

On Digital Disinformation and Democratic Myths

December 10, 2019

What is the sound of a thousand social media bots clapping?

For two years, from 2015 to 2017, the Russian Internet Research Agency (IRA) operated a “Blacktivist” Facebook page. The Blacktivist account posed as an organic, grassroots online activism site, amplifying and contributing to the Black Lives Matter movement. It shared news, memes, and perspectives about racial injustice in America. Blacktivist was the most prominent of the IRA-generated Facebook pages, [with over 6.18 million “interactions” recorded across its top 500 posts](#). The page had 360,000 likes, more than the verified Black Lives Matter account on Facebook (O’Sullivan and Byers [2017](#)). In congressional hearings about Russian online disinformation and propaganda campaigns, poster-sized versions of Blacktivist posts and memes were on prominent display. Blacktivist stands as a warning sign of the sheer volume of online disinformation activities by foreign actors in US politics.

What should we make of these numbers, though? On the surface, it appears the IRA’s Blacktivist account was more popular than the Black Lives Matter Facebook account. But these numbers are inflated to an unknown degree. We do not know what portion of the page’s interactions (shares, likes, and comments) were from unsuspecting American citizens, compared to the portion that was from automated or deceptive accounts, operated by click-farming IRA employees. The easiest route to high Facebook interactions is to build a network of bots sharing, liking, and commenting back and forth to one another. The second-easiest route is simply to echo the same content and frames that are emanating from real-world social movements. What, after all, separates the IRA/Blacktivist-created meme “There is a war going against black kids” from a similar post from a genuine Black Lives Matter activist who is motivated by earnest outrage at the death of 12-year-old Tamir Rice?

The Blacktivist Facebook page is clear evidence that the Russian government sought to amplify and exploit racial strife in US politics. But strategic intent is not strategic impact. And the ease with which researchers can now assemble and visualize data on these influence operations can mask the difficulty in assessing what the numbers actually indicate. Some (likely significant) portion of Blacktivist’s shares, likes, and comments came from the IRA’s own click farmers in St. Petersburg. Those click farmers are densely clustered. They share, like, and comment on one another’s posts—pretend Ohioans

promoting the content of pretend Michiganders, increasing exposure on the Facebook newsfeeds of pretend South Dakotans. But Russian click farmers do not cast ballots. They do not turn out to public hearings. When a separate IRA-backed account used Facebook to promote offline anti-immigration protests, it was hailed as proof of the dangers posed by these foreign disinformation operations. Yet it is also worth noting that [barely anyone showed up to those offline protests](#).

Generating social media interactions is easy; mobilizing activists and persuading voters is hard. Online disinformation and propaganda do not have to be particularly effective at duping voters or directly altering electoral outcomes in order to be fundamentally toxic to a well-functioning democracy, though. The rise of disinformation and propaganda undermines some of the essential governance norms that constrain the behavior of our political elites. It is entirely possible that the current disinformation disorder will render the country ungovernable despite barely convincing any mass of voters to cast ballots that they would not otherwise have cast.

Much of the attention paid by researchers, journalists, and elected officials to online disinformation and propaganda has assumed that these disinformation campaigns are both large in scale and directly effective. This is a bad assumption, and it is an unnecessary assumption. We need not believe digital propaganda can “hack” the minds of a fickle electorate to conclude that digital propaganda is a substantial threat to the stability of American democracy. And in promoting the narrative of IRA’s direct effectiveness, we run the risk of further exacerbating this threat. The danger of online disinformation isn’t how it changes public knowledge; it’s what it does to our democratic norms.

Why direct effects are so rare

How many votes did Cambridge Analytica affect in the 2016 presidential election? How much of a difference did the company actually make?

Cambridge Analytica has become something of a Rorschach test among those who pay attention to digital disinformation and microtargeted propaganda. [Some](#) hail the company as a digital Svengali, harnessing the power of big data to reshape the behavior of the American electorate. [Others](#) suggest the company was peddling digital snake oil, with outlandish marketing claims that bore little resemblance to their mundane product.

One thing is certain: the company has become a household name, practically synonymous with disinformation and digital propaganda in the aftermath of the 2016 election. It has claimed credit for the surprising success of the Brexit referendum and for the Trump digital strategy. Journalists such as [Carole Cadwalladr](#) and [Hannes Grasseger and Mikael Krogerus](#) have published longform articles that dive into the “psychographic” breakthroughs that the

company claims to have made. Cadwalladr also exposed the links between the company and a network of influential conservative donors and political operatives. Whistleblower Chris Wylie, who worked for a time as the company's head of research, further detailed how it obtained a massive trove of Facebook data on tens of millions of American citizens, in violation of Facebook's terms of service. The Cambridge Analytica scandal has been a driving force in the current "techlash," and has been the topic of congressional hearings, documentaries, mass-market books, and scholarly articles.

