

How Digital Disinformation Turned Dangerous

Dave Karpf

They say history doesn't repeat itself, but it rhymes.

Writing for WIRED magazine in January 1997, Tom Dowe reflected on the spread of online rumors, conspiracy theories, and outright lies that Bill Clinton had faced in the 1996 election. His article, titled "News You Can Abuse," will spark a sense of *déjà vu* for any reader familiar with the digital misinformation practices that surfaced throughout the 2016 election:

The Net is opening up new terrain in our collective consciousness, between old-fashioned "news" and what used to be called the grapevine – rumor, gossip, word of mouth. Call it paranews – information that looks and sounds like news, that might even be news. Or a carelessly crafted half-truth. Or the product of a fevered, Hofstadterian mind working overtime. It's up to you to figure out which. Like a finely tuned seismograph, an ever more sophisticated chain of Web links, email chains, and newsgroups is now in place to register the slightest tremor in the zeitgeist, no matter how small, distant, or far-fetched. And then deliver it straight to the desktop of anyone, anywhere who agrees with the opening button on the National Enquirer Web site "I Want to Know!"¹

The parallels to today's digital news controversies are so obvious that they ruin the punchline. It would appear as though online misinformation, disinformation, and "fake news" has been spreading about Democratic candidates named Clinton since the very first internet-mediated election. And even back in 1997, Dowe was raising some of the same concerns that we face today: "When the barriers come down, when people cease to trust the authorities," he writes, "they – some of them, anyway – become at once more skeptical and more credulous. And on the Net right now – hell, in America – there's plenty of evidence of that."

Is Dowe's "paranews" really all that different from the weaponized disinformation campaigns that we witnessed in 2016? A cynic might conclude that the key difference between the two cycles is that 1996's Clinton won and 2016's Clinton lost. (How different, after all, would the contents of this volume be if the election had narrowly swung the other way?) But such cynicism is both unwarranted and unproductive. The online rumor mills of the early Web are substantially different from the industrialized digital disinformation and misinformation operations that trouble us today. The real value of reflecting on the paranews of 1996 is that it provides a helpful point of comparison to see just how much the digital context has changed.

The Internet is not new media any longer. The World Wide Web has over a twenty-five-year history. Digital media is no longer our looming technological future. It has a track record from which we can make observations and draw lessons. We need no longer make static comparisons between mainstream/mass media and digital/social media. We can instead make apples-to-apples comparisons within the digital era, identifying commonalities and differences between today's digital landscape and the digital media of past decades.

The purpose of this chapter is to explore how the digital media landscape has changed over time, and how these changes impact the status of fake news, misinformation, and disinformation. The chapter focuses on three major developments that make today's digital disinformation and propaganda more dangerous than it was in decades past. First, rumors and misinformation spread at a different rate, and by different mechanisms. Second, there is both more profit and more power in online disinformation today than there was two decades ago. Third, online misinformation has now been with us long enough to alter elite permission structures. The chapter concludes by discussing what digital platforms, policymakers, and journalists can do to confront these changing circumstances in the years ahead.

MECHANISMS OF DIFFUSION

Both the Internet and the broader media system were substantially different in 1996 and 2016. The Internet that Tom Dowe was describing was populated by different technologies with different affordances, encouraging different behaviors. It was an Internet of desktop computers and America Online CD-ROMs, an Internet of dial-up modems and search engines that were laughably bad at providing accurate

search results. The “new media” of 1996 was characterized by the expansion of cable television and the growth of conservative talk radio. Fox News Channel debuted in October 1996, attempting to copy CNN’s successful business model. Today’s disinformation can spread more quickly because of a set of structural changes to the overall media system.

Consider how online rumors and disinformation spread in the mid-1990s Internet: one could (a) spread salacious gossip through email forwarding chains, or (b) post made-up stories on a website, or (c) make false claims in an online chatroom. Each of these options is self-limiting for the spread of online rumors.

