently created by these immensely profitable companies—and a much more elegant solution than probably futile attempts to break them up or regulate them through government agencies.

What's Good for Facebook…

In the coming months and years, there will be a profoundly important debate about how best to contend with the unintended consequences of organizing the once decentralized worldwide web around a few network platforms. The big technology companies themselves will offer many different defenses, no doubt. But only one of them is likely to be effective. Testifying before a congressional committee in 2018, Mark Zuckerberg unveiled that winning argument: "I think", he said, "that anything that we're doing to constrain [the big tech companies] will, first, have an impact on how successful we can be in other places. I wouldn't worry in the near term about Chinese companies or anyone else winning in the U.S., for the most part. But there are all these places where there are day-to-day more competitive situations—in Southeast Asia, across Europe, Latin America, lots of different places."[134]

It was more than 60 years ago that the CEO of General Motors, Charles Wilson, made his famous claim that what was good for General Motors was good for the United States. Zuckerberg's version might be paraphrased as "What's bad for Silicon Valley is good for China." Regulate big tech at your peril, in other words, for you may unintentionally help the Chinese competition. It is an argument that represents a radical break with Facebook's past ambition to build a global community—including in China, if possible. The rapid deterioration of relations between the United States and China since 2016 (to say nothing of the growing burden of EU regulation) creates a new possibility: that Facebook and its peers must reconfigure themselves as American corporations, whose interests are aligned with those of the nation-state within which they have their headquarters, and on whose stock market they are listed.

The debate on the regulation of big tech, in other words, is inseparable from a debate about national cybersecurity. The assumption of a decade ago—that aggression in cyberwarfare would primarily consist of the use of computer viruses to disable the software running critical infrastructure—has proved wrong. Instead, we find ourselves the targets of an asymmetrical "information war," pioneered by the Russians but not exclusively waged by them. Memes, not malware, are the weapon of choice.[135] Yet this in turn may soon be superseded as artificial intelligence begins to be deployed by China as a tool of inter-state competition. The American technology companies have spent many years seeking to ingratiate themselves with China's rulers, in the hope of increasing their access to that country's vast market for hardware, software and web services. In the words of a recent and damning report on Chinese influence operations:

> Facebook has been notably solicitous of the Chinese government in an effort to enter the Chinese market, reportedly developing a tool that could be used by a third party to censor content. Despite being blocked in China, Facebook nonetheless generates significant advertising revenues from Chinese companies seeking to reach foreign consumers. As it seeks to reenter the Chinese market, Google's willingness to facilitate that country's national artificial intelligence priorities stand in contrast to its decision to end limited AI cooperation with the US Department of Defense. In June 2018, Tsinghua announced that Google's AI chief would serve as an adviser to that university's new center for artificial intelligence research. The company is already involved in research at Peking University and the University of Science and Technology of China, among others.[136]

To say the least, such initiatives are not easy to square with Zuckerberg's claim that what is good for Facebook is good for the United States.

Conclusion

In a series of brilliant essays, the former Facebook product manager Sam Lessin has challenged much received wisdom about the Internet. As he puts it, the Internet "increasingly represents a strange hybrid of the public and private spheres." The problem revealed in 2016, he argues, was not fake news or feed ranking but the fact that "a public candidate can for the first time effectively talk to each individual voter privately in their own home and tell them exactly what they want to hear." Moreover, the disappearing message feature of Snapchat and Instagram has made tracking of all advertisements harder. "Lots of data, and systems which can react properly to the interests, beliefs, and feelings of different people lead to a world where technology and brands tell us exactly what we want to hear in a way that can't be tracked or audited." Far from worrying about "echo chambers," we need to worry about "personal, private and disappearing messaging that can be powerful but can't be broadly traced or audited."[137]

For Lessin, the Internet is no more likely to produce two opposing ideological camps than it was to produce a global community. Polarization is just a phase on the way to a much more complete atomization as "the pressure built into the internet's DNA, accelerated by things like artificial intelligence, … threatens to undermine our ability to understand one another [and] see reality the same way." Lessin poses a series of difficult questions:

> Is Anonymity a Feature to Be Protected, or a Bug to Be Quashed? (Looks like the latter.)
>
> Should Anyone Be Able to Reach Everyone? (Not clear. Who gets through the email gateways?)
>
> Is Money a Form of Speech? (If extreme views are stickier than moderate ones, they are cheaper to disseminate online.)
>
> Who Decides What Algorithms and Human Policies Control Our View of Reality? (Answer: network platforms, i.e. private corporations in conjunction with ~200 governments. Means rising costs.)
>
> Will We Tolerate Unregulated Escape Hatches for Free Speech? (Answer: if not, the global village blows up.)[138]

His worry is that the Internet is a kind of fission bomb. First it shrank the world into a global village; now it threatens to start "a chain reaction that, just like a nuclear weapon, will cause the world to violently explode." The problem is that, with the consolidation of the public sphere in the network platforms and the rapid collapse of the private sphere, we are losing the ability to let off steam because nothing is any longer "off the record."[139]

Nothing is off the record, but no record is wholly reliable. As technological advances make it easier to create believable fake photos, videos, and audio recordings, we may soon find ourselves unmoored from the modes of rational verification we have evolved since the Scientific Revolution. The network platforms flourished in an anarchic Internet because they appeared to create trust. But our trust in them was misplaced, because their business models incentivized them to send us fake news from fake accounts if it was sufficiently engaging. Our biggest vulnerability turns out to be that while our appetite for entertainment—branded as fiction—is finite, our appetite for "edutainment"—entertaining content branded as educational or informative—seems almost infinite. As a result, when we go online, we find ourselves a "world where we are connected to less trustworthy people and organizations than we ever would have been in the physical world. The value of the human network has almost reversed itself—from being an incredibly good way to 'clean' content and refine information to a system that packages together good and bad information and leaves the two indistinguishable."[140] Worse, we are no longer able to distinguish when we are connected to another human being or to a bot or other device. Those who talk of the "Internet of Things" seem to believe we shall be able to tell the 8 billion people from the 30 billion devices that will be connected to the Internet by 2020. That will be hard when many of those "things" will appear to be people.[141]

To some commentators, the future looks exceedingly bleak. The advance of artificial intelligence, they argue, dooms mankind to a new totalitarianism, rendering liberal democracy and free-market economics "obsolete." According to Yuval Noah Harari, "once we begin to count on AI to decide what to study, where to work, and whom to date or even marry, human life will cease to be a drama of decision making. … We are now creating tame humans who produce enormous amounts of data and function as efficient chips in a huge data-processing mechanism." We shall soon be to data what cows are to milk.[142] Lessin is only slightly more optimistic. He envisages two viable futures (and one untenable one that tries to lie between the two extremes). In his authoritarian future, national governments end up regulating all private and public speech, and the Internet effectively breaks up. In such a scenario, all countries end up going down the Chinese road. However, there is an alternative world—of "technologically guaranteed free speech," based on blockchain technology. In this future, each person would have a public and private key-pair, as with today's cryptocurrencies such as Bitcoin. But this would be a blockchain-based system of communication and record-keeping:

> Everyone could be identified by their public key. Anyone could write public or private "claims" about any other key (signing the information with their key and the keys of intended recipients). And, critically, all the "speech" shared over time could be stored in an append-only distributed database … that would make it impossible for anyone to rewrite history or limit speech. … anyone could say anything to anyone else. Anyone could remember anything they wanted. No one's identity could be blocked or eliminated. We would have true freedom of speech, memory and identity.

Yet this would be no libertarian paradise because "people could have as many different identities as they want … fraudulent identities spewing falsehood would abound. … there would be a huge amount of junk, lies and attempted manipulation … the most abhorrent speech [would be] unstoppable. Bullying would abound."[143] Subversive interventions by foreign governments would presumably abound, too.

None of these futures is very appealing, but each becomes more plausible the longer we leave the Internet in its present state, with generally incompetent or repressive regulation by governments or bogus self-regulation by the network platforms. I began this essay be identifying five distinct problems:

- Monopoly: the sheer scale and market share (national and global) of a few network platforms.
- Pollution of national and international discourse: the tendency of the network platforms to disseminate fake news and extreme views, with destabilizing effects on democracy.
- Harm to individuals with impunity: the network platform's immunity from liability for harms arising from their violations of privacy, spread of defamatory content and addictive applications.
- Vulnerability to outside manipulation: the ease with which the network platforms can be used as assets, if not weapons, by hostile state and non-state actors.
- Censorship: the tendency of the network platforms to respond to public criticism by restricting free speech by modifying their terms of service and the guidelines the issue to their rapidly multiplying content monitors.

The events of 2016 have made it clear that the status quo is indefensible. The network platforms are too powerful and too reckless, too addicted to making us addicted, to be left to regulate themselves. We do not wish to follow China down the road to the surveillance state. It is doubtful that we should passively follow Europe as it seeks, at one and the same time, to live off the network platforms, by taxing and fining them, and to delegate the power of public censorship to them. It would be far better if the United States, before succumbing to the terrors of science fictional futures, sought to learn from history.

The network platforms are new, but they have recognizable antecedents. They combine the pricing power of the big railroad companies with the monopolistic tendencies of Standard Oil; the political influence of William Randolph Hearst with the technological leadership of AT&T and IBM; the convenience of Microsoft with the addictiveness and harmfulness of cigarettes. That we should trust them not to "be evil" is out of the question. We should always have expected them both to move fast and to break things that we did not want to have broken. It will not be enough to tilt the balance of power a little in the favor of the ISPs. Antitrust alone is not likely to be a swift enough remedy, and it would be a chronic optimist who believed a federal government agency would solve all the problems described above. (The recent record of financial regulators offers almost no encouragement.) Far better, surely, to oblige the network platforms to deal directly with citizens and competitors through the lawcourts, by removing the privilege on which they have grown so fat—the exemption from liability for content, established by Section 230—and at the same time to impose on them an obligation that is appropriate for entities that now so dominate the public sphere: a requirement to conduct themselves (as Harvard University conducts itself) as if the First Amendment applies to them.

This combination of remedies seems consonant with the traditions of American government. It recognizes the difficulty of reviving antitrust law in the age of network economics. It avoids the pitfall of entrusting too much power to bureaucrats. Yet by simultaneously increasing the network platforms' costs and reducing their power to distort the democratic process, the measures proposed here would, if implemented, stand a reasonable chance of reining in the new Rockefellers—not to mention the Caesars—of Silicon Valley. These reforms might also ensure that, like David Rockefeller in the 1940s, the platforms put themselves at the disposal of the government that is their ultimate protector, if called upon to serve.

[1] This is a revised version of a paper presented at the Hoover Institution symposium on "The Information Challenge to Democracy," on November 13, 2018. In addition to participants in the symposium, I would like to thank Renee DiResta, Alex Feerst, Jack Goldsmith, Sam Lessin, Carlos Medeiros, and Eric Schmidt for their comments and suggestions.

[2] Joann S. Lublin, "David Rockefeller's Rolodex Was the Stuff of Legend. Here's a First Peek," *Wall Street Journal*, December 5, 2017.

[3] Alexis C. Madrigal, "The Education of Mark Zuckerberg," *Atlantic*, November 20, 2017.

[4] Greg Ip, "The Unintended Consequences of the 'Free' Internet," Wall Street Journal, November 14, 2018.