The reasons for concern are numerous. The company's own marketing materials boasted about radical breakthroughs in psychographic targeting—developing psychological profiles of every US voter so that political campaigns could tailor messages to exploit psychological vulnerabilities. Those marketing claims were paired with disturbing revelations about the company violating Facebook's terms of service to scrape tens of millions of user profiles, which were then compiled into a broader database of US voters. Cambridge Analytica behaved unethically. It either broke a lot of laws or demonstrated that old laws needed updating. When the company shut down, no one seemed to shed a tear.

But what is less clear is just how different Cambridge Analytica's product actually was from the type of microtargeted digital advertisements that every other US electoral campaign uses. Many of the most prominent researchers warning the public about how Cambridge Analytica uses our digital exhaust to "hack our brains" are marketing professors, more accustomed to studying the impact of advertising in commerce than in elections. The political science research community has been far more skeptical. An investigation from *Nature* magazine documented that the evidence of Cambridge Analytica's independent impact on voter behavior is basically nonexistent (Gibney [2018](#)). There is no evidence that psychographic targeting actually works at the scale of the American electorate, and there is also no evidence that Cambridge Analytica in fact deployed psychographic models while working for the Trump campaign. The company clearly broke Facebook's terms of service in acquiring its massive Facebook dataset. But it is not clear that the massive dataset made much of a difference.

At issue in the Cambridge Analytica case are two baseline assumptions about political persuasion in elections. *First*, what should be our point of comparison for digital propaganda in elections? *Second*, how does political persuasion in elections compare to persuasion in commercial arenas and marketing in general?

There is universal agreement that data-driven political campaigns are able to more efficiently target their communications. The political scientists who take issue with the Cambridge Analytica narrative are more narrowly questioning the *marginal additional value* of social media and psychographic data. Eitan Hersh, for instance, studied campaigns' use

of data in his [2015](#) book, *Hacking the Electorate*. He found that campaigns are primarily interested in estimating two variables: (1) likelihood of voting and (2) candidate preference. Those are the variables that campaigns use to build persuasion and turnout models. And for both variables, the lion's share of explanatory power comes from the voter file, which includes an individual's voting record and (in most states) their party registration. In states that included party registration data, Hersh found that adding consumer data and social media data added effectively nothing to the campaign's persuasion and turnout models. In states that lacked party registration data, Hersh found that consumer and social media data were used as a proxy to fill in the gap, but that they were a poor proxy, leaving campaigns partially blind. These characteristics—party registration and vote history—are far more predictive for the variables that political campaigns care about than any personality profile based on a collection of Facebook likes.

At issue here is a stark reality of contemporary American elections. Political persuasion is exceptionally hard. In a recent meta-analysis of field experiments in American elections, Joshua Kalla and David Broockman ([2018](#)) found that “the best estimate of the effects of campaign contact and advertising on Americans' candidate choices in general elections is zero.”

Additionally, they note that “when a partisan cue and competing frames are present, campaign contact and advertising are unlikely to influence voters' choices.” In other words, Republicans and R-leaning independents vote for the Republican presidential nominee; Democrats and D-leaning independents vote for the Democratic presidential nominee. Persuasion campaigns can play a role in primaries (which lack a partisan cue distinguishing between the candidates), but the multibillion-dollar industry of electoral campaign advertisements (TV, radio, and digital) has no statistical impact on the outcome of general elections.

Political persuasion is systematically different from other forms of marketing and propaganda. Imagine if there were only two soft drinks in America—Coke and Pepsi. Further imagine that American citizens could only buy a single soft drink once every four years, and that their brand attachments began forming during their childhood and remained mostly stable over time. Also assume that Coke and Pepsi spent several billion dollars over the course of a year to reinforce Americans' pre-existing attachments to these soft drinks. Under these conditions—high consumer awareness, strong existing preferences, an extremely narrow purchasing window, and high advertising volume—we would presumably expect new advances in advertising to have only the slightest effects on the ultimate outcome. Model and slice the public however you like, in the end the Coke drinkers will purchase their quadrennial Coke and the Pepsi drinkers will purchase their quadrennial Pepsi.

Microtargeted marketing and persuasion have a greater marginal influence under the more relaxed conditions that we find in other forms of marketing and culture—low consumer awareness, weak existing brand attachments, repeat opportunities to purchase the product, and/or low advertising volume. Direct impact on voter behavior, in other words, is a higher bar than you would find in any other type of behavioral-change campaign.