Chain emails are traceable and relatively costly. You know who forwarded them to you, and you probably have some experience with the veracity of the stories they share. Email forwarding is a relatively high-bar activity in the digital landscape. In today’s terms, it takes more work to forward an email to 100 friends than it does to “like” or retweet a post, sharing it with everyone in your network who is then algorithmically exposed to your social media activity. These are structural characteristics of email forwarding chains. Conspiracy theories via email, in other words, are spread by the known conspiratorial thinkers in one’s network; they can be discounted by recipients accordingly.

Conspiratorial websites in the mid-1990s also had a sharply limited audience. This was the pre-Google Internet, where search was time consuming and difficult. Online writers sought to build traffic by forming “web rings” with fellow travelers, and by filling their websites with keywords that might be typed into the Yahoo/Alta Vista search box. Incidental exposure to conspiratorial websites was thus limited. If you wanted to find information about all manner of Clinton conspiracies in 1997, there were websites to indulge your interests. But you would have had to look pretty hard. Again, these are structural characteristics of the World Wide Web of the 1990s that matter for how gossip, propaganda, and disinformation spread through the system.

Chat rooms face a parallel set of constraints. Chats are segregated by topic and occupied by small groups, making them a poor vector for incidental exposure to misinformation and disinformation. The Internet of 1997 provided a virtual space where adherents to all sorts of Clinton conspiracy theories could gather and swap tall tales. But if they entered a random AOL chatroom to post their screeds, they would not find much of an audience. Disinformation efforts via chatroom are liable to fail because they will appear as off-topic ramblings, inserted into an online

conversation among a small group of participants who can just move elsewhere.

The result is that conspiracy theories on the web of the 1990s had quite a lot in common with conspiracy theories in previous media. Dowe's reference to the *National Enquirer* is instructive. Salacious gossip and misinformation did not begin with the Internet. They were spread through tabloids, and through radio programs, and through newsletters. The early Web made misinformation easier to find. It made it easier to interact with like-minded conspiratorial thinkers. But it was a difference in degree, rather than a difference in kind.

By comparison, let's consider how these limiting conditions of the early Web compare to the industrial production of misinformation in the 2016 election. As Samanth Subramanian documents in his WIRED article, "Inside the Macedonian Fake News Complex," the 2016 election featured entire websites set up with the semblance of reputable new outlets.² These websites invented salacious stories, engineered to maximize social sharing and public exposure. They advertised cheaply on Facebook, boosting their visibility in news feeds. NewYorkTimesPolitics.com was one such fake news website, designed to resemble the real *New York Times* website, and featuring plagiarized articles on American politics. Unlike the chain emails of 1997, these stories were shared through social media, spreading faster and farther while presenting fewer signals of their (lack of) source credibility.

Meanwhile, employees of Russia's Internet Research Agency (IRA) piloted swarms of automated and semi-automated social media accounts with fake, US-based profiles. These accounts sought to influence the public dialogue and amplify disagreement and discontent in online discourse. They liked, shared, and retweeted social media posts. They attacked authors and spread misinformation in comment threads, manufacturing the appearance of broader social distrust of Hillary Clinton's candidacy. Where the chatrooms of 1996 were a terrible vector for spreading disinformation, their capacity for amplification was limited; in 2016, however, the deliberate amplification of conspiracy theories and mistrust helped propel topics deemed harmful to Hillary Clinton into the broader public sphere.

Alongside the different affordances of the modern Internet, we also have to reckon with the Internet's changing status within the broader media ecosystem. As Yochai Benkler makes clear in Chapter 2 of this book, American political journalism has changed drastically over the past few decades. Newspapers have been hollowed out. Conservative outlets

from Fox News (founded, incidentally, in October 1996) to Breitbart now play a central role in fueling the spread of conservative propaganda and strategic misinformation. More broadly, as Andrew Chadwick suggests in his 2013 book, *The Hybrid Media System*, digital media has changed the rhythms of news production, converting traditional news cycles into what Chadwick terms “political information cycles.”³ Episodes of political contention now move back and forth between social media, television, radio, and newsprint. Online conspiracy theories do not remain isolated online – trending hashtags and artificially boosted clickbait stories can become the topic of the nightly newscast, dramatically expanding the reach of rumors and misinformation.