[5] Scott Galloway, "Silicon Valley's Tax-Avoiding, Job-Killing, Soul-Sucking Machine," *Esquire*, February 8, 2018

[6] Evan Osnos, "Can Mark Zuckerberg Fix Facebook Before It Breaks Democracy?" *New Yorker*, September 10, 2018.

[7] David Leonhardt, "Companies Avoiding Taxes," *New York Times*, October 18, 2016.

[8] Stacy Cowley, "Google CEO on Censoring: 'We Did an Evil Scale," InfoWorld, January 27, 2006.

[9] Adam J. White, "Google.gov," *New Atlantis*, Spring 2018.

[10] Allum Bokhari, "LEAKED VIDEO: Google Leadership's Dismayed Reaction to Trump Election," Breitbart, September 12, 2018.

[11] The allusion is to Eli Pariser, *The Filter Bubble* (2011).

[12] White, "Google.gov."

[13] David Streitfeld, "'The Internet is Broken.' ev Is Trying to Salvage It," New York Times, May 20, 2017.

[14] Even the term "fake news" originated in Silicon Valley. Sherryl Attkisson of the Google-funded non-profit First Draft first used the term on September 13, 2016. President Obama used the term himself soon after. But it backfired on its creators when Donald Trump accused the established liberal media of themselves purveying fake news.

[15] Tim Wu, *The Attention Merchants: The Epic Scramble to Get Inside Our Heads* (Knopf, 2016).

[16] eMarketer, September 2017.

[17] According to Verto Analytics, Facebook users in the United States spent 18 hours, 24 minutes on Facebook in October 2017, compared with 32 hours, 43 minutes, recorded by Verto in October 2016.

[18] Adam D. I. Kramer and Cindy K. Chung, "Dimensions of Self-Expression in Facebook Status Updates," Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (2011).

[19] Elisa Shearer and Jeffrey Gottfried, "News Across Social Media Platforms 2017," Pew Research Center, September 6, 2017.

[20] Jon Gingerich, "Google Overtakes Facebook for Referral Traffic," O'Dwyer's, June 8, 2018.

[21] Jack Nikas, "How YouTube Drives People to the Internet's Darkest Corners," *Wall Street Journal*, February 7, 2018.

[22] Title V is also known as the Communications and Decency Act (CDA), for anti-obscenity provisions that have since been struck down by the Supreme Court.

[23] Ashley Gold and Joanna Plucinska, "U.S., Europe Threaten Tech Industry's Cherished Legal Shield," Politico, October 8, 2018.

[24] James Pethokoukis, "Should Big Tech be Held More Liable for the Content on their Platforms? An AEIdeas Online Symposium," March 20, 2018. See also Ashley Gold, "Tech's Next Big Battle: Protecting Immunity From Content Lawsuits," The Information, Jan. 11, 2019.

[25] Kate Klonick, "The New Governors: The People, Rules, and Processes Governing Online Speech," *Harvard Law Review* (2018). See also Jeff Kosseff, *The Twenty-Six Words That Created the Internet* (Cornell University Press, forthcoming).

[26] Klonick, "New Governors."

[27] Danielle Citron and Quinta Jurecic, "Platform Justice Content Moderation at an Inflection Point," Aegis Series Paper No. 1811 (Stanford: Hoover Institution, 2018).

[28] Osnos, "Can Mark Zuckerberg Fix Facebook?"

[29] Gabriel J.X. Dance, Michael LaForgia and Nicholas Confessore, "As Facebook Raised a Privacy Wall, It Carved an Opening for Tech Giants," *New York Times*, Dec. 18, 2018; Rana Foroohar, "Facebook Has Put Growth Ahead of Governance For Too Long," *Financial Times*, Dec. 23, 2019; Alexis Madrigal, "Facebook Didn't Sell Your Data; It Gave It Away," *Atlantic*, Dec. 19, 2018.

[30] Gary King, Benjamin Schneer and Ariel White, "How the News Media Activate Public Expression and Influence National Agendas," *Science*, 258 (2017), 776-780.

[31] Osnos, "Can Mark Zuckerberg Fix Facebook?"

[32] Soroush Vosoughi, Deb Roy and Sinan Aral, "The Spread of True and False News Online," *Science*, 359, 6380 (2018), 1146-1151. See also David M. J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, Jonathan L. Zittrain, "The Science of Fake News: Addressing Fake News Requires a Multidisciplinary Effort," *Science*, 359, 6380 (March 2018), 1094-1096.

[33] Renee DiResta, "There are Bots. Look Around," *Ribbon Farm*, May 23, 2017.

[34] Matthew Hindman and Vlad Barash, "Disinformation, and Influence Campaigns on Twitter," Knight Foundation (October 2018).

[35] Caleb Jones, "Visualizing Polarization in Political Blogs," AllThingsGraphed, October 9,2014.

[36] Pascal Wallisch, "Illumination Assumptions Account for Individual Differences in the Perceptual Interpretation of a Profoundly Ambiguous Stimulus in The Color Domain: 'The Dress'," *Journal of Vision*, 17 (4), 5 (2017), 1-14.

[37] M. D. Conover, J. Ratkiewicz, M. Francisco, B. Gonc¸alves, A. Flammini, F. Menczer, , "Political Polarization on Twitter," Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (2011).

[38] William J. Brady, Julian A. Wills, John T. Jost, Joshua A. Tucker and Jay J. Van Bavel, "Emotion Shapes the Diffusion of Moralized Content in Social Networks," www.pnas.org/cgi/doi/10.1073/pnas.1618923114 (2017),

[39] Daniel Preotiuc-Pietro, Ye Liu, Daniel Hopkins, and Lyle Ungar, "Beyond Binary Labels: Political Ideology Prediction of Twitter Users," 10.18653/v1/P17-1068 (2017), 729-740.

[40] Adam Hughes and Onyi Lam, "Highly Ideological Members of Congress Have More Facebook Followers Than Moderates Do," Pew Research Center, August 21, 2017.

[41] Levi Boxell, Matthew Gentzkow and Jesse M. Shapiro, "Is the Internet Causing Political Polarization? Evidence from Demographics," working paper (March 2017).

[42] Zeynep Tufekci, "YouTube, the Great Radicalizer," New York Times, March 10, 2018.

[43] Paul P. Murphy, Kaya Yurieff and Gianluca Mezzofiore, "Exclusive: YouTube Ran Ads from Hundreds [of] Brands on Extremist Channels, CNN, April 20, 2018.

[44] Marlon Mooijman, Joe Hoover, Ying Lin, Heng Ji and Morteza Dehghani, "Moralization in Social Networks and The Emergence of Violence During Protests," *Nature Human Behaviour*, 2 (June 2018), 389-396.

[45] Stefano DellaVigna and Ethan Kaplan, "The Fox News Effect: Media Bias and Voting," *Quarterly Journal of Economics*, 122, 3 (2007), 1187-1234.

[46] Sasha Issenberg, "How Obama's Team Used Big Data to Rally Voters," MIT Technology Review, December 19, 2012: https://www.technologyreview.com/s/509026/how-obamas-team-used-big-data-to-rally-voters/.

[47] Robert Epstein and Ronald E. Robertson, "The Search Engine Manipulation Effect (SEME) and its Possible Impact on the Outcomes of Elections," *PNAS*, August 4, 2015: www.pnas.org/cgi/doi/10.1073/pnas.141982811.

[48] Robert Epstein, "Taming Big Tech: The Case for Monitoring," Hacker Noon, May 17, 2018.

[49] Robert Epstein, Roger Mohr, Jr. and Jeremy Martinez, "The Search Suggestion Effect (SSE): How Search Suggestions Can Be Used to Shift Opinions and Voting Preferences Dramatically and Without People's Awareness," Paper presented at the 98th annual meeting of the Western Psychological Association, Portland, OR, April 26, 2018.

[50] Matthew Rosenberg, "Professor Apologizes for Helping Cambridge Analytica Harvest Facebook Data," *New York Times*, April 22, 2018.

[51] Craig Timberberg, Tony Romm and Elizabeth Dwoskin, "Facebook Said the Personal Data of Most of its 2 Billion Users Has Been Collected and Shared with Outsiders," *Washington Post*, April 4, 2018.

[52] Matthew Rosenberg, Nicholas Confessore and Carole Cadwalladr, "How Trump Consultants Exploited the Facebook Data of Millions," *New York Times*, March 17, 2018.

[53] Carole Cadwalladr, "Why Britain Needs its own Mueller," *New York Review of Books*, Nov. 16, 2018.

[54] Yuriy Gorodnichenko, Tho Pham and Oleksandr Talavera, "Social Media, Sentiment and Public Opinions: Evidence from #Brexit and #USElection," NBER Working Paper No. 24631 (May 2018).

[55] Vijaya Gadde and Yoel Roth, "Enabling Further Research of Information Operations on Twitter," October 17, 2018.

[56] Kate Starbird, "Information Wars: A Window into the Alternative Media Ecosystem," Medium [?], March 14, 2017.

[57] Renee DiResta, Kris Shaffer, Becky Ruppel, David Sullivan, Robert Matney, Ryan Fox, Jonathan Albright, Ben Johnson, *The Tactics and Tropes of the Internet Research Agency* (New Knowledge, 2018).

[58] Hindman and Barash, "Disinformation"

[59] Jane Mayer, "How Russia Helped Swing the Election for Trump," *New Yorker*, October 1, 2018.

[60] Kate Starbird, "The Trolls Within: How Russian Information Operations Infiltrated Online Communities in 2016," Medium, October 20, 2018.

[61] Hindman and Barash, "Disinformation."

[62] Gadde and Roth, "Enabling Further Research."

[63] Lazer et al., "Science of Fake News," 1095. See also https://www.nytimes.com/2018/01/02/health/fakenewsconservativeliberal.html.

[64] Brendan Nyhan, Andrew Guess and Jason Reifler, "Selective Exposure to Misinformation: Evidence from the Consumption of Fake News During the 2016 U.S. Presidential Campaign," working paper, January 9, 2018.

[65] Jorg L. Spenkuch and David Toniatti, "Political Advertising and Election Results," working paper (March 2018): https://www.kellogg.northwestern.edu/faculty/spenkuch/research/advertising.pdf.

[66] Yuriy Gorodnichenko, Tho Pham and Oleksandr Talavera, "Social Media, Sentiment and Public Opinions: Evidence From #Brexit and #Uselection," NBER Working Paper 24631 (May 2018).

[67] Robert Epstein, Roger Mohr, Jr. and Jeremy Martinez, "The Search Suggestion Effect (SSE): How Search Suggestions Can Be Used to Shift Opinions and Voting Preferences Dramatically and Without People's Awareness," Paper presented at the 98th annual meeting of the Western Psychological Association, Portland, OR, April 26, 2018. Google denies the charge: See Gillian Tett, "Facebook or Google—Which Should Worry Us More?", *Financial Times*, May 2, 2018.

[68] Robert Epstein, "Taming Big Tech: The Case for Monitoring," Medium, May 17, 2018.

[69] Osnos, "Can Mark Zuckerberg Fix Facebook?"

[70] Connie Loizof, "'When You Spend $100 Million on Social Media,' It Comes with Help, Says Trump Strategist," TechCrunch, November 2017.

[71] Robert McMillan, "Facebook Logs Text, Call Histories for Some Android Users," *Wall Street Journal*, March 26, 2018.