Consider an alternate hypothetical example: gym membership sales. If Cambridge Analytica's fanciful claims about breakthroughs in the science of psychometric targeting were really true, then we ought to see those same techniques producing massive value when applied to the marketing of pricey gym memberships. One can measure the outcome variable on a weekly basis (how many people were tempted into signing up for a new membership plan). Different messages should affect people with different psychological profiles and vulnerabilities. The marketplace is large, but it is not completely saturated with advertising. It has been several years since Cambridge Analytica supposedly deployed these techniques on an unwitting populace in the United States and the United Kingdom. Why have the same techniques not been used for gym membership sales?

Moreover, why would a marketing firm want to use a presidential election to test out and develop these techniques? Presidential elections are an exceptionally hard test case for psychometric advertising, and the ultimate outcome variable that campaigns care about (candidate choice on election day) can only be measured a single time. Even if one believes in the potential of targeting digital advertisements on the basis of psychological profiles, a presidential campaign would be the *worst* place to test and refine this type of advertising.

I have seen over the past few years a chilly divide forming within the research community. On one side are the digital media researchers who see genuine cause for alarm in the rise of digital disinformation and propaganda. They argue that digital media presents new propaganda and disinformation threats. They hold up Cambridge Analytica as proof of how quickly the landscape is changing. They warn that mainstream political scientists have their collective heads in the sand. On the other side are the political science researchers who see a hype bubble forming and want no part of it. They point to a literature on campaign persuasion and mobilization that shows tiny-to-null effects across the board. They argue that Cambridge Analytica's marketing collateral should not be accepted at face value. They warn that this new group of researchers lack clear counterfactuals and appear to be chasing the latest headlines. It is a divide that could fundamentally undermine this emerging field of research.

And it is a divide that we can bypass if we set aside the narrow focus on direct impacts of disinformation and propaganda in elections. Disinformation and propaganda are not dangerous because they effectively trick or misinform otherwise-attentive voters; they are

dangerous because they disabuse political elites of some crucial assumptions about the consequences of violating the public trust.

The myth of the attentive public

Why should we care about disinformation and propaganda if we cannot show direct persuasive impacts? Discussions of media, propaganda, and online disinformation often rest upon the premise that a well-informed public is a necessary condition for a functional democracy. How, after all, are citizens supposed to hold politicians accountable if they lack accurate information with which to evaluate governmental performance?

There is an awkward flaw in this line of reasoning, however: American democracy has *never had* a well-informed public. As Michael Schudson documents in his [1998](#) book, *The Good Citizen: A History of American Civic Life*, we have not declined from some past golden era of well-informed, attentive citizens. The citizens of any particular decade have never lived up to our expressed civic ideals. This same insight is further supported across a range of empirical literatures, including Michael Delli Carpini and Scott Keeter's ([1996](#)) *What Americans Know about Politics and Why It Matters*, John Zaller's ([1992](#)) *Theory and Nature of Mass Opinion*, and Susan Herbst's ([1998](#)) *Reading Public Opinion*. Ever since the rise of public polling in the early twentieth century, it has been abundantly clear that most American citizens do not follow political and civic affairs closely. Schudson extends our historical memory even further, making clear that the years preceding public polling were no better. There is no bygone era of a well-informed, attentive public.

What we have had in lieu of a well-informed citizenry is what might be termed a “load-bearing” myth—the myth of the attentive public. This myth has had adherents, to a greater or lesser extent, among the country's media and political elites. It has influenced their behavior, buttressing norms that prevent some of the worst excesses of unchecked power. I use the term myth here in the same way that Vincent Mosco does in his [2004](#) book, *The Digital Sublime*.

Myths, he writes, “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” (Mosco 2004, 29). Myths, in this sense, need not be rooted in historical fact; they cannot be proven or disproven. Rather, myths are *social facts*—living myths influence the behavior of their believers, structuring how they interpret the world, how they act, and what they expect of one another. Dead myths cease to influence social behavior, having been discarded in favor of an alternate shared social expectation.

The myth of the attentive public has long held favor among American media and political elites. It states that there is a public trust between politicians and their constituents, that those constituents are aware (or might at any time become aware) of politicians that stray from their public promises or violate the public trust, and that those who are found to violate the public trust or break their public promises will incur a cost. The myth of the attentive public is, in many ways, what separates democracy from all other forms of governance. It imbues voting with the importance of an active referendum upon the state of policymaking, where in nondemocratic countries voting serves just as a civic ritual that signifies the ongoing assent of the masses to continue to be governed.

We can see the myth of the attentive public at work when politicians struggle to justify how their votes on legislation square with their campaign promises. We can see it when journalists try to catch politicians contradicting themselves. We can see it in the rise of the fact-checking industry (Graves [2016](#)). The media as the watchdog or “fourth estate” of American government is premised upon the belief that the media plays a vital role in keeping the mass public suitably informed. Political elites take this interaction seriously because they believe in the myth of the attentive public.