Conspiracy theories on the early Web were treated by the broader media system much like rumors in the pages of the *National Enquirer* or other tabloids. They did not set the mainstream news agenda. They were not incorporated into newsgathering routines. They were at best an oddity, or a whisper that might lead to a story pitch. But digital news was not yet a competitor, either for eyeballs or for advertising revenue. This was a pre-blogsphere Internet. Conspiracy theorists could not influence news routines by swamping comment threads on news websites. News organizations were not yet monitoring clicks or hyperlinks to judge the news value of a given story. The digital challenges to traditional journalism were not yet viewed as a looming threat by newsrooms. As Paul Starr notes in Chapter 3 of this volume, the Internet of the 1990s was characterized by a sense of naïve technological optimism, particularly amongst its vocal advocates and early adopters who believed the technology would soon usher in a new era of rational and critical civic discussion. The Web was decentralized and barely populated. The dotcom boom was still in its first year. Conspiracy theories online were an odd sideshow, rather than an outright social ill.

By 2016, in contrast, major news organizations have adapted to the hybrid media system, modifying their news routines to incorporate trending topics and viral stories into their agenda-setting process. The fact of a viral story is itself news, regardless of the underlying veracity of the story itself. The conservative ecosystem of media organizations (both digital, television, and radio) stokes these stories, decrying the lack of coverage in mainstream outlets and demanding coverage of “both sides” of the manufactured controversy. The hybrid media system is much more vulnerable to strategic misinformation and disinformation than the industrial broadcast media system that still dominated American politics in 1996.

What has changed, then, is both the exposure rate, the traceability, and the lateral impact of misinformation. Digital misinformation has become progressively less traceable, less costly, and more spreadable, while developing a more substantial role in traditional news organizations. When you were handed a John Birch Society newsletter, you could see quite clearly where the newsletter came from and who gave it to you. Those newsletters were filled with disinformation, but they did not travel far and they did not set the agenda for the nightly news. The early Web had many of the same qualities. Today's social media has become unmoored from those limitations.

And the reason why it has become so unmoored leads to my second observation.

PROFIT AND POWER

To state it plainly, fake news in the 1990s was a hobby. Today it is an industry.

As Subramanian notes in his article on the Macedonian fake news industry, "Between August and November, [young Veles resident] Boris earned nearly \$16,000 off his two pro-Trump websites. The average monthly salary in Macedonia is \$371." The mechanics of this money-making operation are entirely determined by how online advertising revenue is generated through Google and Facebook. The purveyors of these manufactured stories would pay Facebook to promote their content in the news feed. Scandalous headlines generated clicks, comments, and shares, and each visitor to the website generated profit through Google AdSense. Though there is now some controversy as to the nature of the relationship between Veles residents and Russian information operations, the Macedonians claimed contemporaneously that they were not particularly interested in supporting Trump or opposing Clinton – they just found that anti-Clinton fake stories generated more traffic (and thus, more advertising revenue).⁴

The incident is a testament to a broader phenomenon in today's hybrid media system. The dynamics of mass attention and of advertising profitability are overwhelmingly shaped by the algorithmic decisions of two corporations: Google and Facebook. As journalist Joshua Micah Marshall describes in his 2017 essay, "A Serf on Google's Farm," about Google's involvement with his digital news site, Talking Points Memo (TPM); "Google has directly or indirectly driven millions of dollars of revenue to TPM over more than a decade. . . . few publishers really want to

talk about the depths or mechanics of Google's role in news publishing." He details the degree to which Google is implicated in the news, owning as it does: "1) The system for running ads [DoubleClick], 2) the top purchaser of ads [AdExchange], 3) the most pervasive audience data service [Google Analytics], 4) all search [Google.com], 5) our [TPM's] email."⁵ Marshall goes on to describe how Google's sheer market power can dictate the success or failure of digital news organizations. TPM was blacklisted by Google for violating the company's ban on hate speech. This was a false positive – TPM was reporting on incidents of white supremacist violence, and the reporting was coded as hate speech – but it was a potential economic catastrophe for the news site, because Google is the center of the digital advertising economy.