[72] Frank Bajak, "AP FACT CHECK: Does Facebook Collect Data on Non-members?" *Washington Post*, April 11, 2018.

[73] Brian X. Chen, "Google's File on Me Was Huge. Here's Why It Wasn't as Creepy as My Facebook Data," *New York Times*, May 16, 2018.

[74] Federica Liberini, Michela Redoano, Antonio Russo, Ángel Cuevas Rumin and Ruben Cuevas Rumin, "Politics in the Facebook Era: How Reading Political Ads on Facebook Affects our Voting Behaviour (… and helped Trump to win the presidential election)," VOX, CEPR Policy Portal, November 7, 2018.

[75] Niall Ferguson, *The Square and the Tower: Networks and Power from the Freemasons to Facebook* (New York: Penguin, 2018).

[76] George Soros, speech at the Davos World Economic Forum, January 25, 2017.

[77] Sandy Parakilas, "Facebook Won't Protect Your Privacy," *New York Times*, November 19, 2017.

[78] "Facebook's Trust Problem," eMarketer, January 11, 2018.

[79] Knight Foundation/Gallup, "American Views: Trust, Media and Democracy" (2017), 10.

[80] Kim Hart, "Exclusive: Public Wants Big Tech Regulated," *Axios*, February 2018.

[81] Harvard Kennedy School Institute of Politics, Survey of Young Americans' Attitudes toward Politics and Public Service, 35th Edition: March 8-March 25, 2018.

[82] Silicon Valley, We Have a Problem," *Economist*, January 18, 2018.

[83] Deborah M. Gordon, "Local Links Run the World," *Aeon*, February 1, 2018.

[84] Holly B. Shakya and Nicholas A. Christakis, "The More You Use Facebook, the Worse You Feel," *Harvard Business Review*, April 10, 2017.

[85] David Ginsberg and Moira Burke, "Hard Questions: Is Spending Time on Social Media Bad for Us?" Facebook, December 15, 2017.

[86] James Bridle, Something is Wrong on the Internet," Medium, November 6, 2017.

[87] Franklin Foer, *World Without Mind: The Existential Threat of Big Tech*.

[88] Jennifer King, "Are Young Adults Growing Tired of Constant Social Connectivity?" eMarketer, March 17, 2018.

[89] For a critique, see Dominic Cummings, "On Referendum #24L: Fake News from the Fake News Committee, Carole, and a Rematch Against the Public," July 27, 2018.

[90] Hannes Grassegger, "The Button: Did Facebook's "Voter Button" Influence Elections in Iceland?" *Das Magazin* (2018).

[91] Osnos, "Can Mark Zuckerberg Fix Facebook?"

[92] Paul Mozur, "Myanmar's Military Said to Be Behind Facebook Campaign That Fueled Genocide," *New York Times*, October 15, 2018.

[93] Max Read, Does Even Mark Zuckerberg Know What Facebook Is?" *New York Magazine*, October 2, 2017.

[94] Based on Danielle Keats Citron, "Extremist Speech, Compelled Conformity, and Censorship Creep," forthcoming in the *Notre Dame Law Review*.

[95] John Gapper, "Alibaba's Social Credit Rating is a Risky Game," *Financial Times*, February 21, 2018.

[96] Lisa Linn and Josh Chin, "China's Tech Giants Have a Second Job: Helping Beijing Spy on Its People," *Wall Street Journal*, November 30, 2017.

[97] Freedom House, *Freedom on the Net 2017: Manipulating Social Media to Undermine Democracy* (Washington, DC: Freedom House, 2018).

[98] Sheera Frenkel, Nicholas Confessore, Cecilia Kang, Matthew Rosenberg and Jack Nicas, "Delay, Deny and Deflect: How Facebook's Leaders Fought Through Crisis," *New York Times*, Nov. 14, 2018.

[99] Charlie Warzel, "How People Inside Facebook Are Reacting To The Company's Election Crisis," BuzzFeed, October 20, 2017.

[100] Ibid.

[101] Max Read, "The Decline and Fall of the Zuckerberg Empire," *New York Magazine*, November 20, 2018; Deepa Seetharaman, "With Facebook at 'War,' Zuckerberg Adopts More Aggressive Style," *Wall Street Journal*, November 19, 2018.

[102] Alexis C. Madrigal, "When the Facebook Traffic Goes Away," *The Atlantic*, October 24, 2017.

[103] Kevin Roose and Sheera Frenkel, "Mark Zuckerberg's Reckoning: 'This Is a Major Trust Issue,'" *New York Times*, March 21, 2018.

[104] "Mark Zuckerberg in His Own Words," CNN, March 21, 2018.

[105] Hannah Kuchler, "Facebook Moves to Make Privacy Policies More Transparent," *Financial Times*, March 29, 2018.

[106] Facebook post, September 8, 2018.

[107] Sarah Provan, "Facebook Removes Hundreds of 'Fake' Accounts Linked to Russia," *Financial Times*, January 17, 2019.

[108] White, "Google.gov."

[109] "Mark Zuckerberg in His Own Words," CNN, March 21, 2018.

[110] Ezra Klein, "Mark Zuckerberg on Facebook's Hardest Year, and What Comes Next," Bloomberg, April 2, 2018.

[111] Frenkel et al., "Delay, Deny and Deflect."

[112] See e.g., Timothy Garton Ash, Robert Gorwa, Danaë Metaxa, *GLASNOST! Nine Ways Facebook Can Make Itself a Better Forum for Free Speech and Democracy* (Oxford, 2019).

[113] In 2012 the FTC fined Google $22.5 million for violating the terms of another consent decree. It remains to be seen if Facebook will be penalized more or less harshly. At the time of writing, Facebook is also under investigation by the Federal Bureau of Investigation, the Securities and Exchange Commission and the Department of Justice. Also investigating Facebook are the attorney generals of New York, Massachusetts and Connecticut.

[114] Noah Smith, "The Battle Over Monopoly Power Is Just Beginning," Bloomberg, December 10, 2018; April Glaser, "Antitrust in the House," Slate, January 16, 2019; Lauren Feiner, "Attorney General Pick Barr Hints He Would Look at Antitrust in Tech," CNBC, January 15, 2019.

[115] Jamie M. Fly and Laura Rosenberger, "How Silicon Valley Can Protect U.S. Democracy," *Foreign Affairs*, February 22, 2018; Kiran Stacy, "How Washington Plans to Regulate Big Tech," *Financial Times*, January 6, 2019.

[116] Jonathan Swan, "Trump Hates Amazon, Not Facebook," Axios, March 28, 2018.

[117] Jonah Engel Bromwich, "YouTube Cracks Down on Far-RightVideos as Conspiracy Theories Spread," *New York Times*, March 3, 2018.

[118] Ezra Klein, "The Controversy Over Mark Zuckerberg's Comments on Holocaust Denial, Explained," Vox, July 20, 2018.

[119] Brad Parscale, "Big Tech Is Meddling with Free Speech… and Elections," Breitbart, October 23, 2018.

[120] Jonathan Albright, "The 2018 Facebook Midterms, Part I: Recursive Ad-ccountability," Medium, November 5, 2018; "The 2018 Facebook Midterms, Part II: Shadow Organizing," November 5, 2018; "The 2018 Facebook Midterms, Part III: Granular Enforcement," November 6, 2018.

[121] Renee DiResta, "Free Speech in the Age of Algorithmic Megaphones," Wired, October 12, 2018.

[122] Michelle Malkin, "The Authoritarianism of Silicon Valley's Tech Titans," *National Review*, November 28, 2018; Jeremy Carl, "Why We Need Anti-Censorship Legislation for Social Media," *The Federalist* (n.d., 2018).

[123] Robert Epstein, "Another Way Google Manipulates Votes Without Us Knowing: A 'Go Vote' Reminder Is Not What You Think It Is," unpublished paper (2019).

[124] Scott Galloway, *The Four: The Hidden DNA of Amazon, Apple, Facebook and Google* (New York: Portfolio/Penguin, 2017)

[125] See e.g., Jonathan Tepper with Denise Hearn, *The Myth of Capitalism: Monopolies and the Death of Competition* (forthcoming).

[126] Lina M. Khan, "Amazon's Antitrust Paradox," *Yale Law Review*, 126, 710 (2017), 710-805.

[127] Roger McNamee, "How to Fix Facebook—Before It Fixes Us," *Washington Monthly*, (January-March 2018).

[128] This is a judgment supported by Carlotta Perez's more recent work on technological development, Carl H. Fulda's study of inland water-carriers, aircraft, and motor carriers, and conservative economist and jurist Richard A. Posner's work on natural monopoly regulation.

[129] Mark Warner, "Potential Policy Proposals for Regulation of Social media and Technology Firms," Draft White paper (2018).

[130] U.S. District Court, Southern District of New York, ruling in May 23, 2018.

[131] Danielle Citron and Quinta Jurecic, "Platform Justice Content Moderation at an Inflection Point," Aegis Series Paper No. 1811 (Stanford: Hoover Institution, 2018).

[132] Jonathan Zittrain, "CDA 230 Then and Now: Does Intermediary Immunity Keep the Rest of Us Healthy?" August 31, 2018: https://blogs.harvard.edu/jzwrites/2018/08/31/cda-230-then-and-now/.

[133] By a factor of between 25 and 30 in each case, according to one experienced lawyer I consulted. But of course there would be many more cases and perhaps also class-action lawsuits. Increased costs would mean more settlements out of court. I have been amused by how many people working in Silicon Valley regard this is an insuperable argument against what I am proposing.

[134] Osnos, "Can Mark Zuckerberg Fix Facebook?"

[135] P.W. Singer and Emerson T. Brooking, *LikeWar: The Weaponization of Social Media* (Eamon Dolan/Houghton Miffin Harcourt, 2018). See also Neil Verrall and David Mason, "The Taming of the Shrewd," *RUSI Journal* (2018), 1-9, DOI: 10.1080/03071847.2018.1445169.

[136] Larry Diamond and Orville Schell (co-chairs), *Chinese Influence and American Interests: Promoting Constructive Vigilance: Report of the Working Group on Chinese Influence Activities in the United States* (Stanford: Hoover Institution, 2018).

[137] Sam Lessin, Free Speech and Democracy in the Age of Micro-Targeting," The Information, December 5, 2016.

[138] Sam Lessin, "The Tower of Babel: Five Challenges of the Modern Internet," The Information, November 6, 2017.

[139] Ibid.

[140] Sam Lessin, "How the Internet Broke and What to Do About It," The Information, September 4, 2018.

[141] Philip Howard, *Pax Technica How the Internet of Things May Set Us Free Or Lock Us Up* (New Haven / London: Yale University Press, 2015).

[142] Yuval Noah Harari, "Why Technology Favors Tyranny," *The Atlantic*, October 2018.

[143] Lessin, "How the Internet Broke."

Niall Ferguson is the Milbank Family Senior Fellow at the Hoover Institution and a senior fellow of the Center for European Studies at Harvard University, where he served for twelve years as the Laurence A. Tisch Professor of History. He the author of fifteen books, most recently The Square and the Tower.

Protecting Democracy in an Era of Cyber Information War

By Joseph S. Nye, Harvard University

The early years of the Internet were marked by a libertarian optimism about its decentralizing and democratizing effects.[1] Information would be widely available and undercut the monopolies of authoritarian governments. Big Brother would be defeated. President Clinton believed that China would liberalize and that Communist Party efforts to control the Internet were like trying to "nail jello to the wall."[2] The Bush and Obama administrations shared this optimism and promoted an Internet Freedom Agenda that included subsidies and technologies to assist dissidents in authoritarian states to communicate.