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 being sued. The force of norms is felt through social pressure—violate norms and you will be ostracized. The myth of the attentive 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 a mythology and find that it can be violated without consequence, the myth is imperiled and the norm ceases to operate.[\[1\]](#)

For a democracy to remain functional, political elites—elected officials, judges, political appointees, and bureaucrats—must behave as though (a) they are being watched and (b) if they betray the public trust, they will face negative consequences. Otherwise there is little to prevent outright corruption. If political elites behave as though there is no cost to outright lying or procedural hypocrisy, then preventing the slide into corruption is a Sisyphean task.

The indirect effect of rampant online disinformation and propaganda, then, is that it undermines the myth of the attentive public. If the news is all “fake,” then there is no need to answer reporters’ questions honestly. If the public is made up of easily-duped partisans, then there is no need to take difficult votes. If the public simply doesn’t pay attention to

policymaking, then there is no reason to sacrifice short-term partisan gains for the public good.

And while opinions within the research community are starkly divided over the direct effect of disinformation and propaganda, there is essentially unanimous agreement that these indirect trends are toxic and abundant.

Conclusion: The dangerous myth of the digital propaganda wizard

In my years of studying digital politics and online activism, I have noticed two competing stories that researchers, practitioners, and journalists tell ourselves about the role of “big data” in politics. The first is a story of digital wizards—an emerging managerial class of data scientists who are capable of producing near-omniscient insights into public behavior. Cambridge Analytica is just the latest iteration of this tale. (This time, it is Cambridge Analytica that is cast as the nefarious digital propaganda wizards, subtly reshaping the electorate from the shadows. Last time, it was the Obama data wizards, and they were celebrated instead of demonized.)

There is something fundamentally comforting and appealing about the digital wizards. It tells us that there are experts, somewhere, who have everything under control. You can expose those experts. You can hire those experts. You can glean business insights from those experts. Through the right graduate program, perhaps you can even become one yourself. It is a story that fits neatly within a Silicon Valley pitch deck, a TED talk, a graduate school brochure, or a journalistic feature story on the latest digital revolution reshaping our society.

The second story we tell ourselves is more mundane. It features “Build-Measure-Learn” cycles and constant iteration. It is a story in which data scientists and communication gurus lack any grand plan or transcendent vision. What they have instead is the capacity to try things out, measure performance, identify problems, patch them, and repeat. It is a tale full of messy workflows, incomplete datasets, and endless trial and error. There are no wizards in this second story, no omniscient geniuses enacting perfectly designed plans. There are just people—some talented, others less so—figuring a few more things out along the way.

The first story is engaging and appealing. But it has little basis in reality. Simply put, we live in a world without wizards. It is comforting to believe that we arrived at this unlikely presidency because Donald Trump hired the right shadowy cabal. That would mean Democrats (or other Republicans) could counter his advances in digital propaganda with advances of their own, or that we could regulate our way out of this psychometric arms race. It is a story with clear villains, clear plans, and precise strategies that might very well

be foiled next time around. It is a story that keeps being told, because it is so easy to tell.

But we pay a price for the telling and retelling of this story. The problem is that the myth of the digital propaganda wizard is fundamentally at odds with the myth of the attentive public. If the public is so easily duped, then our political elites need not be concerned with satisfying their public obligations. If real power lies with the propagandists, then the traditional institutional checks on corruption can be ignored without consequence.

It is easy for researchers to contribute to the myth of the propaganda wizards. Cambridge Analytica was made famous by well-meaning people trying to raise an alarm about the company's role in reactionary political networks. But I would urge my peers studying digital disinformation and propaganda to resist contributing to hype bubbles such as this one. It is not just that we risk creating an unnecessary and unproductive internal rift within the interdisciplinary research community. It is also that, in pursuit of high-impact research, we might be further eroding the very norms that we need to preserve.

The first-order effects of digital disinformation and propaganda, at least in the context of elections, are debatable at best. But disinformation does not have to sway many votes to be toxic to democracy. The second-order effects undermine the democratic myths and governing norms that stand as a bulwark against elite corruption and abuse of power. In amplifying the myth of the digital propaganda wizard, we run the risk of undermining the load-bearing norms that desperately need to be reinforced.

[1] 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 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 none of those norms carried 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. The Trump administration has radically accelerated this phenomenon. Online disinformation and propaganda play an indirect, amplifying role, providing a steady stream of social proof that the myth of the engaged public can be cast away with no immediate personal consequences.

Expert Reflections are submitted by members of the MediaWell Advisory Board, a diverse

group of prominent researchers and scholars who guide the project's approach. In these essays, they discuss recent political developments, offer their predictions for the near future, and suggest concrete policy recommendations informed by their own research. Their opinions are their own, with minor edits by MediaWell staff for style and clarity. You can find other Expert Reflections [here](#).