Facebook, likewise, has arguably become *the* central vector for the social sharing of news and information. Changes to Facebook's algorithmic weighting can create or destroy the market for particularly forms of journalism. As I discuss in *Analytic Activism*, this was the major public lesson of Upworthy.com, a social news site that specialized in developing Facebook-friendly headlines to drive attention to stories and videos with social impact.⁶ In 2013, Upworthy was the fastest growing website in history. Then Facebook debuted a new Facebook video feature, and penalized websites that linked to videos outside of the Facebook ecosystem. Upworthy immediately lost roughly two-thirds of its monthly visitors.

Herein lies the problem with the "marketplace of ideas" arguments that frequently appear in current debates over the negative consequences of online speech. The Web of the 1990s could arguably be thought of as a neutral marketplace of ideas, one in which anyone with a dial-up connection and a bit of training in HTML could write online and potentially find a modest audience. The "Safe Harbor" provision of the Communications Decency Act (Section 230) was designed to help protect free speech by making websites non-liaible for the content that visitors posted to them. That was a reasonable and appropriate provision at the time. But in the intervening years, the Internet has recentralized around a handful of quasi-monopolistic platforms. And in the meantime, online advertising has experienced massive growth, while the advertising markets that supported the industrial broadcast news system have been cannibalized.⁷

Consider how these changes have impacted the status of online rumors and disinformation. Dowe's 1997 article quotes digital pioneer Esther Dyson, who tells the author, "the Net is terrible at propaganda, but it's

wonderful at conspiracy.” This is a remarkable statement, viewed in retrospect. The Internet of 2016 is clearly quite good at propaganda – at least as good as the mass media of decades past! Part of this change is because the broader public has come online. It was a terrible propaganda channel in 1997 because there was not yet a mass audience to be propagandized. Fake news in the 1990s was a hobby because the Internet in the 1990s was confined to hobbyists. Digital media today is everywhere, always on, and always with us.

Alongside this secular expansion in Internet use, the technologies of digital ad targeting have also advanced greatly in the intervening twenty years.⁸ As the masses came online, the Web became more valuable as a substitute for and complement to mass media. The Web has also become more valuable as data, providing insights into what we read, what we purchase, and where we are physically located at all times. Cookie-based and geolocal tracking provide a wealth of data, which in turn has funneled additional investments into online media. While today’s digital advertising is still far less precise than its marketers routinely claim⁹ (Google and Facebook do not actually know you better than you know yourself), the digital advertising economy now determines which speech is profitable, and thus which types of journalism, propaganda, public information and disinformation will receive broad dissemination. The platform monopolists are too big to be neutral; their algorithmic choices are market-makers, with an indelible impact upon the marketplace of ideas.

The result is a situation in which there can be strong economic incentives for misinformation and disinformation campaigns. The online marketplace does not reward the best ideas, or the most thorough reporting. It rewards the stories that perform best on Facebook, Twitter, Google, and YouTube. It rewards user engagement, and social sharing, and time-on-site. Meanwhile there are also compelling strategic incentives for misinformation and disinformation campaigns. The Russian Internet Research Agency is not designed to make money.¹⁰ It is designed to spread mistrust and discontent online. And the logic of troll farms like the IRA is, that now so much of the public is online, disrupting online media can be a high-value propaganda goal. The marketplace for speech will permanently malfunction if lies are made more profitable than truths.

This is not an inherent problem to the Internet or social media. It has gotten worse because of specific policy decisions that have protected and rewarded bad social behaviors. It can be fixed through different policy decisions – the fake news industry in Macedonia disappeared after the 2016 election, as Google implemented new policies that excluded the fake

news websites from the AdWords program. In 2011, Google likewise dramatically curtailed “content farms” through the quasi-regulatory act of adjusting the company’s search algorithms.¹¹ Regulation ought to come from the government, but in the absence of government oversight, the platform monopolies play an uncomfortable, quasi-regulatory role. To be clear, Facebook and Google are not going to create voluntary rules that do much to curtail their own power or profit. But they can, and do, slowly respond to the worst abuses of their platform in order to safeguard their reputation.