Today, in the face of successful Chinese control of what citizens can see and say on the Internet and Russian use of the Internet to interfere in the 2016 American election, the United States (and allied democracies) find themselves on the defensive. The expected asymmetries seem to have been reversed. Autocracies are able to protect themselves by controlling information flows, while the openness of democracies creates vulnerabilities that autocracies can exploit via information warfare. Ironically, one cause of the vulnerabilities has been the rise of social media and mobile devices in which American companies have been the global leaders. Citizens voluntarily carry Big Brother and his relatives in their pockets. Along with big data and artificial intelligence, technology has made the problem of defending democracy from information warfare far more complicated than foreseen two decades ago. And while rule of law, trust, truth, and openness make democracies asymmetrically vulnerable, they are also critical values to defend. Any policy to defend against cyber information war must start with the Hippocratic oath: first, do no harm.

Information Warfare: What's New and What's Not

The use of information as an instrument of conflict and manipulation in international politics has a long history. Britain manipulated information to move American opinion in the direction of war with Germany both in 1917 and 1941. The United States and the Soviet Union both used broadcasts, covert organizations, and funds to interfere in foreign elections during the Cold War.[3] And more narrowly, in battlefield situations in Iraq or in the campaign against ISIS, information was an important tool. In recent years, Russia's hybrid war against Ukraine has encompassed both cyber attacks and manipulation of information. Information operations are a critical component of modern warfare.[4]

Russia has used propaganda to express preferences for candidates in American elections since at least 1964, but new technologies have amplified their impact enormously.[5] According to former CIA Director Michael Hayden, Russian interference in the 2016 election was "the most successful covert influence campaign in recorded history."[6] For example, Russian operatives used Facebook to publicize 129 staged events, drawing attention of 340,000 users; 10 million people saw ads paid for by Russian accounts; and 126 million Americans saw posts by 470 accounts affiliated with the Russian Internet Research Agency.[7] A study by Twitter reported that 50,000 Russia-linked accounts were automated and tweeted election related content.[8] Reports released by the Senate Intelligence Committee estimate that the Russian campaign reached not only the 126 million people on Facebook but another 20 million more on Instagram.[9] Some Russian messages were crafted to support particular candidates while others were designed to create a general sense of chaos. Still others were micro-targeted to suppress voting by particular demographic groups such as African-Americans or younger voters. While skeptics argue that Russian efforts were a small percentage of the total content on the Internet, "for sub-groups of targeted Americans, the messaging was perhaps ubiquitous."[10]

Before the Internet, such operations involved costly training and movement of spies across borders, establishment of foreign bank accounts, and transfers of cash. Now similar effects can be accomplished remotely at much lower cost. It is much easier to send electrons across borders than human agents. Ransoming a failed spy can be costly, but if no one clicks on a phishing e mail, it is simple, deniable, and virtually free to send another. In 1983, when the KGB seeded the rumor that AIDS was the product of U.S. government experiments with biological weapons, the rumor started with an anonymous letter to a small New Delhi newspaper and then was propagated globally but slowly over several years by widespread reproduction and constant repetition in conventional media. It took four years to reach full fruition. [11] In 2016, an updated version of the same technique was used to create "Pizzagate," the bizarre rumor that Hillary Clinton's campaign manager ran a child sex ring in a Washington restaurant. It spread instantly on the Internet. What's new is not the basic model; it's the speed with which such disinformation can spread and the low cost of spreading it.

With its armies of paid trolls and botnets, along with outlets such as Russia Today (RT) and Sputnik, Russian intelligence, after hacking into the e-mails of the Democratic National Committee and senior Clinton campaign officials, could distract and disrupt news cycles week after week without setting foot in the United States. And it could also count on the witting and unwitting help of organizations like Wikileaks. Russian messages aimed at priming, framing, agenda setting, and contagion were accelerated by U.S. media that were too quick and unreflective in using the Russian phrasing and frames.[12] American voters are subject to many influences, and there were many potential causes of the narrow outcome of the 2016 election. It is far too simple just to blame manipulation of social media. As social scientists say, the outcome was "overdetermined." But whatever its effects on the particular election outcome, Russia was able to accomplish its deeper goal of sowing disruption and discrediting the democratic model. It successfully undercut American soft power.

Soft Power and Sharp Power

Many Russian books and articles claim that "the death blow to the Soviet Union came not from NATO conventional forces but from an imperialist information war that Russia lost."[13] From the Kremlin's perspective, color revolutions in neighboring countries and the Arab Spring uprisings in 2011 were examples of the United States using soft power as a new form of hybrid warfare. "Authoritarian governments do not just fear that their citizens will use the Internet to organize and rebel; they also believe that democracies use the Internet to advance pro-democracy narratives to undermine their regimes."[14] While that may not have been the intention of the Obama Administration in Ukraine, Russia felt it needed to respond. The concept of soft power was incorporated into Russia's 2013 Foreign Policy Concept, and in March 2016, Russian Chief of General Staff Valery Gerasimov stated that since responding to such foreign threats using conventional troops is impossible; they must be counteracted with the same hybrid methods.[15] However, Russian and American views of soft power differ.

Power is the ability to affect others to get what you want, and that can be done through coercion, payment, and attraction. Some think soft power means any action other than military force, but this is wrong. Soft power is the ability to get what you want through attraction and persuasion rather than coercion or payment. While it relies in part on information, it differs from the coercive manipulation of information because it rests on the voluntarism of the subject. The soft power of attraction can be used for offensive purposes, but if the degree of manipulation is so deceptive that it destroys voluntarism, the act becomes coercive and is no longer soft power. This manipulative use of information has recently been dubbed "sharp power."[16] Countries have long spent billions on public diplomacy and broadcasting in a game of competitive attractiveness—the "battle for hearts and minds." Soft-power instruments like the Marshall Plan and the Voice of America helped to determine the outcome of the Cold War through attraction. But the United States also used deceptive sharp power in the form of covert support for publications and political parties.

After the Cold War, Russian elites believed that European Union and NATO enlargement, and Western efforts at democracy promotion, were designed to isolate and threaten Russia. In response, they tried to develop Russian soft power by promoting an ideology of traditionalism, state sovereignty, and national exclusivity. This attracted support in countries like Hungary, where Prime Minister Victor Orbán promoted "illiberal democracy," as well as among the diaspora along Russia's borders, in impoverished countries of Central Asia, and among right-wing populist movements in Western Europe. But Russian soft power was quite limited. What Russia lacks in soft power, however, it has made up with its sharp power manipulation of social media.

In 2007, President Hu Jintao told the 17th Congress of the Chinese Communist Party that China needed to increase its soft power, and it has been spending billions on broadcasting, exchange programs, and Confucius Institutes to teach Chinese language and culture.[17] In addition China's impressive economic performance has added to its attractiveness. David Shambaugh estimates that China spends $10 billion a year on its soft power instruments, but it has earned a modest return on its investment. The "Soft Power 30" index ranks China 27th (and Russia 28th) out of 30 countries assessed, far below the United States and European democracies.[18] But China also goes beyond soft power and tries to exercise discourse control and export censorship beyond its borders by manipulation of visas, threatening loss of access to its markets, control of its information companies, covert broadcasting, and payments to foreign groups and politicians. While China has not tried to disrupt the American political process to the extent that Russia has, it has used cyber and other means to intervene in politics in other countries. As Eric Rosenbach and Katherine Mansted point out, democratic civil society actors are often "the primary agents for much of the soft power appeal of the U.S. system of government. This dynamic means that authoritarian states do not just view control of their information environments as a domestic matter; they increasingly believe that offensive action might be required to counter what they perceive as foreign information incursions."[19]

Other authoritarian countries such as North Korea and Iran manipulate information to undercut American power, but Russia and China are the most important. Russian interference in European democracies' domestic politics is designed to reduce the attractiveness of NATO, the embodiment of Western hard power, which Russia views as a threat. In the 19th Century, the outcome of contests for mastery of Europe depended primarily on whose army won; today, it also depends on whose story wins. As Singer and Brooking argue, "these new wars are not won by missiles and bombs, but by those able to shape the story lines that frame our understanding…"[20] In addition to formal public diplomacy organizations like Russia Today and Sputnik, Russian intelligence units and their proxies generate false information that can later be circulated and legitimated as if it were true. And it is easy and cheap to send such disinformation across borders.

Authoritarian sharp power has disrupted Western democratic processes and tarnished the brands of democratic countries, but it has done little to enhance the soft power of its perpetrators—and in some cases it has done the opposite. For Russia, which is focused on playing a spoiler role in international politics, that could be an acceptable cost. China, however, is a rising power that requires the soft power of attraction to achieve its objectives as well as the coercive sharp power of disruption and censorship. These two goals are hard to combine. In Australia, for example, public approval of China was growing until the revelation of its use of sharp power tools, including meddling in Australian politics, set it back considerably. In other words, Chinese deceptive sharp power undercut its soft power.

Although sharp power and soft power work in very different ways—attraction vs. coercion—the distinction between them can sometimes be hard to discern in particular instances and that complicates the response to authoritarian sharp power. Attraction and persuasion involve choices about how to frame information. When that framing shades into deception, which limits the subject's voluntary choices, it crosses the line into coercion. Openness and limits on deliberate deception distinguish soft from sharp power. When RT or Xinhua broadcast openly in other countries, they are employing soft power. Similarly, properly labeled advertising in American media are legitimate exercises of soft power. If their messages are too blatantly propagandistic, they will not attract support and thus fail to produce soft power, but democracies can deal with open information. When the authoritarian states covertly back radio stations in other countries, or secretly promote news on social media, that deception crosses the line into sharp power. Transparency and proper disclosure is necessary to preserve the principle of voluntarism that is essential to soft power.

As democracies respond to sharp power, we have to be careful not to undercut our own soft power by imitating the authoritarian model. Much of American soft power comes from our civil society—Hollywood, universities, and foundations more than from official public diplomacy efforts—and closing down access or ending openness would undercut our crucial asset. Authoritarian countries such as China and Russia have trouble generating their own soft power precisely because of their unwillingness to free the potential talents in their civil societies—witness Chinese censorship of its film industry or the harassment of the artist Ai Weiwei which undercut its soft power overseas. Moreover, shutting down legitimate Chinese and Russian soft power tools can be counter-productive. For example, if China and the United States wish to avoid conflict, exchange programs that increase American attraction to China, and vice versa, can be good for both countries. And on transnational challenges which pose a shared threat such as climate change, soft power can help build the trust and networks that make cooperation possible. But the programs have to be open and transparent to pass the test of soft power.

It would be a mistake, therefore, to prohibit Russian and Chinese soft power efforts simply because they sometimes shade into sharp power. Congress has required that RT be registered as a foreign entity, but it would be a mistake to go further and ban its broadcasts. At the same time, it is important to monitor the dividing line carefully. Take the 500 Confucius Institutes and 1,000 Confucius classrooms that China supports in universities and schools around the world to teach Chinese language and culture. Government backing does not mean they are necessarily a sharp power threat. Only when a Confucius Institute crosses the line and tries to infringe on academic freedom (as has occurred in some instances) should it be treated as sharp power intrusion and be closed.