The more urgent issue is that government regulators in the United States have essentially abandoned their posts. At the time of writing, the Federal Election Commission (FEC) does not have enough commissioners to even make quorum.¹² Thus, the main regulatory agency tasked with determining what forms of electoral communication are supported by law is no longer capable of regulating. The Federal Trade Commission (FTC) has levied fines against Facebook and Google, but is so drastically understaffed that it mostly enforces violations of decades-old laws rather than crafting new regulatory regimes for today’s Internet. While there have been congressional hearings into the role of “Big Tech” in spreading disinformation and propaganda, those hearings have mostly been turned into partisan spectacles. The hearings have even become a vector for their own set of conspiracy theories, with a few Republican politicians advancing the baseless claim that Facebook, Google, and Twitter are suppressing conservative content to support a progressive ideological agenda. In the near term, if the marketplace for disinformation is going to be seriously regulated, those regulations will likely be created and enforced by the platforms themselves, rather than by elected officials.

And this in turn leads to my third observation: the greatest threat posed by online misinformation is the lateral effect it has on the behavior of political elites.

ONLINE DISINFORMATION AND THE DISSOLUTION OF LOAD-BEARING NORMS

Online disinformation and propaganda were clearly a bigger problem in the 2016 election than in the 1996 election. But it still bears exploring just what the nature of the disinformation problem is. Why, really, does it matter that online gossip, propaganda, and strategic untruths are spreading faster and farther than ever before? Where is the impact of digital disinformation most keenly felt? I would argue, perhaps counter-intuitively, that the direct

impact of digital disinformation is quite limited, particularly within the context of a presidential election. There is, however, a second-order effect which is quite threatening to the foundations of democratic governance. Political elites are learning just how much they can get away with in the absence of a well-informed public.

The literature on persuasive effects in US general election campaigns is overwhelmingly clear: even for the most sophisticated, large-scale campaigns, it is *tremendously* difficult to change voters' minds. In a recent meta-analysis of field experiments in American elections, published in the *American Political Science Review*, Joshua Kalla and David Broockman conclude "the best estimate of the effects of campaign contact and advertising on Americans' candidate choices in general elections is zero."¹³ In particular, they find that "when a partisan cue and competing frames are present, campaign contact and advertising are unlikely to influence voters' choices." In effect, they are arguing that the sheer volume of campaign communications in US elections, combined with the established partisan preferences of the mass electorate, reduce the marginal effect of campaign persuasive tactics to practically nil. "Voters in general elections appear to bring their vote choice into line with their predispositions close to election day and are difficult to budge from there."¹⁴

Kalla and Broockman's research is not specifically focused on disinformation or on the 2016 presidential election, but the implication is clear: if well-funded, sophisticated voter persuasion efforts launched by seasoned campaign professionals in collaboration with social scientists have little-to-no effect in general elections, we ought to remain skeptical that less well-funded disinformation efforts launched by Russian trolls, Macedonian teens, or the Trump campaign itself would have substantial impacts on voter behavior. Persuasion in a general election is unlike commercial branding or marketing efforts, where consumer awareness is low and consumer preferences are weak. There is no reason to believe the direct impact of microtargeted digital propaganda and misinformation is larger than the direct impact of microtargeted campaign outreach and persuasion campaigns.

At a more foundational level, discussions of media and disinformation are often premised upon the assertion that a well-informed public is a necessary component of a functioning democracy. Misinformation, disinformation, and propaganda are viewed as toxic to a healthy democracy, because they weaken the informational health of the body politic. But there is a contradiction in this premise that we too often ignore. As Michael Schudson documents in *The Good Citizen: A History of*

American Civic Life, American democracy cannot require a well-informed public, because no such public has existed in American history.¹⁵ Though we routinely hearken back to memories of a past golden era in which citizens were better-informed, civically minded, and more engaged in public life, our lived reality has always been messier. The engaged, attentive public is one of the grand myths of American civic life.

The fundamental tension here is that the myth of the attentive public is itself a necessary precondition for a functional democracy. As Vincent Mosco writes in *The Digital Sublime*, myths

... are neither true nor false, but living or dead. A myth is alive if it continues to give meaning to human life, if it continues to represent some important part of the collective mentality of a given age, and if it continues to render socially and intellectually tolerable what would otherwise be experienced as incoherence.¹⁶

American democracy does not require a well-informed public. What it requires are political elites (including media elites) who behave as though an attentive public is watching, rewarding or penalizing them for their actions. In the absence of this myth, there is little preventing political elites from outright graft and corruption.

The great irony of our current moment is that digital misinformation's most dangerous impact comes not through directly deceiving voters and altering their vote choice, but through indirectly exposing to political elites that voters are inattentive and therefore will not keep misbehaving politicians in check. A politician can run on a platform of deficit reduction and then propose legislation that explodes the deficit. A politician can vote for health care legislation that removes the protections for preexisting conditions and then run advertisements claiming the exact opposite. A politician can spend years strategically refusing to ever work with the opposition party on any legislation, specifically so he can blame his opponents for the lack of bipartisan collaboration. If the public is not paying attention, and if traditional media gatekeepers no longer serve as arbiters of political reality, then there is no incentive for engaging in the difficult, messy, and risky work of actual governance. The well-informed public is a myth, but it is a load-bearing myth. Faith in this mythology is a necessary component of a well-functioning democracy.

We are governed both by laws and by norms. The force of law is felt through the legal system – break the law and you risk imprisonment or financial penalties. The force of norms are felt through social

pressure – violate norms and you will be ostracized. The myth of the well-informed public anchors a set of norms about elite behavior: politicians should not lie to the press; they should keep their campaign promises; they should consistently pursue a set of goals that are justifiable in terms of promoting the public good, not merely in terms of increasing their own odds of winning the next election. And while laws change formally through the legislative process, norms change informally and in haphazard fashion. When someone breaks a long-held norm and faces no consequence, when they test out part of the mythology and find that it can be violated without consequence, the myth is imperiled and the norm ceases to operate.

The conspiracy theorists of 1996 were confined to small corners of the Web, just as the conspiracy theorists of 1976 were ostracized from polite society. Things were very different in 2016. During the 2016 presidential race, Donald Trump appeared on conspiracy theorist Alex Jones's radio program and told him "your reputation is amazing." Trump also made Steve Bannon, executive chairman of Breitbart News (a far-right website trafficking in conspiracy theories, misinformation, and disinformation), White House chief strategist.

This is a trend that predates the modern social Web. It can be traced back to at least the 1990s, gaining traction in the aftermath of Newt Gingrich's 1994 "Republican revolution." It coincides with the rise of the World Wide Web, but I would caution against drawing the conclusion that the Internet is what is driving it. Rather, it is a noteworthy accident of history that the rise of the Web immediately follows the fall of the Soviet Union. Governing elites in the United States no longer had to fear how their behavior would be read by a hostile foreign adversary. They almost immediately began testing old norms of good governance and bipartisan cooperation, and found that the violation of these norms did not carry a social penalty. Our politicians have learned that they can tell blatant lies on the Senate floor and in campaign commercials, and neither the media nor the mass public will exact a cost for their actions. In the meantime, online misinformation has provided ongoing additional evidence that the mass public was not paying close attention and that the myth of the well-informed public could be blithely cast away with little immediate consequence.

Social trust in government and the media is eroding. Technology plays a part in all of this. But changing media technology is more of an ensemble cast member than a headlining star in the narrative. The threat we face today is not that the political knowledge of the citizenry has declined due

to online misinformation. The direct effects of misinformation on social media are small, just as the direct effects of all other forms of propaganda have been small. The great danger is that the current digital media environment is exposing the myth of the attentive public, increasing the pace at which political elites learn they can violate the norms of governance with impunity.