Democracies must also be careful about our own offensive information actions. It may make sense to establish an American "political warfare" capability and strategy in an age of hybrid warfare, but a good strategy must be carefully designed and implemented.[21] Public diplomacy and broadcasting should be public. It would be a mistake to imitate the authoritarians and use major programs of covert information warfare as we did in the Cold War. Such actions will not stay covert for long and when revealed would undercut our soft power as we saw in the 1970s when many CIA covert cultural operations were disclosed. Some argue that in the information struggle against authoritarian systems, democracies should use every weapon available and not worry about nice distinctions between soft and sharp power. However, the two types of power are hard to combine successfully in the long term, and some apparent arrows in the quiver of political warfare may turn out to be boomerangs. In the long term, central manipulation of information can make authoritarian states brittle, and openness may make democracies more resilient.

In the realm of defensive measures, democratic governments must counter the authoritarians' aggressive information warfare techniques (as we shall see below), but openness remains the ultimate defense of liberal societies. The press, academics, civic organizations, government, and the private sector should focus on exposing information warfare techniques, inoculating the public by exposure. Openness is a key source of democracies' ability to attract and persuade. As Henry Farrell and Bruce Schneier point out, information plays a very different role in legitimizing the political order of autocracies than in democracies.[22] Even with the mounting use of sharp power, we have little to fear in open competition with autocracies for soft power. If we succumb to temptation and lower our standards to the level of our authoritarian adversaries, democracies will squander our key advantage.[23]

Technology, New Tools, and Remedies

The authoritarian threat to democracy takes a number of forms ranging from the corruption of election machinery to the manipulation of voters through fake news, through the targeting and destruction of particular candidates, the creation of inauthentic groups to generate or exacerbate conflict, and the creation of chaos and disruption to discredit the democratic model. In the 2016 American presidential election, for example, Russians scanned election systems in at least 22 states; hacked into individual emails and leaked out the contents; and created fake accounts, trolls, and disinformation to disrupt the political process.

Hacking Electoral Systems

The most direct way to corrupt democracy is to manipulate the electoral systems and alter the calculations of voting. This can be accomplished through hacking into voting machines or into the rolls of registered voters. This is a particular problem with older voting machines which do not have a paper backup, but now 80 per cent of Americans vote on machines that incorporate paper ballots or backups. Since 2016, many state voter registration data bases have been hardened against outside attacks. A number of civic organizations have developed programs to alert and train local election officials. State election officials are gaining security clearances to permit access to federal threat information, and in 2018 all 50 states and more than 1000 localities opened a center to exchange data. While hacking election systems remains possible, it is increasingly difficult to rig enough decentralized devices and records to change the outcome of a national election. The Department of Homeland Security has declared that election systems are part of the national critical infrastructure and that makes it now easier to share threat information with state and local officials. Russia does not need to hack into machines to create mistrust about election results. Some of the damage is self-inflicted by American politicians and media, but the press seemed more alert to this danger in the 2018 midterm elections than it had been in 2016. Creating and publicizing a good process is essential. While important, hacking into machines may be the most straightforward and easiest part of the puzzle to solve.[24]

Disseminating Fake News

The term "fake news" has become a political epithet, but as an analytical term, it describes deliberate disinformation that is presented in the format of a conventional news report.[25] Again, the problem is not completely novel. In 1924, Harpers' Magazine published an article about the dangers of "fake news," but today two-thirds of American adults get some of their news from social media where algorithms can easily be gamed for profit or malice. What is different about social media is that they rest on a business model which lends itself to outside manipulation. Many organizations, both domestic and foreign, amateur, criminal, and governmental are skilled at reverse engineering the ways that tech platforms parse information. To give Russia credit, it was one of the first governments to understand how to weaponize social media.

The Internet has flooded the world with information and when people are overwhelmed with the volume of information confronting them, they find it hard to know what to focus on. Attention rather than information becomes the scarce resource to capture. Friends become pointers and filters. Big data and artificial intelligence allow micro-targeting of communication so that people's information is limited to a "filter bubble" of the like-minded. The so-called "free" services of social media are based on a profit model in which the user or customer is actually the product, and their information and attention is sold to advertisers. Algorithms are designed to learn what keeps users engaged so that they can be served more ads and produce more revenue. Emotions such as outrage stimulate engagement, and false news which is outrageous has been shown to engage more viewers than accurate news. A study of demonstrations in Germany, for example, found that "YouTube's algorithm systematically directs users toward extremist content… It looks like reality, but deforms reality because it is biased toward watch time."[26] False news is often more outrageous than accurate news, and one study found that falsehoods on Twitter were 70 percent more likely to be retweeted than accurate news.[27] Fact checking by conventional news media is often unable to keep up in the race for attention, and sometimes can be counterproductive by drawing more attention. The nature of the social media profit model can be exploited as a weapon by states and non-state actors alike.

Recently Facebook chief executive Mark Zuckerberg wrote that "in 2016, we were not prepared for the coordinated information operations we regularly face. But we have learned a lot since then and have developed sophisticated systems that combine technology and people to prevent election interference on our services."[28] Such efforts include automated programs to find and remove fake accounts; featuring Facebook pages that spread disinformation less prominently than in the past; issuing a transparency report on the number of false accounts removed; verifying the nationality of those who place political advertisements; hiring 10,000 additional people to work on security; and improving coordination with law enforcement and other companies over suspicious activity.[29] Even so, the arms race will continue between the social media companies and the states and non-state actors who invest in ways to exploit their systems.[30] Artificial intelligence cannot alone solve this problem. It turns out to be far easier to develop an algorithm that identifies nudity than one that identifies hate speech. In 2018, Facebook reported that only 38 percent of hate speech was flagged by its internal systems compared to 96 percent of adult nudity.[31]

Ironically, because it is often more sensational and outrageous, fake news travels further and faster than real news and that makes it profitable. False information on Twitter is retweeted by many more people and far more rapidly than true information, and repeating false information, even in a fact-checking context, may increase an individual's likelihood of accepting it as true.[32] The Internet Research Agency in St Petersburg "spent more than a year creating dozens of social media accounts masquerading as local American news outlets."[33] Sometimes the reports favored a candidate, but often they were designed to give an impression of chaos, disgust and to suppress voter turn-out.

When Congress passed the Communications Decency Act in 1996, social media companies were treated as neutral telecoms providers where customers could communicate with each other, but this model of pipes that ignores content is clearly outdated. Under political pressure, the major companies have begun to police their networks more carefully and take down obvious fakes, including those propagated by botnets, but the question of limits on free speech is a problem. While machines and foreign governments have no First Amendment rights (and private companies are not bound by the First Amendment in any case), some abhorrent domestic groups and individuals have free speech rights in our democracy, and they can serve as intermediaries for foreign influencers. Foreign manipulation of accurate news about polarized American actors may have more impact than fake news. The damage done by foreign actors may be less than the damage we do to ourselves through polarized political rhetoric and tactics.

The social media companies have now encountered political controversy about their censorship of hate speech and conspiracy theorists. Companies want to avoid regulation, but legislators criticize them for both sins of omission and commission. This part of the problem will not be easy to solve because it raises trade-offs among our important values. Experience from European elections suggests that investigative journalism and alerting the public in advance can help inoculate against disinformation campaigns, but the problem of fake news is likely to remain a cat and mouse game between companies and fakers (both foreign and domestic) as part of the continual background noise of elections.[34]

Manipulating False Actors and Creating Astroturf Groups

Artificial actors can be created and manipulated to create chaos, social conflict, and disrupt the political process. Successful infiltration of the political process requires the creation of fake social media profiles that appear to be authentic, and then their coordination into false grass roots groups. For instance, in May 2016 there was a confrontation in Houston between demonstrators for and against a mosque screaming at each other (and being videoed for the Internet), but both the pro and anti-mosque protests had been planned and promoted by trolls in Russia.[35] A Russian-created account @Blacktivist had 360,000 likes on Facebook—more than the verified BlackLivesMatter account on Facebook. A Russian created group posted authentic video of black and white people hitting each other to exacerbate racial animosity. "Not only did the Kremlin create individuals and organizations on both sides of wedge issues, they also used targeted advertising to reach the audiences that they believed would be most receptive."[36] Other actions were to harass candidates or influential people with organized trolling, including botnets, to the point that they dropped offline. Again, companies can monitor their platforms and public exposure can help, but it is difficult to prevent external manipulation of domestic divisions by analysis of big data and micro-targeting sensitive groups. On the other hand, taking down fake accounts and artificial actors is less likely to encounter the thorny censorship and free speech problems that plague the fake news problem.

Using Artificial Intelligence and Deep Fake Videos

Computers have long been used to generate and manipulate images, but fake images were often detectable by shifts in lighting and voices were often slightly off in cadence or tone. Now with artificial intelligence and deep machine learning, it is "possible to doctor images and video so well that it is difficult to distinguish manipulated files from authentic ones." And with "generative adversarial networks, the algorithm works by getting really good at fooling itself."[37] Distributed ledger technologies may help in verification, but blockchain solutions may not be quick enough to prevent deep damage to political reputations.

When introduced late in a campaign, deep fakes may remain credible for long enough to alter an election result, particularly if they are embedded as brief offhand offensive remarks in otherwise authentic video. While companies are investing in research on counter measures such as digital watermarks, it is far from clear that the defensive technologies of detection will advance as rapidly as the offensive technologies of deception. Nonetheless, artificial intelligence may eventually help the defense as much as the offense if we invest in it.

A Strategy for Response

The defense of democracy in an age of cyber information war cannot rely on technology alone. It will require a strategy with several strands, and will have to involve many government departments, close coordination with the private sector, and will best be coordinated from the White House. The key elements will include domestic resilience and defense, deterrence, and diplomacy.

Domestic Resilience

Some steps are underway; others remain to be taken. Progress has been made on training and support of local election officials and upgrading the security of election infrastructure including paper backups.[38] Political parties, candidates and staffs have become more alert to the importance of basic cyber hygiene such as encryption and dual authentication, but phishing is always a danger and volunteers are often untrained. Various civic organizations are focused on the problem and investigative journalism and independent fact-checking has helped to alert the public and inoculate against some of the cruder forms of sharp power.

Laura Rosenberger of the Alliance for Assuring Democracy has suggested a number of further steps such as an honest ads act which would apply the same rules to online political advertising as apply to such ads on TV; a rule requiring social media companies to disclose any bots on their platform and prohibit candidates and parties from using bots; creating a better legal framework for protection of data privacy; and enhancing better mechanisms for information sharing among government agencies and with the private sector.[39]

The social media platforms are crucial to coping with this problem, but rather than heavy-handed regulation, a process should be set up for continual consultation and sharing of information between the companies and government. Rather than turn the companies into purely private censors, *The Economist* has recommended making the companies more publicly "accountable for their procedures: clarify the criteria applied to restrict content; recruit advisory bodies and user representatives to help ensure that these criteria are applied; give users scope to appeal against decisions. They also need to open their algorithms and data to independent scrutiny, under controlled conditions."[40] Independent bipartisan boards or commissions might enhance algorithmic accountability without revealing proprietary information in a damaging way. Rather than try to break up the companies or deprive them of all autonomy, it would be better to have them monitor their systems more effectively and in a publicly more accountable manner. As Alex Stamos has argued, the companies will "need to act in a quasi-governmental manner, making judgments on political speech and operating teams in parallel to the U.S. intelligence community, but we need more clarity on how these companies make decisions and what powers we want to reserve to our duly elected government."[41] But given the transnational scale of the companies, there will have to be provisions for cultural differences about values like privacy and fairness, and companies will have to obey local laws. Few other countries share American "First Amendment absolutism," and that includes allied democracies like Germany and France.[42]

More generally, a successful strategy would have to focus on raising the general level of cyber hygiene in society and government. This would not solve the problem, but it could remove the most vulnerable low hanging fruit and make the tasks of attackers more costly. Stronger cyber defense measures are like vaccinations in term of creating public goods, and legal frameworks could be developed to encourage this. The 2016 problems of hacking and doxing political emails could be made more difficult if dual factor identification and encryption were more widespread. Much could be done to encourage development of higher standards in software by revising liability laws and encouraging the development of the cyber insurance industry as the number of points of vulnerability to cyber intrusion expands exponentially with the Internet of Things.[43] Similarly, more can be done to improve the quality of cybersecurity in government agencies both by new investment and by raising standards.

Deterrence

Some skeptics believe that deterrence does not work in cyber conflict, at least not in the gray zone of hybrid warfare below the level governed by the law of armed conflict. They often cite the case of the 2016 election where President Obama personally warned President Putin to desist in September but to no avail, and where American intelligence officials have told the Congress that Russian interference continues. But the case is not definitive because American responses were inhibited by domestic politics in both parties. In September 2018, President Trump signed an executive order enabling sanctions (which include asset freezes and prohibitions from doing business) and defined foreign interference as efforts to "influence, undermine confidence in, or alter the result or reported result" of an election or "undermine public confidence in election processes or institutions." That broad definition would cover anything from state-sponsored social media campaigns to altering vote tallies. Its effectiveness remains to be seen, but deterrence must be a crucial part of a successful strategy.

Understanding deterrence in cyberspace is often difficult because our minds are captured by Cold War images of deterrence as threatening massive retaliation to a nuclear attack by nuclear means. The analogy to nuclear deterrence is misleading, however, because the aim with nuclear explosions is total prevention. In contrast, many aspects of cyber behavior are more like other behaviors, such as crime, that governments strive only imperfectly to deter. Moreover, cyber deterrence need not be limited to cyber responses, but can cross domains. There are four major mechanisms to reduce and prevent adverse actions in cyberspace: threat of punishment, denial by defense, entanglement, and normative taboos. None of these four mechanisms is perfect, but together they illustrate the range of means by which it is possible to reduce the likelihood of adverse acts causing harm. They can complement one another in affecting actors' perceptions of the costs and benefits of particular actions.[44]

Deterrence by defense involves many of the steps we already wish to take to enhance our resilience, and by hardening ourselves as a target, we affect the ratio of costs to benefits that an attacker expects. If the targets are soft and the costs are low, the temptations are greater. That need not be the case. Deterrence by entanglement refers to situations where an attacker holds back because the interdependence is so great that damaging the target may damage oneself. That level of interdependence does not exist between the United States and Russia, Iran or North Korea. And despite some progress in the UN Group of Government Experts that reported in 2015 on development of norms restricting damage to civilian targets, cyber taboos are not as strong as they are, for example, in biological weapons. It is interesting, however, that after the events of 2016, the United States added electoral processes to a list of 16 critical civilian infrastructures as a signal.

Deterrence by threat of retaliation remains a crucial but underutilized aspect of deterrence of cyber attack. There has been no attack on our electrical systems despite the reported presence of Chinese and Russians on the grid. Pentagon doctrine has announced that we will respond to damage with any weapon of our choice, and deterrence seems to be working at that level. Presumably it could be made to work in the gray zone of hybrid warfare as well if we had not been so pusillanimous in our responses to 2016 and 2017. Since American intelligence is reported to carry out espionage in Russian and Chinese networks, one could imagine that we discover embarrassing facts about the hidden assets of foreign leaders which we could threaten to disclose or bank accounts we could shut. Similarly, we could go further in applying economic and travel sanctions against authoritarians' inner circles. The diplomatic expulsions and indictments that have occurred thus far are only a first step toward strengthening our deterrent threat of retaliation.

Diplomacy

Negotiating treaties for cyber arms control involves a number of problems, but this does not make diplomacy impossible. In the cyber realm, the difference between a computer program that is a weapon and a non-weapon may come down to a single line of code, or the same program can be used for legitimate or malicious purposes depending on the intent of the user. Thus it will be difficult to anathematize the design, possession, or even implantation for espionage of particular programs. In that sense, cyber arms control cannot be like the nuclear arms control that developed during the Cold War. Verification of the absence of a stockpile of zero-day exploits would be virtually impossible, and even if it were assured, the stockpile could quickly be recreated. Unlike physical weapons, for example, it would be difficult to reliably prohibit possession of the whole category of cyber weapons.

But if traditional arms control treaties are too difficult, it may still be possible to set limits on certain types of civilian targets, and to negotiate rough rules of the road for behavior that limits conflict. For example, the United States and Soviet Union negotiated an Incidents at Sea Agreement in 1972 to limit naval behavior that might lead to escalation. The United States and Russia might negotiate limits to their behavior regarding each other's domestic political processes, in which we would draw a line between activities that constitute soft power and those that cross the line into sharp power. Even if they cannot agree on precise definitions they could exchange unilateral statements about areas of self-restraint and establish a consultative process to prevent escalation. Such a procedure of exchanging unilateral statements could protect democratic non-governmental organizations' right to criticize authoritarians while at the same time creating a framework that limits governmental escalation.

Skeptics object that such an agreement is impossible because of the difference in values between our two societies, but even greater differences did not prevent agreements related to prudence during the Cold War. Skeptics also say that since elections are meaningless in Russia, they would have no incentive to agree, but this ignores the potential threat of our retaliation across domains as discussed above. Others object to the implied moral equivalence, but since our democracy is more open and we have more to lose in the current situation, that should not hold us back from pursuing our self interest in developing a norm of restraint in this gray area.

As Jack Goldsmith puts it, "The United States needs to draw a strong principled line and defend it. That defense would acknowledge that the United States has interfered in elections itself, renounce those actions and pledge not to do them again; acknowledge that it continues to engage in forms of computer network exploitations for various purposes it deems legitimate; and state precisely the norm that the United States pledges to stand by and that the Russians have violated."[45] This would not be unilateral disarmament on our part since we would draw the line between soft and sharp power; overt programs and broadcasts would continue to be allowed. We would not object to the content of others' political speech but to how they pursue it through covert coordinated inauthentic behavior. Non-state actors often act as proxies of the state in varying degrees, but the rules of the road would require their open identification. Even if adherence to such rules of the road were imperfect on the part of authoritarian states, a reduction of their level of interference could make our defense of our democracy more feasible.[46]

Conclusions

Democracy depends upon open information that can be trusted. Authoritarian states can exploit and weaponize this openness. Information warfare is not new, and it has always presented a challenge to democracy, but technology has transformed the nature of the challenge. What's new is the speed with which such disinformation can spread and the low cost and visibility of spreading it. The Internet has expanded the information attack surface and the instruments that can exploit it. Electrons are cheaper, faster, safer, and more easily deniable than human spies. What is more, the business models of the large American social media companies can be readily manipulated by malign actors for criminal or political purposes. But as democracies respond to such challenges, they run the risks both of doing too little, but also too much. Measures that curtail openness and trust would become self-inflicted wounds. This will be true of both the defensive and offensive measures that democracies undertake. Imitation of the authoritarian practices would be a defeat.

In the case of Russian interference in the 2016 presidential election, the United States was poorly prepared and inadequate in its response. Different vectors of attack require different measures. Hacking and doxing of political actors requires greater awareness and practice of cyber hygiene. Hacking of electoral machinery and voter rolls requires more robust machines and audit trails as well as improved federal, state, and local cooperation. Thwarting and removing false accounts, botnets, sockpuppets, and astroturf actors requires strong action and cooperation among social media companies. Dealing with fake news designed to polarize, disrupt, and suppress voting also requires action by the companies but with procedures for protecting transparency in algorithms and processes that reveal difficult trade-offs regarding free speech. None of this will be solved easily. In some cases, artificial intelligence will help the offense, in other cases the defense. The game of cat and mouse does not end; it must be continually monitored.

At a more general level, a national strategy for defending democracy in the cyber age must include all three dimensions of resilience, deterrence, and diplomacy. American actions have been inadequate on all three dimensions but some useful steps have begun, and this discussion has suggested more that can be done. We are only at the beginning of a long process of protecting democracy in an era of cyber information war.

---

[1] I am grateful to Jack Goldsmith, Eric Rosenbach and Michael Sulmeyer for comments on an early draft.

[2] William J. Clinton, "Remarks at the Paul H. Nitze School," (Washington, March 8, 2000) http://www.presidency.ucsb.edu/ws/index.php?pid=87714

[3] Kenneth Osgood, "The CIA's Fake News," *New York Times*, October 14, 2017, pA19.

[4] Herb Lin and Jackie Kerr, "On Cyber-Enabled Information/Influence Warfare and Manipulation," *Oxford Handbook of Cyber Security*, Paul Cornish ed., Oxford University Press 2018, Christopher Paul and Miriam Matthews, *The Russian "Firehose of Falsehood" Propaganda Model: Why It Might Work and Options to Counter It* (Rand Corporation 2016).

[5] Sir David Omand, "The threats from modern digital subversion and sedition,"*Journal of Cyber Policy*, 2018, pp 1-19

[6] Michael Hayden quoted in Davis V. Gioe, "Cyber Operations and useful fools: the approach of Russian hybrid intelligence," Intelligence and National Security. https://doi.org/10.1080/02684527.20181479345. Hayden also is reported to have said "I would not want to be in an American court of law and be forced to deny that I never did anything like that as director of NSA." "Suing Spies," The *Economist*, September 15, 2018, p29.

[7] Suzanne Spaulding, "Countering Adversary Threats to Democratic Institutions: An Expert Report" Washington, Center for Strategic and International Studies, 2018, p 4.

[8] Kate Conger and Adam Satariano, "Twitter Clamps Down, But Rogue Accounts Turn the Pressure Up," *New York Times*, November 6, 2018

[9] Tony Romm, "New report on Russian disinformation, prepared for the Senate, shows the operation's scale and sweep," *Washington Post*, December 17, 2018. See also Scott Shane,

"Five Takeaways from Reports on Russian Campaign, *New York Times*, December 18, 2018, pA14

[10] Renee DiResta, "Russia's Information Warfare," *New York Times*, December 18, 2018, pA23.

[11] For details of "Operation Infektion", see P.W. Singer and Emerson T. Brooking, *Like War: The Weaponization of Social Media*, New York, Houghton Mifflin, 2018, p104.

[12] See Alexander Klimburg, "Hacking the Presidency," Review of Kathleen Hall Jamieson's CyberWar, *Nature*, Vol 562, October 11, 2018. See also Alexander Klimburg, *The Darkening Web*, New York, Penguin, 2017,2018,

[13] Tim Maurer and Garrett Hinck, "Russia: Information Security Meets Cyber Security," in Fabio Rugge, ed. *Confronting an Axis of Cyber?* Milano, Ledi Publishing, 2018, p 39.