CONCLUSION

Writing in 1997, Tom Dowe remarked, “When the barriers come down, when people cease to trust the authorities, they – some of them, anyway – become at once more skeptical and more credulous.” Over the intervening twenty years, the barriers have been in a perpetual state of decline. Trust in all sorts of authority has slipped as well. The credulous skeptics have only gotten more vocal and prominent. The early Web, as Esther Dyson states in Dowe’s article, was “terrible at propaganda, but wonderful at conspiracy.” Today’s Internet excels at both. And though digital propaganda may not directly change many voters’ minds, its second-order effects hasten the erosion of the very foundations of American democracy. What, if anything, can be done to reverse this trend?

The path to repairing our load-bearing democratic myths and constructing a healthier information ecosystem is neither simple nor straightforward. No single political leader, tech company, or journalistic organization can fix these issues on their own. But there is a role to be played by each. Here is what I imagine those roles might look like.

First, there are the platform monopolies – Google, Facebook, and Twitter.¹⁷ In the immediate future, it seems the platforms are going to shoulder an uncomfortable burden. The US government is facing a crisis of competence; the regulatory state is in disarray: the FEC no longer operates. Other government agencies are mired in scandals, run by political appointees whose main qualifications tend to be their personal ties to the Trump organization. Google, Facebook, and Twitter *should not* be determining how we regulate disinformation and propaganda. Such regulatory decisions are beyond what is appropriate to their role and beyond their expertise – the boundaries of acceptable political speech should not be determined by a handful of profit-maximizing firms. But in the near future, there is little hope of genuine regulatory oversight. The platforms will be blamed for the ways in which they are misused in the next election, so they will need to take an active role in

determining and enforcing the boundaries of appropriate behavior. In the long-term, it is an untenable situation, but in the short-term, the platforms stand in as self-regulators-of-last-resort.

Next, there are the political elites. We are going to need our politicians to start believing in the myth of the attentive public again – not because the public is in fact closely watching, but because American democracy only works when our elected officials behave as though they are under close and meaningful scrutiny. Disinformation and propaganda can reduce the public sphere to endless static and noise. It can drown out the very notion of an overriding public interest. But it can only do so if our political elites choose to behave as though it does. If American democracy is to survive, we are going to need public officials who take the public compact seriously. If the regulatory state is going to reclaim its important role, we are going to need to start repairing our regulatory capacity.

Finally, there are the journalistic organizations. As other authors in this volume have noted, the past twenty years have been a time of rapid change within the journalism industry. Much of that change has been more negative than was once predicted. Today's journalism not only has to defend itself against being labeled "fake news" and "the enemy of the people," it also has to compete with partisan propagandists in the struggle for relevance, attention, and revenue. Today's media organizations should hold tight to journalistic principles and editorial judgment. That is what makes them different from the propagandists. The temptation to chase every controversy in service of more eyeballs and more clicks is neither healthy nor productive. Disinformation and propaganda campaigns thrive by creating controversies which then become news stories by virtue of their virality. Media organizations are at their strongest when they prioritize issues of public importance, and when they fulfill their role as watchdogs of political elites. They should focus on this mission not just because it is morally right, but also because it is what distinguishes them from the cheap content farms and partisan propagandists.

Today's misinformation is not identical to the misinformation of the early Web, nor has it proceeded in a linear fashion. Rather, as the Internet has changed and the decades have passed, the quality and character of online misinformation has changed as well. Today's misinformation travels further and faster. It is less traceable and harder for well-meaning individuals to evaluate on their own. Today's misinformation is a strategic asset, at least for campaigns and particular digital media companies. Public mistrust is good for (some) politicians, at least those

who traffic in authoritarian populist appeals. Jettisoning the myth of the well-informed public has worked out very well for some political elites. But it is also worth reminding ourselves that today's Internet is not a finished product. The current version of the social Web does indeed seem to further accelerate public mistrust. This was not always true of the Internet. It is still changing. It is still governable.

The great conundrum we face is that our current political moment routinely and repeatedly reveals that the myth of the well-informed, attentive public can be easily rejected without immediate consequence. Myths are not true or false, but living or dead. Twenty years of online misinformation at an ever-accelerating pace threatens to kill this myth, and there will be consequences. The norms and assumptions governing elite behavior are everywhere tested, and everywhere proven to be easily violated without consequence. We can see, through digital trace data, that misinformation and lies are more clickable than policy details and truths. We can see, through high-profile examples, that political elites can adopt win-at-all-cost strategies and face no social penalty.