[14] Eric Rosenbach and Katherine Mansted, "Can Democracy Survive in the Information Age," https://www.belfercenter.otg/publication/can-democracy-survive-information-age October 2018

[15] Singer and Brooking, cited, p 106

[16] Christopher Walker and Jessica Ludwig, *Sharp Power: Rising Authoritarian Influence*, Washington, National Endowment for Democracy, 2017.

[17] David Shambaugh, *China Goes Global,* Oxford, Oxford University Press, 2013

[18] Portland and USC Center on Public Diplomacy, *The Soft Power 30: A Global Ranking of Soft Power*. London, 2017

[19] Rosenbach and Mansted, cited.

[20] Singer and Brooking, cited, p21.

[21] See Charles Cleveland, Ryan Crocker, Daniel Egel, Andrew Liepman and David Maxwell, "An American Way of Political Warfare: A Proposal," *Perspective*, July 2018. Santa Monica,RAND, 2018.

[22] See Henry Farrell and Bruce Schneier**, "**Common-Knowledge Attacks on Democracy *Berkman Klein Center Research Publication No. 2018-7 November 2018* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3273111

[23] See Suzanne Spaulding and Eric Goldstein, "Countering Adversary Threats to Democratic Institutions," Washington, CSIS. February 2018.

[24] Michael Wines and Julian Barnes,"Fear Trolls, Not Hacked Voting Machines," *New York Times*, August 3, 2018, p16

[25] David Lazer et. al, "The Science of Fake News," *Science*, March 9, 2018, Vol 359, Issue 6380,

[26] Max Fisher and Katrin Bennhold, "Germans, Seeking News, Find You Tube's Far Right Tirades," *New York Times*, September 8, 2018, pA4.

[27] Sheera Frenkel, Mike Isaac, and Kate Conger, "With Growth, Social Media Spread Harm," *New York Times*, October 30, 2018, pA1

[28] Mark Zuckerberg, "Preparing for Elections," September 12, 2018, https://www.facebook.com/notes/mark-zuckerberg/preparing-for-elections/10156300047606634/

[29] Sheera Frenkel and Mike Isaac, "Facebook, After Reforms, Is Now Better Prepare to Ward Off Skulduggery," *New York Times*, September 14, 2018, p B2

[30] Kevin Roose, "Capitalizing on a Mass Killing to Spread Fake News Online," *New York Times*, October 3, 2017, p19

[31] Sheera Frenkel, Mike Isaac and Kate Conger, "With Growth, Social Media Spread Harm," *New York Times*, October 30, 2018.

[32] Lazer et al., cited above, p 1095

[33] Jared Cohen, "Confronting Hybrid Warriors and the Disinformation Tactics They Use," paper delivered at the Aspen Strategy Group, August 2018.

[34] Erik Brattberg and Tim Maurer, "Russian Election Interference: Europe's Counter to Fake News and Cyber Attacks," Washington, Carnegie Endowment, 2018. https://carnegieendowment.org/2018/05/23/russian-election-interfererence-europe-s-counter-

[35] Farhad Manjoo, "The Drama of Reality TV, Brought to You by Russia," *New York Times*, November 9, 2017, pB1. See also, Darrell West, "How to combat fake news and disinformation: a report.," Washington, Brookings, December 2017.

[36] Jared Cohen, cited above.

[37] Chris Meserole and Alina Polyakova, "The West is ill prepared for the wave of 'deep fakes' that artificial intelligence could unleash," Washington, Brookings Institution, May 25, 2018

[38] See for example, Defending Digital Democracy Project, *Election Cyber Incident Communications Coordination Guide*, Cambridge, Harvard Kennedy School, February 2018.

[39] Laura Rosenberger, "Countering Technologically-Driven Information Manipulation," paper delivered at the Aspen Strategy Group, August 2018. See also Laura Rosenberger and Jamie Fly, "Shredding the Putin Playbook," *Democracy Journal*, Winter 2018, pp 51-63.

[40] "Truth and Power," *The Economist*, September 8, 2015, p14

[41] Alex Stamos, "Yes, Facebook made mistakes in 2016. But We Weren't the Only Ones," *Washington Post,* November 17, 2018.

[42] See Frederick Schauer, "The Politics and Incentives of Legal Transplantation," in Joseph Nye and Jack D. Donahue, eds, *Governance in a Globalizing World*. Washington, Brookings,2000, p253

[43] Ariel Levite, Scott Kannry, Wyatt Hoffman, "Addressing the Private Sector Cybersecurity Predicament: The Indispensable Role of Insurance," Washington, Carnegie Endowment, November 2018.

[44] J.S. Nye, "Dissuasion and Deterrence in Cyber Space" , *International Security*,41 (3) 47-71

[45] Jack Goldsmith, "Uncomfortable Questions," https://www.lawfareblog.com/uncomfortable-questions-wake-russia-indictment-20-and-trumps-press-conference-putin. See also Jack Goldsmith, "The Failure of Internet Freedom," *Knight Foundation Emerging Threats Series* (June 2018), available at https://knightcolumbia.org/content/failure-internet-freedom, and Jack Goldsmith and Stuart Russell, "Strengths Become Vulnerabilities: How a Digital World Disadvantages the United States in Its International Relations," *Hoover Institution Aegis Papers, No. 1806* available at https://www.hoover.org/sites/default/files/research/docs/381100534-strengths-become-vulnerabilities.pdf.

[46] Alex Grigsby, "Russia Wants a Deal with the United States on Cyber Issues: Why Does Washington Keep Saying No? New York, Council on Foreign Relations. Net Politics. August 27, 2018

Joseph S. Nye is the university distinguished service professor, emeritus, at Harvard University and the former Dean of the Harvard's Kennedy School of Government.

Observations from the Roundtable

The information and communications revolution has complicated governance everywhere. It has broken down traditional borders: people can communicate, organize, and act both with their fellow citizens and across country boundaries. The age-old challenge of governing over diversity grows more difficult by the day.

Prior to the spread of the internet, as Niall Ferguson writes, "only the elite could network globally." Now nearly two and half billion people do so on social media platforms. At the advent of this networked age, most practitioners pictured a better educated, more knowledgeable populace enlightened by the democratization of information. The public would have up-to-the-minute access to information and the ability to communicate globally. Voters would have more ways to learn about candidates and engage in political speech than ever before.

That Panglossian view of the internet proved simplistic. The spread of information and new means of communication—particularly social media and other network platforms—created new vulnerabilities in democratic states. Joseph Nye explains that while authoritarian regimes can manipulate or even control information flows, democracies, in their commitment to transparency, find themselves on the defensive.

Foreign actors can manipulate information, particularly in the cyber domain, to undermine trust in institutions, sow domestic discord, encourage partisanship, or otherwise complicate electoral and governance processes, as the Russian's demonstrated in their interference with the 2016 election. But such behavior is not the sole purview of foreign entities. Private citizens and corporations alike have the power to influence election outcomes in new and powerful ways.

If the 2016 presidential election showed the American public the potential of network platforms as tool of political manipulation, it also taught us how thorny the problem is. Russia's interference in the election achieved an important goal: it helped to undermined faith in the American electoral and political process and in this country's democratic reputation.

The papers prepared for this program address two separate but related issues: 1) the domestic problem of managing the highly powerful network platforms and 2) the international problem of information warfare enabled by these new communications technologies. The former demands a reconsideration of U.S. policy at home: the current status quo of "self-non-regulation" by the network platforms has proved wanting. The latter requires both U.S. policy corrections and multinational engagement: as Joseph Nye argues, we can improve our resilience to foreign information campaigns and our capacity to deter them while also engaging in diplomacy to define new rules of the road.

What follows is not a definitive statement of what the United States must do to address these two facets of the governance challenge. Instead, we endorse certain recommendations regarding information warfare and propose a set of potential corrections—informed by the roundtable discussion of these papers—to the domestic problem so well-defined by our colleague Niall Ferguson.

Cyber Information Warfare

Information warfare is an old form of competition, and one practiced by friends and foes alike. As Joseph Nye explains, the British cut Germany's overseas communications cables at the outset of World War I, but they also fed the United States the Zimmermann Telegram to encourage U.S. engagement later in the war. As old as it may be though, new technologies have made information operations faster, more effective, and cheaper.

Russia's interference in the 2016 election comes to mind again: 126 million Americans saw posts generated by the St. Petersburg-based "Internet Research Agency". Whereas the Soviet Union's Operation Infektion conspiracy about AIDs took four years to spread into mainstream media, the recent Comet Pizza conspiracy theory spread across the country in a matter of hours. And the cost of creating a Facebook post or generating other online content is microscopic compared to that of traditional human intelligence operations.

In his paper, Joseph Nye draws the distinction between sharp and soft power. Soft power, he writes, rests on persuasion, while sharp power involves deception or coercion. Soft power is exercised openly, sharp power covertly.

The openness of the American system makes it more vulnerable to sharp power than more closed systems are, and states and non-state actors alike have a host of tools available for disrupting democratic processes: manipulation of voters through fake news, targeting candidates anonymously or under false names, creation of inauthentic groups to generate conflict, and sowing of chaos and disruption. Russia, China, North Korea, Iran, and others all wield these tools against the United States.

However, while Russia proved adept at exercising sharp power, Russia and other authoritarian states, including China, are less adept at soft power. Russia's actions in the 2016 election, for example, fall under the umbrella of sharp power. The Russian news channel RT, on the other hand, generally engages in the above-board exercise of soft power. An American soft power analog would be Radio Free Europe and Voice of America, which were powerful tools of information warfare during the Cold War.

New technologies—chiefly social media and other network platforms—are fertile ground for the exercise of cyber information war. They can promote polarization and spread fake stories. The business of Facebook, YouTube, or Twitter is to maximize the attention of their users.. The more time and attention, the more advertisements seen, the more ad money for the platforms. False stories, outrage, and emotion capture attention far better than sober-minded articles or videos. It is unsurprising then, that YouTube's algorithms, for example, tend to suggest videos that push viewers towards more extreme ends of the political spectrum. You may begin at the center, but the suggested content will push you to the extreme.

As modern technology, including artificial intelligence, make the manipulation of images and videos easier, bad actors can create fake or altered content that are increasingly difficult to distinguish from authentic ones. Introduced late in a campaign, such altered images could spread quickly enough—before Facebook's operators, say, could take them down—to influence the outcome of that election.

Russia learned to weaponize social media and use it as a tool against the United States. The Internet Research Agency and similar operators created fake accounts, catalyzing polarization in American society, and amplifying extreme voices on both sides of the aisle in the United States.

How to Protect Our Democracy from Foreign Interference?

The openness of the United States may be a vulnerability, but it is also a great value. We must be careful not to sacrifice it. In other words, the U.S. government should not try to stop transparent information campaigns—legitimate exercises of soft power—in its effort to secure itself against illicit interference. Nor should it look to technology companies to solve the challenges. Facebook has taken steps to address the problem, after not seeing it coming in 2016, hiring new employees and applying artificial intelligence to find and remove hate speech, bots, and false accounts. But the enormous quantity of content, the entanglement of foreign- and domestically-generated content, and the mix of human and bot actors complicate the problem; the vast majority of Russian posts during the 2016 election amplified existing content created by Americans. Moreover, just as AI can help monitor and police content, it can also be used to generate new, harder to identify false content. The technical contest between the network platforms and foreign agents is likely to remain a cat and mouse game.

Joseph Nye proposes a three-fold approach, which we believe is wise: The United States should look to increase resilience and strengthen deterrence at home, while engaging in diplomacy with foreign powers.

Resilience: The United States must take steps to harden its electoral and political systems against cyber information warfare.

The U.S. government and non-governmental institutions, such as the academy, could upgrade the security of U.S. election infrastructure by training local election officials and improving basic cyber hygiene, such as using two-factor authentication. Given how much political campaigns and electoral offices rely on interns, volunteers, and other part-time workers, it may be difficult to train everyone, but even some training and better resilience would make a difference.

We should also encourage development of higher standards in software by revising liability laws and encouraging development of the cyber insurance industry as the number of points of vulnerability to cyber intrusion expands exponentially with the internet of things. Election laws could also change to force candidates to put their names on online political ads just as they do for television ads and to ban the use of bots by political parties or campaigns. As in other areas of cybersecurity, improved information sharing between government and industry would contribute as well.

Deterrence: Deterrence can be established in four ways: through the threat of punishment, denial by defense(resilience), entanglement, and establishment of normative taboos. Effective punishment would, of course, depend on reliable attribution. In nuclear strategy, the aim of deterrence is total prevention of nuclear attack, by maintaining an assured ability to retaliate with a devastating strike. Deterrence of cyber information warfare, at the other end of the spectrum of conflict, need not be perfect to be useful but ought to raise the cost for malevolent actors. Cyber intrusion could be treated as we do crime. When seeking to deter criminal activity, the certainty of getting caught matters more than the severity of the punishment, so better, faster attribution and action will be key. The Trump administration's September 2018 executive order promising sanctions in response to election interference is a step in the right direction. Entanglement complements punishment; if an attack on the United States hurts the attacker, that reduces the incentive for malicious behavior. Deterrence won't solve the problem but could increase the cost and difficulty. Defensive measures could then be better focused on those attacks that do still occur.

Diplomacy: This arena is not conducive to arms control. A Twitter account is inherently "dual-use": a tool of disinformation or a means of innocuous social networking; the key variable is the user and the user's intentions. Instead, as Joseph Nye proposes, we ought to establish rules of the road to limit certain malicious behavior.

We are not proposing a treaty but a set of agreements or understandings, which will depend on the values of the involved parties. Just as the United States and the Soviet Union came to the 1972 Incidents at Sea agreement to reduce the risk of inadvertent crises, so too could the United States and Russia conceivably commit not to interfere covertly in elections, while allowing overt broadcasting and transparent information. Each side could unilaterally propose and share its own expectations of conduct, tracking and communicating how the cyber behaviors it observes over time do or do not comply with those expectations. The United States does not have to act alone here. It could work with its allies and partners—fellow liberal democracies—to coordinate collective action; sharing defensive recommendations, mutually shoring up electoral and political processes, and collaborating on diplomatic agreements.

In other words, work to establish upper-bounds of cyber information activities—thereby allowing U.S. officials and others to focus their resilience-building efforts on a narrower range of challenges—and prepare for prompt retaliation for activities that exceed the bounds.

What to Do About Network Platforms?

As the United States addresses cyber information warfare, it ought to consider the preeminent and uncontested power of network platforms. Manipulation of information to disrupt our electoral process demands a response, but the information challenge to governance extends beyond cyber information war. The technologies, and our relationship to them, must be addressed.

Niall Ferguson ably describes the current status quo: eight technology companies—including Facebook, Alphabet, and Tencent— dominate global internet commerce and advertising. They are near monopolies and immensely profitable. Network platforms have become a "public good," not just commercial enterprises, trading on the attention of the public. But they are contaminated with fake news and extreme views, some incited by our nation's adversaries. And network platforms, such as Twitter, have transformed governance in the United States.

They may be public goods, but these platforms are essentially self-regulated, or more accurately self-non-regulated. What regulation exists gives the network platforms significant leeway. Under US federal law, they are generally not regarded as publishers nor are they liable for the content they host, or the content they remove.

It is unsurprising, then, that companies curate and customize content on their platforms. As described above, they seek to maximize user attention and have done so to great effect—the average American spends 5.5 hours per day on digital media. Alongside this comes fake news and polarizing content, which spread more quickly and attract more attention than sober-minded alternatives.

With their vast network of users and grasp of user attention, U.S. internet platforms became a key battleground of the 2016 election—one in which the winning campaign was most focused.

Regulation, Firewalls, and Other Proposals

If the status quo is unacceptable, what should be done to change it? Two foreign models for managing internet platforms suggest what not to do:

Europe has adopted a tax, regulate, fine model. As Ferguson writes, Europe "seeks, at one and the same time, to live off the network platforms, by taxing and fining them, and to delegate the power of public censorship to them." China, on the other hand, zigged when the West zagged, adopting "internet sovereignty" in contrast to the U.S.-led internet freedom agenda. It built the great firewall and fostered its own domestic industry through total protectionism (for more discussion of this see "China in an Emerging World" in this series).

Neither foreign model appeals. In the wake of the 2016 election, U.S. network platforms responded—within their own largely laissez-faire business environment—by pledging more strenuous self-regulation at the firm level. Facebook, for example, now requires disclosure of who pays for political ads, uses artificial intelligence to detect bad content and bad actors, removes certain foreign government accounts, and reduces access to user data. These are measurably helpful but remain essentially reactionary steps. It is hard to have confidence that they have solved the next threat.

How to Manage the Network Platforms?

While we do not know the precise solution to these problems, let us consider a few options, drawn from the papers included herein and the roundtable discussion of them. It is easy to focus on the new problems generated through these platforms while taking for granted the informational value and personal satisfaction they also do generate. We therefore wish to redress the more damaging effects of these technologies while continuing to take advantage of their promise.

Niall Ferguson's paper recommends that the U.S. government make network platforms liable for content on their products—essentially scrapping the 1996 Telecommunications Act provisions protecting them—while also imposing First Amendment obligations on them. That is, do not allow the platforms themselves to decide what speech is acceptable by their own rules. His approach would give users and competitors recourse to challenge companies in the courts.

Ferguson's diagnosis of the fundamental flaw at the heart of the current regulatory framework rings true, and he rightfully focuses on how network platforms have come to dominate the public square. But as discussants noted, there are certain internal contradictions. Asking companies to monitor content for which they could be held liable—a task that would necessarily rely on AI—would likely complicate their ability to post everything permitted by the First Amendment. And what of anonymity, which has been so crucial to the internet freedom agenda? How do we protect anonymity while also enforcing liability? Perhaps a first step would entail banning content generated by bots or nonhumans.

Alternative, but unappealing regulatory steps would include a return to net neutrality, which would empower internet service providers to monitor content, but there is little reason to believe they would do a better job given their own profit incentives. Antitrust efforts intended to break up platform companies would also likely be of limited utility: it would be slow, of questionable effect, and run against the natural "winner takes all" direction of network platforms. Moreover, we must remember that historically regulation tends to cement the dominance of the largest players, stifling innovation and competition.

The government and public could ask more of the tech industry. We could press companies to be more transparent about their criteria for managing content on their platforms, while also establishing a recourse to challenge network platforms' actions. Along those same lines, companies could be ordered to make some portion of their algorithms available for public review or to the review of a select court, in the vein of the FISA court.

More generally, both the public and U.S. government officials ought to be cognizant of the immense political power internet companies hold. As Robert Epstein has documented, they can shift election outcomes and public opinion through slight manipulation of search results or suggestions, content feeds, and other user interfaces. So-called "dark patterns" are a well-documented aspect of digital interaction design outside of the political arena and are an emergent threat here too: almost imperceptible tweaks to underlying algorithms can swing voters towards a single candidate. As we consider whether network platforms ought to continue to self-regulate, we would do well to recall their power.

Conclusion

Social media and network platforms have come to dominate the public square, enabling broader and more complex social networks and political organizations than ever before. Information has always been an extremely valuable asset—once costly to obtain and share, now essentially free to all strata of society. But the spread of internet platforms comes at a cost. Malicious actors can engage in information warfare more quickly, decisively, and cheaply than ever before, and fake news, disinformation, and polarizing political speech proliferate. Whereas social media were once seen as a tool to disrupt non-tech-savvy authoritarians, we are increasingly aware of how they can be manipulated to transform democratic elections too.

Democracy—both elections and the process of governance—depends on transparency and the spread of open, trustworthy information. That commitment to openness is a vulnerability, but it is also a great virtue, one deserving of protection. When considering what can be done to redress the information challenge to governance, then, we must commit to first and foremost doing no harm to our democracy.

Fortunately, the United States can counter cyber information warfare without curtailing its own openness—indeed while strengthening that core value. It can, as Joseph Nye proposes, pursue a strategy of improving resilience and deterrence, while engaging in diplomacy to establish rules of the road governing interference in elections. Moreover, it is worth noting that the information battleground is not static. Russia caught the United States unprepared in 2016, but it was punished for doing so, primarily in the form of sanctions. And the U.S. government and network platforms have raised the costs of engaging in such behavior, though there is much work to be done.

In the United States, some may look to the past, when symbols of public trust—Walter Cronkite and Huntley and Brinkley being the canonical examples—gave us the news. Those days have passed, but the importance of trust and reputation remain. Internet companies would be wise to regain the trust of the people through careful stewardship of their platforms, giving priority to accuracy over attention, and willing public-private engagement. The public has an important role to play as well. As both creators and consumers of the content that populates network platforms, individuals can refrain from relying on social media for "news" and be discerning in what they share. The government and the companies do not bear sole responsibility for addressing this challenge.

Finally, what happens after an election? The papers in this program and the discussion of them focused on ways to safeguard and improve the electoral process, but the challenge of governing once in office is also formidable. How do these new means of communication affect the capacity of political officials to govern over diversity? We will continue to address this subject in the course of our project.

Notes

About

New and rapid societal and technological changes are complicating governance around the globe and challenging traditional thinking. Demographic changes and migration are having a profound effect as some populations age and shrink while other countries expand. The information and communications revolution is making governance much more difficult and heightening the impact of diversity. Emerging technologies, especially artificial intelligence and automation, are bringing about a new industrial revolution, disrupting workforces and increasing military capabilities of both states and non-state actors. And new means of production such as additive manufacturing and automation are changing how, where, and what we produce. These changes are coming quickly, faster than governments have historically been able to respond.

Led by Hoover Distinguished Fellow George P. Shultz, his Project on Governance in an Emerging New World aims to understand these changes and inform strategies that both address the challenges and take advantage of the opportunities afforded by these dramatic shifts.

The project features a series of papers and events addressing how these changes are affecting democratic processes, the economy, and national security of the United States, and how they are affecting countries and regions, including Russia, China, Europe, Africa, and Latin America. A set of essays by the participants accompanies each event and provides thoughtful analysis of the challenges and opportunities.

Efforts have been made to locate the original sources, determine the current rights holders, and, if needed, obtain reproduction permissions. On verification of any such claims to rights in the articles reproduced in this book, any required corrections or clarifications will be made in subsequent printings/editions. Hoover Institution assumes no responsibility for the persistence or accuracy of URLs for external or third-party Internet websites referred to in this publication, and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

Hoover Institution,
Stanford University
434 Galvez Mall
Stanford, CA 94305-6003
650-723-1754

Hoover Institution in Washington
The Johnson Center
1399 New York Avenue NW, Suite 500
Washington, DC 20005
202-760-3200

HOOVER
INSTITUTION
ONE HUNDRED YEARS