Online misinformation is not new. But today's online misinformation is different, and dangerous. We can construct policy frameworks that change the Web and incentivize pro-social behavior and penalize misinformation. But it will be a long and winding path, requiring leadership and commitment from platforms, political elites, and journalistic organizations. Disinformation is a threat to American democracy, not because of how well it works, but because of what it reveals and enables.

NOTES

1. Tom Dowe, "News You Can Abuse," *WIRED*, January 1, 1997, www.wired.com/1997/01/netizen-6/, emphasis added.
2. Samanth Subramanian, "Inside the Macedonian Fake-News Complex," *WIRED*, February 15, 2017, www.wired.com/2017/02/veles-macedonia-fake-news/; Craig Silverman and Lawrence Alexander, "How Teens in the Balkans Are Duping Trump Supporters with False News," *BuzzFeed News*, November 3, 2016, www.buzzfeednews.com/article/craigsilverman/how-macedonia-became-a-global-hub-for-pro-trump-misinfo#.fuzokXaeKo. I focus on Subramanian's later story here because of how it runs parallel to Dowe's, which also appeared in *WIRED* magazine.
3. Andrew Chadwick, *The Hybrid Media System: Politics and Power* (New York: Oxford University Press, 2013).
4. Craig Silverman, J. Lester Feder, Saska Cvetkovska, and Aubrey Belford, "Macedonia's Pro-Trump Fake News Industry Had American Links, and Is Under Investigation for Possible Russia Ties," *BuzzFeed News*, July 18, 2018,

- www.buzzfeednews.com/article/craigsilverman/american-conservatives-fake-news-macedonia-paris-wade-libert.
5. Josh Marshall, "A Serf on Google's Farm," Talking Points Memo (blog), September 1, 2017, <https://talkingpointsmemo.com/edblog/a-serf-on-googles-farm>.
 6. Dave Karpf, *Analytic Activism: Digital Listening and the New Political Strategy* (New York: Oxford University Press, 2016).
 7. Robert W. McChesney and Victor Pickard, *Will the Last Reporter Please Turn out the Lights: The Collapse of Journalism and What Can Be Done to Fix It* (New York: The New Press, 2011).
 8. Joseph Turow, *The Daily You: How the New Advertising Industry Is Defining Your Identity and Your Worth* (New Haven, CT: Yale University Press, 2011).
 9. Maciej Ceglowski, "The Internet With A Human Face" (Keynote lecture, Beyond Tellerrand, Dusseldorf, Germany), May 20, 2014. https://idlewords.com/talks/internet_with_a_human_face.htm.
 10. Adrian Chen, "The Agency," *New York Times Magazine*, June 2, 2015, www.nytimes.com/2015/06/07/magazine/the-agency.html.
 11. Dave Karpf, *Analytic Activism*.
 12. Shane Goldmacher, "The Federal Election Commission Needs 4 of 6 Members to Enforce the Law. It Now Has 3," *New York Times*, August 26, 2019, www.nytimes.com/2019/08/26/us/politics/federal-election-commission.html.
 13. Joshua L. Kalla and David E. Broockman, "The Minimal Persuasive Effects of Campaign Contact in General Elections: Evidence from 49 Field Experiments," *American Political Science Review*, 112, no. 1 (2018): 148–166.
 14. Andrew Gelman and Gary King "Why Are American Presidential Election Campaign Polls so Variable When Votes Are So Predictable?," *British Journal of Political Science*, 23, no. 1 (1993): 409–451.
 15. Michael Schudson, *The Good Citizen: A History of American Civic Life* (New York: Free Press, 1998).
 16. Vincent Mosco, *The Digital Sublime: Myth, Power, and Cyberspace* (Cambridge, MA: MIT Press, 2004).
 17. Amazon is a platform monopoly as well, though it has been less central to the particular episodes of propaganda and disinformation discussed in this book.